

What Cannot Be

Signals Faster Than Light, Time Travel, Simulation and Determinism, in Plain Words

Forrest Landry

UVSM PUBLICATIONS
Wytheville, Virginia

© 2026 Forrest Landry. Some rights reserved.

Published by UVSM Publications LLC
965 Reed Creek Mill Rd
Wytheville, VA 24382-8287
USA

uvsmpublications.com
contact@uvsmpublications.com

ISBN 979-8-90710-415-0

DOI: https://doi.org/10.67291/KHIY2463

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (CC BY-NC-ND 4.0).

You may share this work for non-commercial purposes with proper attribution, but you may not adapt, transform, or build upon it.

Full license text: https://creativecommons.org/licenses/by-nc-nd/4.0/

Author ORCID: 0009-0005-6275-0362

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), except as permitted under the license above, without prior written permission of the publisher.

Contents

Preface: Before We Begin v

Part One: The Ground 1

1. How Does Anyone Find Anything Out? 2
2. Two Kinds of Law 26
3. Why One Description Cannot Be Perfect Both Ways 47
4. Process, Pattern and Possibility 70

Part Two: Causes Without Clockwork 93

5. Is Everything Already Fixed? 94
6. What Is a Cause, Then? 112
7. Choice, Change and Cause 141
8. A Click No One Could Predict 164
9. The Free Will Theorem 187

Part Three: Bell and the Escape That Is Not There 207

10. What Bell Found 208
11. The Escape That Is Not There 230

Part Four: Signals and Time 253

12. The Shape of What Can Reach What 254
13. Why Time Has a Direction 273
14. Three Kinds of Time 297
15. Time Machines on Paper 318
16. Why Time Travel Names Nothing 340

Part Five: The Edges of Physics 365

17. Two Theories That Cannot Be One 366
18. The Three Horizons 399
19. The Two Constants That Bound Every Domain 432

20. What Physics Cannot Reach 470

Part Six: The World Is Not a Program 505

21. Are We Living in a Simulation? 506

22. Why No Computation Is Everything 533

Afterword: What Cannot Be, and What Can 563

Appendix: For the Physicist 574

Terms 608

Works Cited 623

Index 641

Before We Begin

Four Things We Are Told Might Be Possible

What is this book, and what will it ask of you? Start with four things that people are often told might one day be possible. The first of the four things is a message sent faster than light, so that a friend near a distant star could read the message before the light of your lamp had gone a fraction of the way. The second of the four things is a trip back to last Thursday, to stand again in a moment that has already gone by. The third of the four things is a world that is someone's program, running on a computer in some other world, with us inside the program. The fourth is a universe that runs like clockwork, in which everything that will ever happen was fixed from the start.

Each of the four is usually discussed as a matter of difficulty: too much energy, too little skill, a machine not yet built, a theory not yet found. Each of the four asks for something that cannot be, and I say exactly what "cannot be" means. It is about the limits that the structure of knowing sets on the laws of physics. The same result excludes four more things that are often called merely difficult: a hidden fixing of the experimenters' own choices, a faithful running back of time, one theory that perfects both of physics' great theories at once, and a program that is a mind.

What This Book Offers

What this book offers is a reading of the physics that already exists: an account of why the deepest results of the last century of physics have the shape they have, and of why certain long-pursued projects of physics can-

not be completed in the form in which they were set. The equations, the measurements and the values of the constants that the book reads are physics' own. In this book I add no new physics: I propose no new equation, predict no new measured value, and derive the value of no constant. What I predict, I predict by kind: what sort of relation physics will find, and never how large.

The reading rests on one result about comparison, the Incommensuration Theorem (incommensurate means having no common measure). Every measurement is a comparison: the act of finding whether two things are the same or different, in some respect. What a comparison finds is the content of the comparison; the setting that makes the two things two, such as two places or two moments, is the context of the comparison. Sameness and difference, applied to content and to context, give four characters that a description can have. Symmetry is the same content in different settings: the same law here and there, today and tomorrow. Continuity is the same content within one setting: change without a break, with nothing jumping. Asymmetry is a content that changes as the setting changes, and discontinuity is a content that jumps within one setting.

The theorem says that, of all the ways of pairing these four characters, only two pairings hold together when both members of the pair are applied absolutely, that is, held exactly, at the foundation of a description, with no exception at any scale. The two are continuity with asymmetry, smooth and going somewhere, and symmetry with discontinuity, the same everywhere and countable. I call those two the *survivors* of the theorem. Perfect symmetry and perfect continuity cannot be had together, and neither can perfect asymmetry and perfect discontinuity. Chapter 3 pictures the theorem with a floor of identical square tiles beside a ramp poured in one piece: the tiles repeat perfectly and always have seams, the ramp is seamless and never repeats itself exactly, and no floor can be perfectly repeating and perfectly seamless at once. The Incommensuration Theorem, with the account of comparison on which it stands, is the whole of the ground. The chapters set out what follows from the Incommensuration Theorem.

The Incommensuration Theorem is a structural fact about the nature of knowing. Every act of knowing is an act of comparison, or is built from acts of comparison, so the theorem, together with the intrinsics of comparison from which it is built, is inherent in every single act of knowing, and it is the deepest and most powerful epistemic basis possible. Epistemic means concerning knowing: how anything comes to be known, checked and held. The intrinsics of comparison are its six parts (sameness and difference, content and context, and the side from which and the side about which a comparison is made), which Chapter 1 counts one by one. Part One builds the theorem from the beginning, one step at a time.

Two words that are usually treated as one are kept apart in every chapter. Whatever makes a comparison, whether a person, an instrument or an atom, is a *comparer*. A cause is a consistency that a finite comparer, acting within time, finds between two things in the world, one following the other. Determinism is the claim that every content is completely fixed for every context, needing no one to find it and no time to pass. The two differ in what they ask of an observer, in how they stand to time, and in the size of things at which they apply. The theorem makes determinism impossible, and it leaves causes in place.

What "Cannot Be" Means Here

Ask someone for the color of the number seven. The answer is not hidden, and it is not too hard to find. The request names nothing: it puts the word "color" where the meaning of that word does not reach. Philosophers call this a category error: a concept applied where it does not belong, so that the request is malformed rather than hard.

When this book says that something cannot be, it means that the request names nothing. It never means that the thing is too hard, too costly, or waiting on a better machine. A signal faster than light asks for a comparison beyond the edge of comparison. A trip back to last Thursday, in the sense people intend, asks one kind of time to be another kind as well.

A clockwork universe asks one description to be perfect in both of the two ways that the theorem says no description can be perfect at once. And a world run as a program inside another world, with real people in it, turns out to be a region of the one world that runs it. Each of these is shown in full in its own chapter. No future machine can satisfy a request that names nothing, so what is excluded in principle is excluded, for that reason, in all future practice.

Three Kinds of Claim

I make three kinds of claim, and I keep the three kinds apart for you, the reader: a claim of one kind, heard as a claim of another kind, would sound stronger or weaker than it is, and each claim is stated at its true strength only while the three kinds are kept apart. Each chapter says which kinds of claim it makes.

The first kind are claims of the structure of comparison itself: the six parts that every comparison has, the four characters, the Incommensuration Theorem, the three axioms that follow from it, the two kinds of law, and the six measures that every physical measurement uses. The claims of the first kind are the ground. They are argued from the structure of comparison and from nothing else.

The second kind are readings of known physics. A result that physics has already established, by its own methods and on its own evidence, is read as an instance of the structure of comparison, and the claim made is that the structure is the reason the result holds. Emmy Noether's theorem of 1918 (which ties each law that stays the same across settings to an amount that is conserved), Bell's theorem (on the linked outcomes of particles measured far apart), the uncertainty relations (the limits on how sharply two paired quantities, such as position and motion, can both be found), the fact that the very small shows itself both as wave and as particle, the fact that the very small comes in countable steps, the failure of determinism as a way of knowing, the incompatibility of general re-

lativity (the theory of gravity as the shape of space and time) with quantum theory (the theory of matter and light, whose findings come in countable steps), and the impossibility of running time back faithfully: each of these holds for a structural reason that the theorem identifies, and naming that reason for each result is the book's claim. In this book I do not establish these results of physics again. I say what each of these results is an instance of. Reading the physics that exists asks the most of a framework, because those results were not made with it in mind, and it must read them without forcing.

The third kind are structural exclusions. No world is perfectly deterministic. Superdeterminism, the claim that the experimenters' own choices were fixed together with what they measure, is a category error. No signal crosses the edge that the speed of light marks. Time has an arrow that comes before the arrow of heat and disorder. And no single theory can find both survivors of the theorem, the two pairings that hold together, perfected at once, the survivor that general relativity finds and the survivor that quantum theory finds. These claims go a step beyond reading. They are consequences of the structure of comparison that reach into physics, into the world in which findings are made, and they say what physics will not find there. They say it by structure, and not by number.

No number in this book is derived in it. Where a number appears, it is quoted from physics, and I say so; the number is said in words, and it is given one of four standings: set by the structure of comparison; set by the structure of comparison once a stated step is taken; measured; or measured, with the reason for its value still a question.

So that you can always tell what a sentence rests on, other statements in the book also say which kind of statement they are. A result from the literature of physics or mathematics is given with the names of the people who found it and the year, and it rests on their work. A conjecture, a statement the literature poses and has not proved, is reported as the literature states it, and it is never used as a step of an argument. A definition fixes

the meaning of a term where the term first appears. Where I identify something in physics with a role in comparison, I say so in the first person, once, where the identification is made.

Knowing the kind of each claim keeps you clear of two misreadings. When an argument from structure lands on a familiar result, it is easy to credit the argument with more than it did, and equally easy to credit it with less. A framework of this kind is most often misunderstood in just those two ways, and I state the kind of each claim to prevent both. I claim what I claim, no more and no less.

How the Book Is Set Out

The book has six parts, and each builds on the ones before it.

Part One, "The Ground", starts from nothing. It shows that every measurement is a comparison with six parts, that every law is one of two kinds, that one description cannot be perfect both ways at once, and that process comes before structure and possibility. Its claims are claims of the structure of comparison.

Part Two, "Causes Without Clockwork", excludes determinism at its root, says what a cause is instead, shows where the new comes from, reads hard randomness, an outcome that nothing before it fixes, as physics finds it in the decay of a single atom, and reads the free will theorem of John Conway and Simon Kochen (2006).

Part Three, "Bell and the Escape That Is Not There", reads what John Bell showed in 1964 and what the experiments since have confirmed, and it shows that superdeterminism, the last proposed escape from what Bell showed, is no escape.

Part Four, "Signals and Time", shows that the light cone, the set of every event that a cause can connect to a given event, is the shape of what can reach what, and that no signal crosses its edge; that time has a direction in every single comparison, so that the world is not symmetric in time, and no reversal of time can be faithful; that physics has three kinds

of time; what the time machines that physicists have found in the equations of general relativity are; and that "time travel", in the sense in which the phrase is used, names nothing.

Part Five, "The Edges of Physics", shows why the two great theories are joined side by side and never made one, and why there are exactly three kinds of acceleration; it follows comparison to its edges, the three horizons of what can be known from here (a horizon is an edge past which nothing reaches us), and to the two constants that mark the two edges of physics: the speed of light, at the edge of the very large, and Planck's constant, a fixed amount of what physicists call action (a quantity made by multiplying an energy by a duration), at the edge of the very small. It presents the correspondences between the two constants and the structure of comparison, it takes the value of each constant from physics, and it argues for no value of either; and it says what physics cannot reach at all.

Part Six, "The World Is Not a Program", asks whether a world could be computed. It gives the argument from the continuum (a whole with no gaps anywhere in it) against simulation, and the argument against the view that everything, minds included, is computation. The world that is someone's program is the third of the four requests, and it is answered last, because its argument uses what the earlier Parts build: the continuum, the bound triple of hard randomness, the arrow of time and subjectivity, and the edges of comparison.

An Afterword gathers what the book has shown. An appendix, "For the Physicist", carries the technical forms of a few results; the main text gives every one of those results in words. A list of Terms, the Works Cited and an Index close the book.

What the Book Covers

Where speed is concerned, the subject of this book is signals: something that passes from one event to another and lets a record at the second depend on an act at the first. I take up signals faster than light. Travel faster than light is a topic for another day.

Where time is concerned, I argue one thing in this book about time travel: that the phrase "time travel", in the sense in which the phrase is used, does not name a concept. The phrase asks that an earlier moment be at once a place in the geometry of space and time, to which a traveler goes, and the very happening it was, given again with the traveler present. An earlier moment taken as a place in the geometry and the same moment taken as the very happening belong to two different times. The two times are geometric time, laid out in the shape of space and time, and event time, the order of happenings as a comparer finds the happenings. By the Incommensuration Theorem, no one time can be found as both geometric time and event time. That argument stands on the theorem and on one principle from the physics of time machines: that a world with loops in time (paths through space and time that come back to their own past) must be consistent with itself. It uses that principle only to examine one reading of the phrase. Physicists have also posed a conjecture, called strong cosmic censorship, about how far the equations of gravity predict, and Chapters 14 and 15 report it. The argument uses no result of strong cosmic censorship, so it holds however that conjecture is settled. The physics of time machines is reported because it is what physics itself says about the machines the phrase imagines, and it checks the conclusion.

A Promise to the Reader

Every idea in this book is built from the beginning, in plain words. Every term is explained where you first meet it. Every number is said in words and given its standing. Nothing depends on seeing a figure or reading an equation, so the book can be heard as well as read.

Each chapter opens with the question it answers, gives a picture before it names an idea, explains each term where it first appears, and closes with a Summary and the question that the next chapter takes up. The chapters are a walk, and Part One fixes the terms that every later chapter uses. A reader or a listener who comes to one chapter on its own will find each earlier term recalled in plain words where that chapter first uses it. What the book asks of you, then, is to walk with it: to take each step from the steps before it, in order.

This book is a reading of physics, and it is the companion of the books of derivation, which take the steps from comparison to the measures and the forms of physics. In this book I do not repeat the steps of derivation. I draw on three foundation volumes. *Comparison; The Necessary Basis* sets out what comparison is, its six parts, and the Root Tautology, that interaction, comparison and relation are one act. *Incommensuration; The Theorem and the Three Axioms* proves the theorem from comparison alone and derives the three axioms from it. Measurement; The Steps from Comparison to Physics takes the ground set out in the first two volumes (comparison, the theorem and the three axioms) step by step to the measures of physics. I state in these pages every result that I take from the earlier volumes, and a reader who has not read the earlier volumes finds in these pages everything needed from those volumes.

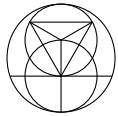
Where We Begin

If what cannot be is settled by what comparison is, then the place to begin is the plainest of all acts of finding out. When a physicist says "we measured it", what actually happened?

Forrest Landry
Wytheville, Virginia

PART ONE

THE GROUND



Chapter 1

How Does Anyone Find Anything Out?

What Happens When We Measure

When a physicist says "we measured it", what actually happened? The question sounds too plain to need an answer. Its answer is the ground of everything this book says about what cannot be.

Start with the plainest measurement there is. Lay a ruler beside a table. The edge of the table is the same as one mark on the ruler, and different from each of the other marks, and the reading is the mark at which the sameness is found. A measurement of length is exactly that: a quantity set against a standard, and a finding of where the two are the same.

Every measurement has the ruler's shape. A *measurement* is a comparison made from inside the world: a count, a reading or a frequency set against a standard. A *comparison* is the act of finding whether two things are the same or different, in some respect. What one act of comparison yields is a *finding*: one outcome, laid down as a record. The mark on the ruler where the table's edge falls is a finding.

Whatever makes a comparison is a *comparer*. A person can be a comparer, and so can an instrument, or an atom that absorbs light. The comparer is a role, the role of relating one thing to another, and a mind, an instrument or a thermostat can fill it. The word comparer names that role, and in this book it never means a mind.

In this chapter I follow the ruler's small act a long way. I show that measuring, comparing and interacting are one act; that every comparison has six parts, which can be counted by trying to take one away; that every measurement brings into being three divisions that were not there before it; and that physics, as the study of measurement, has edges, which are set by what comparison is, before any experiment is done.

Two Particles Meeting

Picture two particles flying toward each other, meeting, and flying apart. Physicists call such a meeting a scattering. Where each particle goes next depends on what each one carried in: its momentum (the amount of motion it carries), its electric charge, and its spin (a built-in turning, which makes a particle such as an electron act as a tiny magnet). In the meeting, those properties are weighed against one another, and the result is written in the paths the two particles take afterward. The outcome of the scattering is a record of how the two particles were the same and how they differed. An interaction, looked at closely, is a comparison.

Now look at the same act starting from comparison instead of from interaction. To compare two things is to bring them into contact. Something must pass between them: some trace of the one must reach the other, some process must connect them. There is no comparison at a distance with nothing crossing the gap. A comparison, looked at closely, is an interaction.

Now look in the same way at relation. To stand in a relation to something is to be placed against it, to be compared with it, which is to have interacted with it. A relation that involves no comparison relates nothing to nothing. Relation is both comparison and interaction.

I call this identification the *Root Tautology*: interaction, comparison and relation are one act, named three ways. A tautology is a statement true by what its words mean, and I mean the word precisely. The Root Tautology is meant literally, and it is true by the meanings of its words, so it is

neither a metaphor nor a guess waiting on evidence. It recognizes that three vocabularies have been naming one act all along: the vocabulary of physics, which says interaction; the vocabulary of the study of knowing (philosophers call that study epistemology), which says comparison; and the vocabulary of mathematics and logic, which says relation. Everything in this book rests on this one identification. It is simple enough to be passed over, and it changes how one reads the whole of physics.

In the Root Tautology, the word relation means the relating itself: one thing actually placed against another, in an act. The word has a second use, which this book keeps apart from the first: a relation written down as structure, such as a law laid out on a page, which stands the same for everyone who reads it. Chapter 4 calls the first a relation in the full sense, which is real, and the second a relation in mathematics, which exists, and shows how the two stand to each other.

The three names are distinct, they cannot be separated, and none can stand in for another. Comparison names the act as a finding; interaction names it as an event, something that happens; relation names it as a standing connection. The three match each other part for part, and they are not one word said three times (a mathematician would say they are isomorphic: alike in form, so that each can be read as the other, without being identical). To call an act a comparison describes it more specifically than to call it an interaction, and to call it an interaction describes it more specifically than to call it a relation. The nesting is a nesting of descriptions, and not of acts: every act that is one of the three is also the other two.

If you doubt the Root Tautology, the test is direct: try to produce an interaction that is not a comparison, a comparison that is not an interaction, or a relation that is neither. The support for the Root Tautology comes from the act of denying it. Any attempt to deny it compares the tautology with something else and finds the two different, so the attempt is already a comparison, with an interaction and a relation inside it. Whoever denies the Root Tautology uses it in the denial. The Root Tautology also carries

the rest of the ground: the six parts of comparison that come next in this chapter, the four ideas that Chapter 2 builds from them, the theorem of Chapter 3 and the three axioms of Chapter 4 all follow from it.

The consequence of the Root Tautology gives this book its subject. Interaction and comparison are one act, so the structure that every comparison has by its nature is the structure that every physical interaction has by its nature. Whatever is necessarily true of comparison as such is necessarily true of physics as such. The Root Tautology leaves physics as it is, and it shows that what physics studies, interaction, is the act that comparison is. Physics is the study of interaction, and interaction has a shape that can be examined before any particular interaction is measured.

Six Parts of Every Comparison

Put an apple on a kitchen scale. The needle settles at a mark, or between two marks: the same as one mark, and different from the others. It reads a weight, in this kitchen, now. And the scale reads from where it stands, about the apple that sits on it. That one small act has six parts, and every comparison has the same six.

The six come in three pairs. The first pair is sameness and difference. *Sameness* is what the compared things share; *difference* is how they diverge. The needle's match with one mark, and its mismatch with the rest, are the verdict of the comparison.

The second pair is content and context. The *content* is what is being compared, what the comparison is about: on the kitchen scale, the weight. The *context* is the frame within which the comparison is made; more exactly, the context is the respect in which the two things compared are two, such as two places or two moments. The context is the setting that makes one finding distinct from another: this kitchen and not another, this moment and not the next.

The third pair is the subjective and the objective. *The subjective* is the stance from which the comparison is made: for the kitchen scale, the scale itself, with its place, its setting and the way it stands. *The objective* is the side about which the comparison is made: the properties of the things compared, which do not depend on who compares the things; on the kitchen scale, the apple. Neither word means a feeling or an opinion. The subjective is the side from which, and the objective is the side about which.

The objective side of one act is a different thing from objectivity. *Objectivity* is the agreement of many readings by many comparers, and it comes in degrees, across many acts: what is objective is what any comparer would find. Scientists call that agreement across many acts reproducibility.

These six are the *six intrinsics* of comparison: sameness and difference, content and context, the subjective and the objective. An intrinsic is a part that belongs to an act by its nature. None of the six can be defined without at least an unspoken reference to the others.

That there are six, and not five or seven, is found and not chosen. It is found by trying to remove one, and watching what happens to the act. Take away the content, and the comparison compares nothing. Take away the context, and the comparison has no frame, and so no way to say what counts as same or different. Take away sameness, and it finds no pattern; take away difference, and it finds no signal. Take away the subjective side, and the comparison has no standpoint from which it is made; take away the objective side, and it has nothing it is about. Remove any one of the six, and the act of comparing falls apart.

Try instead to add a seventh. On examination, the seventh turns out to be built from the six: a sameness held across some context, say, or a content held from some standpoint. Six are necessary, and six are enough.

The pairs hang together in a strong sense. Sameness has a meaning only against the possibility of difference; content has a meaning only within some context; the subjective is always the subjective of some objective.

Where one member of a pair appears, the other is present, at least implicitly. I ask the reader to hold on to the fact that the two members of each pair cannot be separated, since neither has a meaning without the other, because the theorem of Chapter 3 rests on that fact.

One more feature of the six matters for everything that follows. The six intrinsics belong to the act of comparing, and not to the things compared. They belong to the interaction, and so, by the Root Tautology, they belong to physics. When this book says that a result of physics is an *instance* of the structure of comparison, it means exactly this: the result is a place where the six intrinsics, or ideas built from them, show up in the world because they could not fail to show up.

Measuring Is Interacting

Physics is the study of measurement, and measurement is interaction under a particular description. In my book *An Immanent Metaphysics* (first published in 2001) I put the identification in one line: the concept of a measurement is isomorphic with the concept of an observation, a perception and an interaction. I mean that line literally, as a claim at the foundation. It means that every measurement is also an interaction, and that every property of measurement as such is a property of interaction as such. Anything true of one of the four, as a way of knowing, is true of the others, in every case. The ways by which we come to know the world cannot be separated from the ways by which the world acts upon itself.

Measurement is a comparison; so is perception, and so is every interaction. Measurement is comparison made by an instrument against a standard; perception is comparison considered from the side of the self. Calling an act a measurement, or a perception, describes it more specifically, and it remains the same act. Because a measurement is a comparison, it inherits the structure of comparison: the six intrinsics are necessary and suffi-

cient to characterize any comparison, no comparison lacks any of them, and none can be removed without destroying the act of comparing itself. Every measurement, however simple or complex, involves all six.

I identify measurement with interaction because measurement and interaction are both comparisons. A comparison always compares interactions, and never purely static things. Whatever presents itself to be compared does so by acting, by being acted upon, by making a difference somewhere. What we call a thing is a steady pattern in a field of interaction: a content held the same across some range of contexts. Its steadiness is what makes it available to be compared at all.

Three consequences come with the identification, and each shows in physics.

The first consequence is that every observation compares at least two things: the reading and the scale, the sample and the standard, the signal and the background. One cannot observe just one thing. A measurement of one thing alone, with nothing for it to differ from, is the absence of a measurement. That is why every measurement in physics is a ratio, a count against a unit, a difference from a reference.

The second consequence is that an observer always observes at least one something. A detector that registers no event has still registered: the absence of a click is an outcome that the apparatus was set up to be able to produce. A silent detector has still given an answer. Every observation has an outcome, and there is no empty observation. To be an observer at all is to be already in interaction, and there is no position outside interaction from which the interacting could be surveyed.

The third consequence is that every measurement leaves part of what it measures unmeasured. Every interaction passes on some content and leaves other content behind, and the content left behind is more than content not yet measured: it is content that cannot be measured together with what was measured. One cannot completely observe any one thing. A photograph shows what it is for one thing to be observed incompletely. A very short exposure fixes exactly where a runner is, sharp and still, and in

fixing that it gives up the blur that would have shown how fast she was going; a long exposure shows the motion, and gives up the sharp place. That incompleteness belongs to what measuring is, and it is no fault of any instrument. Chapter 18 shows that incompleteness in physics, where physicists call it uncertainty.

In a measurement, one side of the interaction is taken as the measurer, and the other as the measured, and what passes from the measured to the measurer is taken to stand for the measured. The one who measures supplies the labels 'measurer' and 'measured', and the interaction itself carries no such labels. So every limit on interaction is a limit on measurement. And every limit on measurement holds of what measurement is, and not of any one instrument, so every limit on measurement is a limit on interaction.

Does a detector measure when no one is watching? Yes. The comparer is a role, and an instrument fills it. A detector registers a particle, and the interaction in the detector is a comparison whether or not anyone reads the count later.

What a Measurement Brings Into Being

Click a stopwatch as a runner crosses the line. The click makes a record: a before and an after that the record fixes, and the record will never be unclicked. Before the click, the race had no measured time.

Every measurement brings three divisions into being. The first of the three divisions is the division of past from future. A measurement yields a record, and a record is of what has happened; the record is used for a prediction, and a prediction is of what has not yet happened. A reading that could not be told apart from the state the instrument will be in a moment later would be no reading. The measuring interaction has a before and an after, and the after is not the before. The act of measuring makes a division in time that did not exist as a division until the measurement was

made. The one-way running of time is a fact of process, of what is happening here and now, and the fact that a measurement cannot be undone is the mark in time that a comparison has occurred.

The second division that every measurement brings into being is the division of here from there. Every measurement takes place somewhere: the apparatus is here, and the sample is there (or, in the limiting case, here as well); the measurer is not the measured. Even where the two touch, the touching is a boundary, and a boundary has a here on one side and a there on the other. The result of the measurement applies to one bounded region. Where nothing is told apart as elsewhere, nothing has been measured.

The third division that every measurement brings into being is the division of the actual from the possible. A measurement has an outcome, and an outcome is one among alternatives that did not happen. Before the measurement, the state of the system is a range of possibilities; after it, one possibility is actual. A reading that could have been nothing else carries no information: it was a foregone conclusion, and not an outcome. The actual is actual against a range of what was possible, and the measurement is what selects within that range. To measure is to turn the possible into the actual.

These three divisions belong to measurement itself: take them out, and measurement goes with them. They are the three contexts of measurement: time, space and possibility. Time is the context in which past differs from future; space is the context in which here differs from there; possibility is the context in which the actual differs from the possible. I call the three together the *reification triple* (to reify is to make actual), because every measurement makes all three actual.

The three contexts line up with three ways of looking at anything, which Chapter 4 names and sets out: as process, what is happening here and now; as structure, what is laid out the same for every comparer; and as the source of the new, what makes things possible. Time, and the force that acts through it, belong with process. Space, and the pattern that fills

it, belong with structure. Possibility, and the probability (the likelihood of each outcome) weighed within it, belong with the source of the new. So the three divisions are one division for each way of looking, and no measurement can be considered without all three being considered, at least implicitly.

The comparer, in turn, is placed by the three divisions. It is placed in time, bound to one moment, the present; it is placed in space, bound to one place, here; and it is placed in possibility, bound to one actuality, the outcome that occurred.

Each of the three divisions has a subjective side and an objective side. The objective side shows what has happened, what came to be, what is laid out for any comparer to read. The subjective side shows the stance from which the objective side is read: placed in one present, at one here, facing a range of what could be found, what has not yet happened, what might come to be. The object belongs with structure: what exists, what is patterned, what is settled and can be shared. The subject belongs with the source of the new: possibility, choice, what is not yet settled. At the scale of a human life, the way the future, the possible, and what is other than here hang together is the felt meaning behind choice.

The stopwatch shows both sides. The objective side is the time on its face, which anyone at the track can read and which will read the same tomorrow. The subjective side is the timekeeper at the line: in one present, at one place beside the track, facing a next race whose winner is not yet settled.

The structure of measurement, which is the foundation of physics, always involves both a measured and a measurer. The claim that measurement always involves both a measured and a measurer is a statement about the structure of measurement. The measurer can be a detector as well as a person, so the claim does not make physics into psychology. If an account removes the measurer, then that account is left incomplete, in ways that matter in practice and not only in philosophy. If a physics tries to write down only the objective side, then the mathematics of that phys-

ics still keeps the subjective side, unstated: in the choice of variables, in the choice of reference directions (physicists call that set the basis), and in the choice of what is to count as a system.

What Physics Measures

A ship's log records, at each entry, three things about the ship: its load, the push of the wind, and the chance of a storm. It sets each of them against three things about the voyage: where the ship was, when, and which courses lay open. Those six columns are a picture of a second set of six, which belongs to physical measurement in particular.

The six intrinsics are what every comparison has. The *six measures* are what a physical measurement measures: three contents, each measured in its own context. Because measurement is interaction, the six intrinsics of comparison have a physical face, and in *An Immanent Metaphysics* I name that face as six measures: mass, space, force, time, probability and possibility. The six measures are what the six intrinsics look like when the comparison is a physical one.

Set the six measures out as two rows of three. The upper row holds the contents: pattern (which takes in mass, inertia, form, shape and structure; inertia is a body's resistance to being set moving), force, and probability. The lower row holds the contexts: space, time and possibility. Each content is measured in the context beneath it: pattern in space, force in time, probability in possibility. The three columns are the three ways of looking named above: pattern and space belong with structure, force and time with process, probability and possibility with the source of the new. The count, six, is set by the structure: three roles, each taken twice, once as content and once as context.

Each of the two rows of the six measures has its own kind of law. In *An Immanent Metaphysics* (aphorism 1.85-10) I put the rule this way: "Where pattern, force, and probability are the content of a perception, these will be defined by symmetry law. Where space, time, and possibility are the

context of a perception, these will be defined by continuity law". In plain words: the three contents obey laws of sameness across settings, and the three contexts obey laws of unbroken change, change without a jump. Chapter 2 sets out those two kinds of law, which are called symmetry and continuity.

Each of the six measures therefore has two addresses, as a day on a wall calendar has a row, its week, and a column, its day of the week. Its row says which kind of law governs it: the three contents fall under symmetry law, and the three contexts fall under continuity law. Its column says which way of looking it belongs to: structure, process, or the source of the new. The two addresses answer two different questions, so a measure can have both at once with no conflict between them. Space sits in the row of contexts and in the column of structure: as a context it falls under continuity law, and as the setting in which pattern is laid out it belongs with structure. Probability sits in the row of contents and in the column of the source of the new: as a content it falls under symmetry law, and as the weighing of what could happen it belongs with the source of the new.

The two rows also say exactly what determinism would have to fix. *Determinism*, which Chapter 5 examines, is the claim that everything is fully and exactly fixed: every value of every content, for every context, with no observer and no time needed. In the terms of the two rows, determinism would fix every entry of the upper row for every entry of the lower row.

A True Zero

Turn a lamp down, and the light lessens until it is gone. Darkness is the absence of light, and not its opposite: nothing lies beyond the absence. There is no negative light.

Many of the most basic quantities that measurement tracks have this shape. They have a fixed absence at one end, an unbounded positive range at the other, and no negative counterpart. I call such quantities *zero-positive*. Life is present in positive degrees of organization; death is the absence

of those processes, and not a process in its own right. Joy and suffering are each present in some positive degree, or absent, and the absence of one is not the presence of the other.

The zero-positive shape is a feature of what is measured. The zero-positive shape is no feature of language and no convention of measurement. Certain quantities are one-sided in their very make-up, and to give them negative values, or a scale that runs evenly both ways, distorts the structure of what they describe. Measuring such a quantity faithfully needs three things: an *origin*, the zero point, which marks the absence; a *unit*, the standard of comparison, which says what counts as "one more"; and an *extent*, the positive range, which may be unbounded. Origin, unit and extent are the least apparatus for measuring a zero-positive quantity, and no fewer terms will do.

Choosing a unit is a convention, and making a comparison against it is not. For a zero-positive quantity the choice of origin is no convention at all, since only a true zero makes the ratio of two readings a finding. Temperature shows the difference between a true zero and a zero chosen by convention. The Kelvin scale has its zero at the absence of thermal motion (the jiggling of atoms and molecules that heat is), and twice the Kelvin temperature is twice as much of something. The Celsius scale has its zero at the freezing of water, an offset chosen for convenience, and twenty degrees Celsius is not twice ten degrees Celsius. On the Kelvin scale, ten degrees Celsius is about two hundred and eighty-three, and twenty degrees Celsius is about two hundred and ninety-three (the gap between the two zeros is set by the definitions of the two scales), so at twenty degrees Celsius there is only a little more thermal motion than at ten, and not twice as much. The Celsius offset is the kind of distortion that a false zero brings.

Physicists also speak of negative absolute temperatures in certain special systems. There the negative number reflects how temperature is defined, and it names no quantity below the absence of thermal motion, so for ordinary matter the Kelvin zero is the absence described here. The Appendix gives those systems in the physicist's terms.

Origin, unit and extent are the three things any scale of measurement must fix. The origin fixes where the measure starts: a position. The unit fixes how large one step is: a scale. The extent fixes which way, and how far, the measure runs: a direction.

Each of origin, unit and extent also faces one of three scales of description. The unit is the standard the comparer holds, made and kept at the comparer's own scale, so the unit carries the meaning of the middle scale, the *mesoscopic*, the scale of the comparer itself. The origin is the limit approached by ever finer distinctions, the point below which nothing more of the quantity can be found, so the origin marks the boundary of the very small, the *microscopic*. Cut a cake in half, and the half in half again: each cut makes a finer distinction, and the pieces shrink toward nothing. The road toward the zero of any quantity, whether length or warmth, is a road of that kind, ever finer distinctions of less and less, and that is why the origin faces the very small. The extent is the limit approached by ever greater reach, so the extent marks the boundary of the very large, the *macroscopic*. Every measurement therefore carries the very small, the middle and the very large as its origin, unit and extent, and the three scales are present in the build of every reading.

The Same Act Everywhere

The six parts can be seen at work in cases that run from the made to the living, and from the everyday to the quantum. In each case the same six intrinsics appear, and measurement, perception and comparison turn out to be one act, done in different materials and at different scales.

A thermometer turns a physical process into a reading with a meaning. The reading, the height of the column against its marks, is the content; the setting in which it is read, this room at this place and moment, is the context; the graduated scale is the unit, the standard of comparison. The same temperature gives the same reading on different occasions: that is sameness of content across contexts, which Chapter 2 calls symmetry.

Different temperatures give proportionately different readings, a small change giving a small change: that is content keeping in proportion with context, which Chapter 2 calls continuity, and continuity is proportion, not constancy (Chapter 2 shows how proportion fits its definition of continuity). Difference shows wherever a reading departs from a mark. The subjective side is the thermometer, with its scale and its placing, from which the reading is made; the objective side is the air of the room, about which the reading is made. And the zero-positive shape appears as well: the origin is absolute zero on the Kelvin scale, the unit is the degree, and the extent is the range of the instrument.

A single electron spin gives the smallest measurement: one bit of information (one answer between two possibilities), drawn from one quantum system (a system small enough that quantum theory, the theory of matter and light, whose findings come in countable steps, is needed to describe it). The experimenter chooses a direction to measure along, called an axis; the electron gives a result, up or down. All six intrinsics are present even here. The spin is the content; the chosen axis is the context; the result is one of two, and the distinction between the two results is what makes the act a measurement (sameness with one result, difference from the other); the apparatus, with the axis the experimenter set, is the subjective side; the electron registered is the objective side. The single spin belongs in this list for a further reason: a famous result rests on measurements of exactly this six-part kind. The free will theorem of John Conway and Simon Kochen (2006), which Chapter 9 sets out, rests on a case with exactly this structure: a pair of entangled particles of a kind physicists call spin one (particles prepared together, whose results, measured far apart, agree in a pattern; spin one means that their built-in turning is twice the electron's). The theorem's case adds one feature: the agreements between the distant measurements are the measured ground of the theorem's conclusion.

The human eye is a measuring instrument made of living tissue. It compares photons (single particles of light), registering whether they are present or absent, their frequency (which we see as color), and the direc-

tion from which they arrive, and it turns the result into perception. The same six intrinsics apply, and the act of seeing is, in its structure, the act of measuring. Seeing differs from a thermometer's measuring only in material, living rather than manufactured, and in scale, the sensitivity of the retina (the sheet of light-sensitive cells at the back of the eye) rather than the range of a thermometer. The comparisons of seeing can be named one by one. An edge is a difference of brightness or of color between neighboring points of the retina. Contrast is one region found brighter or darker than the region around it, and so the same gray patch looks lighter on a black ground than on a white one. A shape, as seen, is a pattern of such samenesses and differences across the field of view. Seeing an edge, a contrast or a shape needs only sameness and difference across one field of view, and no absolute measure of light.

The smallest act of knowing is a question that admits exactly two answers, yes or no. All six intrinsics are present. The question supplies the context; the answer supplies the content; the asker is the subjective side; what the question is about is the objective side; and the answer is the verdict: sameness with one of the two answers, and difference from the other. A question with several possible answers, or with no list of answers given in advance, widens the range of possible answers without changing that structure: the question still supplies the context, and the answer still finds sameness with one possible answer and difference from the rest. A question with no range of possible answers at all asks for no comparison, and so it is not a question. The six-part structure of comparison reaches beyond physical measurement into knowledge itself: every act of knowing is an act of comparison, and every act of comparison has the same six parts.

Two kinds of faithfulness apply to every instrument above. *Precision* is sameness of content across repeated measurements in the same context: the readings cluster together. In the terms of Chapter 2, precision is continuity. *Accuracy* is sameness of content between the instrument's result and a reference standard read in a different setting (the standard is itself the result of a comparison: the unit, made real in some object or procedure): the readings agree with the standard. In the terms of Chapter 2, ac-

curacy is symmetry. An instrument with precision and without accuracy repeats its error; an instrument with accuracy and without precision scatters around the right answer. Both are needed, and neither reduces to the other.

The cases come together in a single claim: measurement, perception and comparison are one act. They differ in material, in scale, and in the words used to describe them, and they share one structure (the six intrinsics, the three ways of looking, and the three divisions every measurement makes actual), so that no feature of one is absent from the others. What tells them apart is the material in which the act occurs (an instrument, a sense organ, a question put to the world), and so the side from which the act is considered, and never its structure.

Where Physics Has Edges

Physics, on this account, is measurement understood as interaction, bounded above and below, and measured by six quantities that answer to the six intrinsics. Measurement is interaction, interaction is comparison, and comparison has a *domain*: the range within which comparisons can be made. So physics has boundaries. Physics has two, and both are constants that every physicist knows. The knowing self has a third, at the middle scale, and that third boundary carries no constant; Chapter 18 sets out all three.

At the boundary of the very small stands Planck's constant. Planck's constant is a fixed amount of what physicists call action (a quantity made by multiplying an energy by a duration). Some quantities come in pairs that physicists call conjugate, such as the position of a particle and its momentum: neither member of the pair can be made perfectly sharp together with the other, and once the two are pinned down finely enough, sharpening one spreads the other. Below the amount of action that Planck's constant marks, no measurement can tell apart the differences that would fix both members of a pair of conjugate quantities together. And whatever

crosses that boundary into the domain, arriving from below the smallest difference that can be told, can appear only as what we call random, because nothing about it could have been compared before it arrived. Chapter 5 comes back to this edge.

At the boundary of the very large stands the speed of light. Beyond that boundary, no signal can arrive (a signal is anything that passes from one event to another and lets a record at the second depend on an act at the first). What lies past that boundary is unreachable by the structure of comparison itself, and not slow to reach us.

On this account, neither constant is a speed or a size in the first instance. Each constant is an edge of the domain of comparison: Planck's constant on the side of the smallest difference that can be told, the speed of light on the side of the farthest connection that can be made. Between them lies everything that can be measured, which is to say everything that is physics.

The two great theories of twentieth-century physics can be read, in these terms, as statements about access: about what a comparer can receive, and how much. *General relativity*, the theory of gravity as the shape of space and time, sets the limit of signaling. An event is one place at one moment, and the light cone of an event is the set of all events from which a signal can reach it (Chapter 12 draws the cone in words). The light cone defines what an observer at our own middle scale can, in principle, receive information from. *Quantum theory*, the theory of matter and light whose findings come in countable steps, sets the limit of extraction. Its uncertainty relation (the rule that the two members of a conjugate pair cannot both be fixed sharply at once) defines how much information an observer at our scale can, in principle, draw out of a very small system.

Physics has three such limits: no influence faster than the limit of causal signaling; nothing returning across a horizon (a surface, like the edge of a black hole, from beyond which nothing comes back); and no exact joint knowledge of conjugate quantities. The first two limits are two faces of the absolute elsewhere, the region of events that no signal from here can

reach and from which no signal can reach here; the third is the uncertainty relation. Beside them stands the limit in time that every comparison carries: no backward causation, no effect arriving before its cause. The three limits of physics and the limit in time are together the limits on signaling and knowing, and counted as in Chapter 9 the limits on signaling and knowing are three: the absolute elsewhere, the uncertainty relation, and no backward causation. They answer, one to one, to the three boundaries named at the start of this section: the absolute elsewhere to the edge the speed of light marks, the uncertainty relation to the edge Planck's constant marks, and the limit in time to the boundary of the self, where every comparison is made, now.

Each of the two theories states its limit in a deterministic formalism (a formalism is the mathematics in which a theory is written, and a deterministic one is mathematics in which each state fixes the next). Quantum theory, read this way, is a theory of the limits of how much content can be known; relativity is a theory of the limits of how signals can flow. They are two controls on access, and the mathematics in which each is written says nothing about whether what is accessed is itself determined.

The structure of comparison establishes that such limits exist, and that no future theory will remove them. The form of the limits, and the values of their constants, belong to physics, which finds them by measurement. The theorem of Chapter 3 explains why such limits exist at all; their form, the speed of light and the value of Planck's constant come from physics, and the theorem derives none of them.

One question about quantum theory belongs here, because it is the place where the third division, the actual from the possible, meets that theory. Does this account solve the measurement problem of quantum theory? The account says where the measurement problem lies. The measurement problem is the question of how one definite outcome arises from a mathematics that describes only a smooth spread of possibilities. What the smooth mathematics leaves over is the measurement problem itself: the jump to one definite finding, which a smooth, unbroken mathematics does not contain. The reification triple says what every measurement does

at that point: it turns the possible into the actual. The derivation of quantum mechanics from this ground, with its full treatment of the single outcome, is the subject of a companion book, *The Quantum Passage*.

The Map and the Land

A map can be accurate in every detail a traveler checks, and still not be the land. Many maps fit the same journeys; none of them is the ground walked on.

The same holds of every model in science. A model that agrees with a body of measurements is compatible with them, and compatibility is not equivalence: many models may agree with the same data, and no amount of agreement can establish that a model is the reality it describes. That follows from the structure of measurement itself. Every measurement is made from a position, at a scale, by a comparer that cannot stand outside the act of comparing, and its result always reflects both the system measured and the conditions of the measuring. I put the difference in four words, law, exists, measurement and real, joined in one short line: "A law exists; a measurement is real".

A relation laid out as structure, such as a law written down, exists: it stands the same for every comparer who reads it. An interaction is real: it happens, here and now, in an actual act of comparison. Chapter 4 makes the words exist and real exact, and shows that interaction comes before existence, reality and objectivity, since each of the three is confirmed only by the possibility of interacting with what is said to exist, to be real, or to be objective.

Because a relation laid out as structure exists while an interaction is real, mathematics is not physics. Physics is the study of measurement, and a physics that claims to describe something beyond all possible measurement has stopped doing physics. I take that consequence to mark the boundary of the discipline. A physics beyond all measurement has not

thereby started doing metaphysics either. Metaphysics is the study of the structure of comparison that physics already is, and what no comparison could ever reach can be the basis of no assertion, in physics or out of it.

Where the Ground Leads

The ground laid in this chapter reaches every later chapter, and its shape can be seen from here. Chapter 2 builds four ideas from the six intrinsics, by applying sameness and difference to content and to context: symmetry (the same across settings), continuity (unbroken, with no jump), asymmetry (changing as the setting changes) and discontinuity (a jump, a step). Four ideas can be paired in ten ways, counting each idea paired with itself. Chapter 3 shows that of those ten pairings, with both members applied completely and at once, two are forbidden and two survive. The other six drop out for plain reasons that Chapter 3 gives: four pair an idea with itself and say nothing new, and two pair an idea with its opposite and contradict themselves. Symmetry with continuity is forbidden, and so is asymmetry with discontinuity; continuity with asymmetry survives, and so does symmetry with discontinuity. I call that result the Incommensuration Theorem (incommensurate means having no common measure). The Incommensuration Theorem is a structural fact about the nature of knowing.

Because comparison has two survivors and two forbidden pairings, physics has a definite shape at its limits. Where physics asks for perfect symmetry, it finds discontinuity: that is quantization, the coming of things in whole units that cannot be split. Where physics asks for perfect continuity, it finds asymmetry: that is the arrow of time, the one-way running of time. Where physics asks for both at once, it finds that it cannot have them: that is the failure of determinism (Chapter 5), the failure of any attempt to unite the theory of gravity and the theory of the quantum in one perfected theory (Chapter 17), and the failure of any hidden variable (a

hidden setting that would fix each outcome in advance) to restore what John Bell's theorem of 1964 took away, the picture in which every outcome is fixed beforehand by causes close at hand (Chapters 10 and 11).

In this book, symmetry belongs with structure, continuity belongs with the source of the new, and interaction itself belongs with process. This book's placing of symmetry with structure, continuity with the source of the new, and interaction with process puts the two kinds of law themselves among the ways of looking; where each kind of law belongs is a different question from the two addresses, row and column, that each of the six measures has. General relativity and quantum theory are the two survivors of the Incommensuration Theorem. Each of the two theories uses the tool that belongs to one of two ways of looking, structure or the source of the new, in the domain that belongs to the other way of looking, as I show in Chapter 17.

Theories that lay a surface of appearances, which counts for less, over a hidden reality, which counts for more, contradict themselves. The cat of Erwin Schrödinger's thought experiment of 1935 (a cat sealed in a box, whose living or dying is tied to a single atom) is undecided for those outside the box because being decided is a relation that only an interaction can make actual, and not because looking would disturb the cat. Chapter 4 shows both: why such theories contradict themselves, and why the cat is undecided for those outside the box.

The ground is complete for the work of this book: Chapter 4, once it has set out the three axioms, shows that reasoning from them never leads outside them, and that nothing the account needs lies outside the ground. The argument of this book needs nothing beyond what the book itself states.

The arguments built on this ground are exact: each step is fixed by the steps before it, and every definite thing follows from them. What they leave unfixed is unfixed completely, and so randomness, where it appears

in the world, is hard: no prior specification fixes it, however complete, for any comparer whatever. I show in Chapter 8 why randomness, where it appears in the world, is hard.

Every measurement is a comparison made in time. That fact bears on one more question, whether a computation, which runs through its steps by a fixed rule, could be what a mind is. The question belongs to the third request, a world run as a program, and Chapter 22 answers it.

What We Now Know

We now know what happened when the physicist measured. Something was compared with something: a quantity with a standard, a reading with a scale, a signal with its background. The comparison was an interaction, and the interaction was a comparison, and the record it left is a finding. The finding has six parts, and it divided past from future, here from there, and the actual from the possible. Whatever is true of comparison as such is true of every physical interaction as such.

Because whatever is true of comparison is true of every physical interaction, the ground laid in this chapter matters for the question of this book, what cannot be and why. A limit of comparison is a limit of physics, and it can be found before any experiment is done, because it belongs to what an experiment is. The things this book shows cannot be (a message sent faster than light, a return to last Thursday, a world that runs as someone's program, a universe that runs like clockwork) are each a limit of comparison, met in a different place.

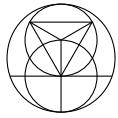
Summary

A measurement is a comparison: the act of finding whether two things are the same or different, in some respect, made from inside the world against a standard. Interaction, comparison and relation are one act named three ways (the Root Tautology), so measurement, observation, perception and

interaction are one act under different descriptions, and whatever is necessarily true of comparison is necessarily true of physics. Every comparison, and so every measurement, has six parts in three pairs (sameness and difference, content and context, a subjective side from which it is made and an objective side about which it is made), and neither member of a pair has a meaning without the other. Every measurement makes three divisions actual: past from future, here from there, and the actual from the possible. Physical measurement measures six measures, three contents each in its own context. Zero-positive quantities need an origin, a unit and an extent, which face the very small, the middle and the very large. Physics is bounded by two constants, Planck's constant and the speed of light, and each of the two constants is an edge of the domain of comparison. The limits of knowing exist because of the structure of comparison, and the constants that mark those limits are found by physics. Every measurement has a subjective side, every perception has an objective side, and every comparison has both content and context. The scope of these findings is every act of comparison, by a person, an instrument or an atom, at every scale.

Of the three kinds of claim, this chapter makes the first two. That interaction, comparison and relation are one act, that every comparison has six parts, that every measurement makes three divisions actual, and that physical measurement has six measures are claims about the structure of comparison itself. That Planck's constant and the speed of light are the two edges of comparison, and that the two great theories state the limits of access, one of signaling and one of extraction, are readings of known physics: each gives the reason of structure for which a known result holds.

If every finding is a sameness or a difference, of some content across some context, what kinds of sameness and difference can a law of physics describe?



Chapter 2

Two Kinds of Law

The Same Law on Tuesday

Why do the laws of physics look the way they do? A law of physics says that something holds: the total amount of energy stays the same; when one body pushes on another, the other pushes back just as hard; a moving body changes its course smoothly, never by a jump. Why laws of these shapes, and not of some other shapes? This chapter answers that question, and the answer starts from the act that the last chapter found beneath every measurement.

Picture an experiment on a laboratory bench. On Monday a student lets a ball roll down a ramp and times it with a stopwatch. On Tuesday she carries the whole apparatus across the room, turns it to face the window, and runs it again. The ball takes the same time. The place has changed, the day has changed and the direction has changed, and the law of the rolling ball has not changed at all. Physicists expect exactly that result: the same law, in another place, on another day, facing another way.

Chapter 1 showed that every measurement is a *comparison*: the act of finding whether two things are the same or different, in some respect. The *content* of a comparison is what the comparison finds. The *context* of the comparison is the setting that makes the two things compared two, such as two places or two moments. In the student's experiment, two runs are compared, Monday's and Tuesday's. The content, the time the ball takes,

came out the same. The context, the place, the day and the direction, came out different. Same content across different context: that is the first of four simple patterns that a comparison can find, and the laws of physics take their shapes from two of them.

Four Ways to Be the Same or Different

Chapter 1 found six parts in every comparison, the *six intrinsics*: sameness and difference, content and context, the subjective and the objective. Every comparison needs all six. The intrinsics can also be applied to one another. Apply the pair sameness and difference to the pair content and context. The content can be found the same or different: two choices. The context can be found the same or different: two choices. Two choices for the content and two for the context give four patterns, and the count of four is set by the structure itself, two times two. I call the four patterns the four *compounds*, because each is made of two findings at once, one about the content and one about the context.

Continuity is sameness of content within sameness of context. Hold the setting fixed, and the content does not jump. A small shift of the setting brings at most a small shift of the content. Turn the knob of a dimmer a little, and the light changes a little. The sameness in the definition is sameness between settings so close that they count as one setting: there, the content found is the same, and so, as the setting moves on a little, the content moves on only a little, and nowhere jumps. That is how the light of the dimmer can change and still be continuous, and it is what Chapter 1 meant in saying that continuity is proportion, not constancy. Every differential equation assumes continuity; a differential equation is the kind of equation physics uses to say how a quantity changes from one point or moment to the next point or moment close beside it. Every smooth field expresses continuity; a field is a quantity spread through space, with a value at every point, as the temperature of a room has a value at every point in the room.

Discontinuity is difference of content within sameness of context. Hold the setting fixed, and the content changes anyway: a jump, a break, a threshold crossed. An atom can hold only certain separate amounts of energy, called its energy levels, with nothing allowed between them, and it passes from one level to another by a jump. A detector clicks, and the click comes whole. One particle is different from two particles, and nothing lies between one and two.

Symmetry is sameness of content across difference of context. Change the setting, and the content is unchanged. Symmetry is the pattern we began with. Move the apparatus across the room, run the experiment on Tuesday instead of Monday, turn the directions from which every position is measured: the same law comes out, and the same conserved quantity, the same unchanging total that I take up in a later section of this chapter. A content that stays unchanged when its setting changes is called *invariant*, and every invariance is a symmetry in this sense.

Asymmetry is difference of content across difference of context. Change the setting, and the content changes with it. The past is not the future; here is not there; the record of a process is not the process run backward, as a footprint in wet sand is not a step taken in reverse. Every direction is an asymmetry. Every gradient is an asymmetry; a gradient is a slope, such as the fall of heat from a hot stove to the cool far side of a room. Every irreversibility, every change that does not run back, is an asymmetry.

These four definitions are built from the intrinsics of comparison alone. Nothing in them is imported from physics, and nothing in them is borrowed from mathematics. That the four definitions are built from the intrinsics of comparison alone matters for everything that follows. Chapter 1 found that interaction, comparison and relation are one act, named three ways, and called that finding the *Root Tautology*. So whatever holds of symmetry and continuity because of these definitions holds wherever comparison holds; by the Root Tautology, that is wherever interaction holds; and wherever interaction holds is everywhere physics reaches.

Two Kinds of Law, and No Third

Wherever there is an idea of lawfulness, its laws are of only two basic types. Every law is either a law of symmetry or a law of continuity. Every law that has ever been written down, in any field, is one, or the other, or a combination of the two, and there is no third type.

A *law of symmetry* is a law of invariance: it says that a content stays the same across a difference of context. Three examples show the pattern, and in each we can name the content and the context.

The *conservation* of energy is a symmetry law. To be conserved is to stay the same in amount through every change. The total energy (the content) is the same across different times and different arrangements (the context).

Newton's third law, which Isaac Newton published in 1687, is a symmetry law. When you push on a wall, the wall pushes back on you with a force of the same size. The force (the content) is the same when we exchange the body that pushes and the body that is pushed (the context).

The principle of relativity is a symmetry law about laws themselves. Albert Einstein made it a founding principle of physics in 1905: the laws of physics are the same in every *inertial frame*, which is any setting that moves steadily, without speeding up, slowing down or turning. Pour a glass of water on a train gliding along a straight, smooth track, and the water pours as it does at home. The content is the lawfulness itself, and the context is the frame.

A *law of continuity* is a law of persistence: it says that a content stays the same within a sameness of context, carrying on without a break. That a thrown stone follows a path with no jumps in it is a continuity law. That a body's state changes smoothly with the forces upon it is a continuity assumption of physics, and the setting in which the body changes, space and time, is governed by continuity. That the self who wakes in the morning is the self who fell asleep is a continuity law. That a stream of experience has no absolute gap within a single experience is a continuity law.

Why only two kinds? Because of what a law is. A law is a statement of what stays the same. Of the four compounds, two are statements of sameness of content, and those two are symmetry and continuity. Discontinuity and asymmetry are statements of what changes, and a law about what changes is, on examination, a law about what stays the same under that change. The thrown stone changes its place at every moment of its flight; the law of its flight states what stays the same through all that changing, the same pull of the Earth at every moment. There is no third kind of sameness for a law to be about. So the two kinds of law are all the kinds there are.

World and Self

Each kind of law has its home. To say where, I need three words that Chapter 4 takes up in full. Where the foundations of any domain rest on three ideas that together are necessary to it and enough for it, those three are its *modalities*, its three basic ways of being. The *immanent* is the mode of process: the relational, the here and now, what is happening between things. The *omniscient* is the mode of structure: what holds throughout, the same for everyone who looks. The *transcendent* is the mode of what makes things possible: the creative, and, as we see in a moment, the single self that persists as itself.

In my assignment, symmetry has the nature of the omniscient modality. Continuity has the nature of the transcendent modality. Interaction has the nature of the immanent modality. I keep that assignment wherever I speak of the two kinds of law; where a later chapter pairs the modalities with something else, that chapter says so where it makes the pairing.

Symmetry has the nature of the omniscient because the omniscient is the mode of the domain as a whole: of what holds throughout, of the same content across every context. The same content across every context is symmetry. The omniscient is also the mode of the object, of existence,

and of what is known. Continuity has the nature of the transcendent because the transcendent is the mode of identity: of the single thing that persists as itself, of the same content within its own context. The same content within its own context is continuity. The transcendent is also the mode of the subject, and of what is understood. It is the mode, too, of objectivity in the strict sense of that word: the standpoint that is no particular standpoint. (In the ordinary sense of the word, the objective is what lies outside a self, and that ordinary sense belongs to the world side, as the account of world and self below shows.) The immanent is the mode of relation: the interaction in which knowing and understanding both happen, the reality in which object and subject meet. Laws, accordingly, have the nature of the omniscient; principles have the nature of the transcendent; axioms have the nature of the immanent. The reason lies in what each mode was just said to be the mode of. A law states what is known, the same for everyone, and the omniscient is the mode of what is known. A principle, such as the principle of relativity that this chapter meets below, states what is understood about every law of a domain, and the transcendent is the mode of what is understood. An axiom states a basic relation from which both knowing and understanding proceed, and the immanent is the mode of relation.

Holding the same for whoever looks, and being objective in the strict sense, can sound like one idea, and a picture shows them apart. A score sheet from a game of chess lies on a table, and everyone who walks in reads the same list of moves. The score sheet is the omniscient: something laid out, already there, the same for whoever looks. Now say that the position on the board is lost for white, a judgment that any strong player would reach. That judgment speaks of every player who might ever look at the board, those who have looked and those who never will. The judgment that the position is lost is a claim about the whole range of possible observers, and so the judgment is a claim about possibility, which is the home of the transcendent. The score sheet holds the same for whoever looks; the judgment says what any possible observer would find. Chapter 4 comes back to this game.

The same assignment can be said in terms of world and self. The basic being, nature and lawfulness of all world is defined by symmetry and founded on it. The basic being, nature and lawfulness of all self is defined by continuity and founded on it. What lies outside the boundary of a self is objective in the ordinary sense, and has the nature of existence; it has also the nature of the omniscient. Its lawfulness is a lawfulness of content: the same content across every context, the same law everywhere. What lies inside the boundary of a self is subjective, and has the nature of creation; it has also the nature of the transcendent. Its lawfulness is a lawfulness of context: the frame that persists as itself, the standpoint that is never exchanged for another. The meeting of the two, the interaction in which a self perceives the world and expresses itself into the world, has the nature of the immanent. That meeting is where measurement happens, and so it is where physics happens.

So: symmetry governs world; continuity governs self; and the interaction between them is where physics is done.

I call this the *universal law*. The universal law is the principle of structure that says which kind of law governs each side of reality; it is not one more law inside some particular domain. Physics, the study of world, finds its content under symmetry law, and finds the setting in which that content is found, space and time, under continuity law. So when this chapter says that symmetry governs world, it speaks of the world's content, what is found in it; the setting of that content is a context, and contexts fall under continuity law, as the lower row of Chapter 1's six measures showed. Psychology, the study of self, is a system of continuity laws, since the laws of a self are laws of persistence, like the law that the self who wakes is the self who fell asleep. Physics and psychology do not compete, because their laws are laws of different things: of world and of self. And, as this chapter goes on to show, they could not have been laws of one and the same thing, both perfected together.

Conservation: Symmetry Said as a Number

Conservation is what symmetry looks like when it is stated as an amount. The usual story runs the other way round: experiments found that certain totals stay fixed, and the formulas came afterward to express what had been found. On the account I am giving, conservation laws have an origin prior to the experiments that found them.

Take the definition. Symmetry is sameness of content across difference of context. Now ask what it is for a *domain*, a range of settings within which one description holds, to be governed by symmetry. It is for there to be some content that stays the same however the context is varied within the domain. Name that content, and you have named a conserved quantity.

I have a bottle. I put some water in it, and I put a cork in it. The water is going to stay in the bottle: it is conserved. Five months might pass, and there is still the same amount of water in the bottle. The amount of water is the content, and it is the same; the moment at which we look is the context, and it is different. The water in the corked bottle is a case of symmetry: sameness of content, difference of context.

Vary the context in time, and the content that does not vary is energy. Vary it in space, and the content that does not vary is *momentum*, the quantity of motion a moving body carries, larger for a heavier body and larger for a faster one. Vary it by turning, and the content that does not vary is *angular momentum*, the quantity of turning motion, which keeps a spinning top spinning. The conservation law is the symmetry itself, said in the form of a number that does not change. It is nothing extra that a symmetric domain happens also to obey.

Here a famous result of mathematics enters. In 1918 Emmy Noether proved a theorem about every system whose laws can be written in one standard form, which physicists call the Lagrangian form (a way of writing all of a system's laws through one quantity, from which the laws follow). For such a system, to every smooth symmetry (a symmetry that can be applied by any amount, as a turn can be made through any angle, how-

ever small) there corresponds a conserved current, a quantity that can flow from place to place and is never made or lost. Noether's theorem is a theorem of mathematics, and its content is exact.

I claim that conservation is a necessity that comes before any particular system is written in the Lagrangian form, and that Noether's theorem is the exact form, inside one particular and very powerful kind of mathematics, of a feature of structure that holds of every symmetric domain as such: wherever a content is kept the same across changes of its context, there is a matching conservation. That holds of every symmetric domain, whether or not it can be written in her form, and whether or not anyone has written the conserved quantity down. The Appendix gives, in the physicist's terms, how far the matching of symmetry with conservation reaches beyond the systems her theorem covers, and how this account and her theorem stand to each other. Her theorem and this account do different work. Her theorem does work of its own: it tells the physicist exactly which conserved quantity goes with which symmetry, and how to compute it. The account given here tells the physicist why there was going to be one.

What a Group Is

Physicists describe symmetries with a piece of mathematics called a *group*. I need the idea for one reason: a symmetry is a set of actions that leave some content unchanged, and the group is the whole set of those actions; Noether's smooth symmetries, which the next section comes back to, are groups whose actions can be made by any amount, as a dial can be turned by any amount. Group theory looks at patterns and symmetries. Nathan Carter, in 2009, gave an intuitive way in, using four conditions on a set of actions. Picture a Rubik's cube, whose actions are the turns of its faces.

First, the actions are well defined: the list of possible moves never changes, and the basic moves from which all the others are built, called the generators, stay the same. Second, the actions are reversible: any move

can be undone, and any sequence of moves can be undone by making the reverse moves in the reverse order. Third, the actions are deterministic in the mathematician's sense, which says only this: the same action on the same arrangement always gives the same result. (Chapter 5 takes up a much larger claim that goes by the name determinism, the claim that everything in the world is fixed in advance; the mathematician's sense here is about moves on a puzzle.) Fourth, the actions are composable: moves can be combined in any sequence, and any sequence of moves is itself a move. That fourth condition is called closure. Mathematicians write the four conditions as the four axioms of a group, and the Appendix gives that formal form. Group theory is about what follows from these four observations.

A group is wherever all four conditions hold. The group is the whole set of actions generated by the generators, and not the whole set of arrangements: for the cube, the group is the set of all the ways of turning it, and not the set of all its colored patterns. Every group can be seen as a collection of ways to rearrange some set of things.

One more remark explains why physical law is simple enough to have conservation laws at all. The simpler a thing is, the more its relations resemble the relations of every other simple thing. Simple things are the same as simple things: they have similar relations, and they behave in similar ways. So the laws that relate simple things are themselves simple, and at the most basic level the variety of the world reduces to a small number of structural types, whose laws reduce in the same way. Because the variety of the world reduces to a small number of structural types, a law can hold everywhere and for everything, and a law that holds everywhere and for everything is a symmetry law. The simplicity of the most basic laws is what one would expect wherever symmetry is the governing type of law, and so it is no lucky feature of this world.

A First Look at the Limit

Now let the two kinds of law meet. Ask for one description that is both: perfectly the same across every change of setting, and perfectly unbroken within every setting. Write out what that asks. Symmetry says: the same content, when the context differs. Continuity says: the same content, when the context is the same. Put the two together, each perfected, and the content is the same whether the context changes or does not change. Then the content does not depend on the context at all. The difference between content and context has done no work in the description, which is to say that the difference has collapsed. Content and context are two of the six intrinsics, and every comparison needs both of them. So a description in which content and context have collapsed into one is no longer the description of any comparison.

That limit has a name: the *Incommensuration Theorem*. Symmetry and continuity cannot both be perfected in one description of one thing, and neither can asymmetry and discontinuity. The Incommensuration Theorem is a structural fact about the nature of knowing. Chapter 3 sets it out whole, with all four compounds and every way of pairing them. We need only the half that we have just written out: perfect symmetry with perfect continuity, together in one thing, is the *forbidden pairing*.

The limit is a limit on perfection. Within any finite domain, symmetry and continuity occur together, each to some finite degree, and most of working physics lives there. You may ask whether Noether's smooth symmetries are the forbidden pairing, since they are both symmetric and smooth. They show how the two can stand side by side without being perfected in one thing. In a smooth symmetry, the symmetry is what is found: the conserved quantity. The smoothness lives in the mathematics, as the smooth dial of the group, the amount of turning, which can be set to any value. The two hold of different things, one of what is found and one of the mathematics in which it is written, and Noether's theorem joins the two levels. So a smooth symmetry is not the forbidden pairing.

From the forbidden pairing follow two consequences, one for each kind of law. What is perfectly symmetric must be discontinuous at its limit. What is perfectly continuous must be asymmetric at its limit.

Four Theorems of Being

Apply those two consequences to self and world. From the universal law and the theorem together, four propositions follow, and they organize the rest of this chapter. They are called the *four theorems of being*: two are about the self, the subjective side, and two are about the world, the objective side; two are about content, and two are about context.

Subjective context is continuous. The self persists. Personal identity endures. The stream of experience has no absolute gap. The continuity of subjective context is the most basic feature of what it is to be a self.

Subjective content is discontinuous. Thoughts, perceptions and experiences are necessarily many and varied. No one can think the same thought forever, or perceive the same thing always. The content of a self's experience must vary, must change, must be many.

Objective content is symmetric. The laws of physics are the same everywhere. Energy, momentum and charge are conserved. The structural content of the world is invariant when its setting is transformed. The invariance of the world's structural content, when the setting of that content is transformed, is the basis of all physical law.

Objective context is asymmetric. Time has an arrow. *Entropy*, the measure of how far a thing's energy has spread out and mixed, has a gradient, a one-way slope; Chapter 13 takes entropy up in full. The past is not the future. The frame within which physical events happen has a direction, and it does not run backward.

Each of the four is an instance of the theorem, and the theorem says more than the four do: it says what second character each of the four must also have. Subjective context is continuous, so it must also be asymmetric: each self is single, each perspective unrepeatable, each moment of experi-

ence its own. Objective content is symmetric, so it must also be discontinuous: matter comes in whole units, energy levels are separate steps, the structural content of the world comes in units. The theorem allows exactly two pairings of the compounds to be perfected together, continuity with asymmetry and symmetry with discontinuity; Chapter 3 shows why exactly these two. They are exactly the two pairings that self and world show: the self is continuity with asymmetry, and the world is symmetry with discontinuity. In the rest of this chapter I take the world side first and then the self side. The new results of the chapter are two, which I call shadows: quantization and uniqueness.

The Shadow of Symmetry: Quantization

I apply the Incommensuration Theorem to the world side, where the content of the world is symmetric. Symmetry and continuity cannot both be perfected together, so the content that symmetry describes, followed to its limit, must be discontinuous. A content that is perfectly invariant comes in units. The coming of a perfectly invariant content in units is *quantization*: coming in whole units that cannot be split.

Let me say it more slowly, because the claim is strong, and I want you to see exactly where it rests. At its limit, a symmetric content is discontinuous, and the reason can be followed step by step. Take a content that is invariant under every change of context its domain allows. Were that content also to vary smoothly within a fixed context, it would be the same whether the context changed or did not change. By the argument we wrote out for the forbidden pairing, the content would then have come loose from its context altogether: the line between content and context would have collapsed, and the comparison that defines the domain would no longer be a comparison. So at the limit a symmetric content is not continuous. It is discontinuous; it is *discrete*, which means that it comes in sep-

arate pieces that can be counted. The invariant comes as a countable thing: so many units of charge, so many units of angular momentum, one electron or two, and never one and a half.

When I say that quantization is the shadow of symmetry, I mean that a perfectly symmetric content, followed to the limit, must come as a countable thing. Quantization follows symmetry as a shadow follows the body that casts the shadow. The quantum character of the microscopic world, its coming in separate, countable steps, is what anything governed by symmetry laws must be at its limit. It belongs to every domain governed by symmetry, and so it is no accidental feature of our particular universe, to be accepted as a brute fact and built into the mathematics by hand. Where physics has asked for perfect invariance, it has found the discrete, and on this account it could not have found anything else.

The claim about quantization goes a step further than the account of conservation. The account of conservation read a known result and gave the reason of structure for which it holds. The claim about quantization says that any domain whose content is governed by symmetry will show discontinuity at its limit, and it says so before looking. The world we have does show quantization: that is consistent with the claim, and, as I read it, it is what the claim predicts.

The claim establishes that a symmetric content comes in units. Where the steps fall, how large each step is, and which quantities come in steps and which do not in a particular theory are not argued here. Which charges are exactly conserved is a settled matter, and that settled answer shows where the claim meets the physics we have: the algebraic derivation in my book Non-Commutative Geometry and the Immanent Metaphysics finds the unbroken symmetries to be those of electromagnetism and of color. Electromagnetism is the force of electric charge and magnetism; color is the charge that binds quarks, the parts of protons and neutrons, to one another.

Where Conservation Ends

The same theorem that gives quantization also sets a bound on conservation. Perfect symmetry requires discontinuity, so no conservation law can hold perfectly across all conditions. There will be discontinuities: the edges of domains; *phase transitions*, which are changes of state, such as water freezing into ice; and conditions under which the conservation fails, or under which the quantity it conserves stops being well defined.

The bound is a structural statement about domains. It says that the symmetry a conservation law expresses is a symmetry within a domain, and that at the edge of the domain the sameness of content across contexts gives way to something else. It speaks of the edges of domains, so it forecasts no loss of energy in a laboratory, which lies well inside the domain of the conservation of energy. At the largest scale there is, the bound shows itself. Physicists find it hard to define a conserved total energy for a universe that is curved and expanding. On this reading that difficulty is an instance of the bound, and no technical embarrassment.

Asymmetry, the complement of symmetry, shows itself as the differences that conservation cannot remove. Every conserved quantity has a side that is not conserved. Energy is conserved; the way energy is spread out is not. Momentum is conserved; the direction of any one body's momentum can change. Go back to the corked bottle. The amount of water stays the same for five months, while the water itself sloshes, warms, cools and settles. The total is invariant, and everything interesting happens in what the total leaves free.

What is conserved in one domain may be the asymmetry of another. What is a matter of structure, and necessary, at one scale may be a matter of circumstance, and variable, at the next. In every chemical reaction, for example, each kind of atom is kept: the carbon that goes in comes out as carbon, a matter of structure for the chemist. In the core of a star, where atoms fuse, one kind of atom becomes another, and the count that chemistry keeps fixed varies with the circumstances of the star. That what one domain keeps fixed can vary in another is the reason one cannot simply

carry a conservation law across the edge of a domain and expect it to hold in the same form. The bound on conservation and the necessity of quantization are the same fact seen from two sides. Symmetry perfected is discontinuous; discontinuity is where conservation ends.

The Shadow of Continuity: Uniqueness

Now the self side. The law of continuity is the law of self. A person's identity persists; consciousness carries on; the stream of experience is undivided. An identity that persists, a consciousness that carries on and a stream of experience that is undivided are what anything that is a self must have, because to be a self is to be governed by continuity. Continuity is sameness of content within sameness of context, which is to say persistence as oneself. They belong to every self as such, and so they are not accidental features of one kind of animal that happened to evolve here.

The theorem then says what the context governed by continuity must also be. It must be asymmetric. Perfected persistence cannot also be perfected invariance across contexts: a self that persists as itself cannot also be the same as every other self. That a persisting self cannot be the same as every other self is *uniqueness*, being one of a kind. Uniqueness is the shadow of continuity exactly as quantization is the shadow of symmetry.

There is no perfectly symmetric self, and picturing a perfectly symmetric self shows the reason. A perfectly symmetric self would be a self identical to another self: the same content across the difference of context that separates the one self from the other. And it would be a self that persists: the same content within its own context, over time. A perfectly symmetric self, identical to another self and persisting as itself, would be symmetry with continuity, applied together and perfected, in one and the same thing. Symmetry with continuity, perfected in one thing, is the forbidden pairing itself. A perfectly symmetric self would violate the Incommensuration Theorem, and so there is none.

So every self is unique, and unique for a reason of structure (as distinct from a reason of circumstance). Two selves differ in their genes, their histories and their surroundings, of course; identical twins, who share their genes, differ in their histories from the day they are born. The reason of structure is this: the continuity which makes a self a self requires the asymmetry which makes it this self and no other. Were every difference of circumstance removed, the difference of structure would remain, because without it there would be no self to be the same. The singleness of each self is the shadow that continuity casts; it belongs to selves as such, and is no peculiarity of biology.

The two shadows are the two pairings that survive the theorem, and they fall on opposite sides of the boundary of self. On the world side: symmetry with discontinuity, invariant content that comes in units. On the self side: continuity with asymmetry, persistent context that is single. The theorem forbids two pairings and permits two: symmetry with discontinuity is the structure of the objective, and continuity with asymmetry is the structure of the subjective. The match of the two permitted pairings with the objective and the subjective is what the universal law says, reached here by way of the theorem.

Two shadows, then, of one theorem. Quantization is the shadow of symmetry: the content of world, perfected in its invariance, must be discrete. Uniqueness is the shadow of continuity: the context of self, perfected in its persistence, must be single. Each shadow is the second character that the Incommensuration Theorem requires in a surviving pairing: discontinuity beside perfected symmetry, and asymmetry beside perfected continuity. Neither is a fact about this world alone.

Two Theories, Two Tools

The assignment and the theorem together say something about the two great theories of twentieth century physics. *General relativity*, which Albert Einstein published in full in 1916, is the theory of gravity as the shape of

space and time. *Quantum theory* is the theory of matter and light, whose findings come in countable steps. Chapter 17 introduces both more fully; here we need only their characters.

General relativity is continuity with asymmetry. General relativity finds a smooth *continuum* of events, a whole with no gaps anywhere in it, in which between any two events there are always others. General relativity is written in a mathematics that stays the same under every smooth change of the coordinates, the numbers used to label places and times; that invariance is a symmetry of the description, of how the theory is written. And its characteristic results are directed: light cones, which mark at each event the events that light can reach from it and the events from which light can reach it; horizons, surfaces that can be crossed in one direction only; the collapse of a star that does not undo itself. General relativity is the first surviving pairing of the theorem. Continuity is the tool proper to the transcendent. But general relativity is applied to the *macroscopic*, the large scale: the cosmos, the whole; and the macroscopic is the omniscient domain.

Quantum theory is symmetry with discontinuity. Quantum theory finds what is invariant: the conserved quantum numbers, which are fixed counts, such as charge, that label each kind of particle, and the charges themselves. Quantum theory also finds the exchange symmetry of identical particles, by which nothing about the world changes when two electrons swap places. Meanwhile, in its mathematics, a smooth family of transformations, called the unitary group, carries the description forward without a jump. And the characteristic results of quantum theory are discrete: the quantum; the eigenvalue, which is one of the definite values a measurement can return; and the click. Quantum theory is the second surviving pairing of the theorem.

Why the large belongs with the omniscient, and the small with the transcendent, can be seen from what each mode is the mode of. Picture a crowd in a square, seen from a balcony. From up there the whole crowd is laid out at once, the same for anyone who stands on the balcony: that is the view of the whole, the view of structure, and the large scale is the om-

niscient domain. Now go down into the crowd. There is only this one person, here, beside you, and then that one: each a single case, only itself. The transcendent, as this chapter said, is the mode of identity, of the single thing that persists as itself, and the very small is where description comes down to the single event, one click, one electron, one atom's decay, each only itself. So the small scale is the transcendent domain.

Symmetry is the tool proper to the omniscient. But quantum theory is applied to the *microscopic*: the very small, the single event, what is only itself; and the microscopic is the transcendent domain.

Each theory, then, applies the tool proper to one modality to the domain proper to the other. I call this the *cross-modal reach*, and it is required, not accidental. By a theory's tool I mean the character it perfects in what it finds; the theory carries the other character too, exactly, at the other level, in its mathematics. The reach of each theory into the domain of the other is dense and never exhaustive: as close as one likes, everywhere, and never the whole. The fractions marked on a ruler give a picture of such a reach: between any two of them there is always another, as close as you like to any point of the line, and still the fractions never fill the whole line.

A tool applied to its own domain only repeats what the domain already says: a character applied to itself adds nothing to it, as a ruler laid along another copy of itself finds only its own marks. To learn anything about the omniscient, one must reach into it with the instrument of the transcendent; to learn anything about the transcendent, one must reach into it with the instrument of the omniscient. That the two theories reach across is the only way either of them could have been a theory at all, so the reach is part of what each theory is, and no defect awaiting repair. From this it follows that no single theory can unify the two in the strong sense, the sense in which both would be limits of one theory, each being what that one theory becomes within its own range: such a theory would have to find both surviving pairings perfected at once, in what one and the same comparison finds, and that is the one thing the theorem forbids. Chapter 17 walks through that argument in full.

Two Kinds, and Only Two

We now know that there are two kinds of law, and only two. A law of symmetry says that a content stays the same when its setting changes. A law of continuity says that a content stays the same, without a break, within its setting. Symmetry governs world, continuity governs self, and physics is done in the interaction where the two meet. Conservation is symmetry said as a number. Quantization is the shadow that symmetry casts at its limit, and uniqueness is the shadow that continuity casts at its limit.

Knowing that there are two kinds of law, and only two, changes the question this book asks. The chapters ahead meet the forbidden pairing in several places. The claim that everything is fixed in advance needs one law perfectly the same everywhere, together with every state specified without a gap (Chapter 5). Running a process perfectly backward needs laws perfectly the same in both directions of time, together with a state carried back without a break (Chapter 13). Making the two great theories one needs both surviving pairings perfected in one comparison (Chapter 17). In each place the question comes down to one question: can one law be both kinds, perfectly?

Summary

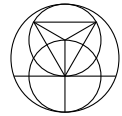
Sameness and difference, applied to content and context, give four compounds: continuity, discontinuity, symmetry and asymmetry. They are built from comparison alone, so whatever holds of them holds everywhere physics reaches. Every law is a law of symmetry or a law of continuity, or a combination of the two, and there is no third type, because a law states what stays the same, and only those two compounds are sameness of content. Symmetry has the nature of the omniscient and governs

world; continuity has the nature of the transcendent and governs self; interaction has the nature of the immanent, and is where measurement, and so physics, happens.

In this chapter I make claims of three kinds. That every law is a law of symmetry or of continuity, and that symmetry governs world and continuity governs self, are claims about the structure of comparison itself. That conservation is a feature of every symmetric domain, with Noether's theorem as its exact form within the Lagrangian mathematics, is a reading of known physics: the reason of structure for which a known result holds. That quantization is the shadow of symmetry, and uniqueness the shadow of continuity, are consequences of the theorem: they say what any symmetric domain and any continuous domain must show at its limit, and they say it before looking.

Quantization is what symmetry must become at its limit, because perfected invariance cannot also be continuous; so the discrete character of the microscopic is the shadow that symmetry casts, and no peculiarity of this world. Uniqueness is what continuity must become at its limit, because perfected persistence cannot also be perfected sameness across selves; so the singleness of each self is the shadow that continuity casts, and no peculiarity of biology. No conservation law holds perfectly across all conditions; every conserved quantity has a side that is not conserved; and conservation ends where discontinuity begins. Each of the two great theories is one of the two surviving pairings, and each reaches with the tool of one modality into the domain of the other. The universal law is the principle that says which kind of law governs each side of reality, and the Incommensuration Theorem is the principle that says what second character each kind of law must have when it is perfected.

We have followed one argument to its end: symmetry and continuity, perfected together in one thing, leave no line between content and context, and so no comparison. Does that limit reach every description whatever: can any description be perfectly the same everywhere and perfectly unbroken at once?



Chapter 3

Why One Description Cannot Be Perfect Both Ways

A Floor and a Ramp

Can a description be perfectly the same everywhere and perfectly unbroken at once? The last chapter ended on that question. This chapter answers that question, and the answer states the one limit that every later chapter of this book uses.

Picture a floor laid with identical square tiles. Every tile is the same as every other, wherever it lies: the same content in every place. Now picture a smooth curved ramp, poured in one piece. Move a little way along it, and its slope changes only a little: there are no jumps anywhere.

Now try to tile the ramp. The tiles can be made smaller and smaller, until the tiled surface follows the curve of the ramp as closely as anyone likes. Yet a tiled surface always has seams, and a curved ramp never repeats itself exactly from one place to the next. A floor can be very nearly repeating and very nearly seamless, both at once. A floor can be perfectly repeating, or perfectly seamless. A floor cannot be perfectly repeating and perfectly seamless at once.

The everyday fact about the floor, that it can be very nearly repeating and very nearly seamless at once, and perfectly either one, but never perfectly both, has an exact form. The exact form is called the *Incommensuration Theorem*, and this chapter is where the book states it in full. It is the

central result of my account of comparison, and of everything that account says about physics; the later chapters will restate it briefly and build on it.

The Four Characters, Again

Chapter 1 showed that every measurement is a comparison: the act of finding whether two things are the same or different, in some respect. Every comparison has a *content*, what is compared, and a *context*, the setting in which it is compared, such as two places or two moments. A comparison can find sameness or difference in the content, and sameness or difference in the context. Sameness and difference, content and context, and the subjective and the objective (the stance from which a comparison is made, and what the comparison is about) are the six intrinsics: the six parts that every comparison must have.

Chapter 2 applied sameness and difference to content and context, and found four characters. In all four characters, the one who compares keeps the same position and the same way of comparing, and only the content and the context vary. *Symmetry* is sameness of content across a difference of context: the same tile in a different place, the same law at a different time. *Continuity* is sameness of content within a sameness of context: move a little and change only a little, with no jump. *Asymmetry* is difference of content across a difference of context: change the setting, and the content changes with it, as the slope of the ramp is different at each place along it. *Discontinuity* is difference of content within a sameness of context: a jump, a click, a seam.

Chapter 2 also found that every law of physics is a law of symmetry or a law of continuity, or a combination of the two. So the question of this chapter can be put exactly. Which of the four characters can be made perfect together, in one description of one thing?

Ten Ways to Pair Four Characters

Four characters can be paired in ten ways. In four of the pairings a character is paired with itself: symmetry with symmetry, continuity with continuity, and so on. In the other six, two different characters are paired. Four and six make ten, and the number is set by nothing more than counting the pairs that four things can make.

Three plain considerations settle all ten pairings.

The first consideration is that pairing a character with itself says nothing new. Symmetry of symmetry says nothing that symmetry alone did not already say, and the same is true of the other three. The four pairings of a character with itself are set aside, and six pairings remain.

The second consideration is that a character and its opposite exclude each other when both are made perfect. Total symmetry with total asymmetry, of the same content across the same settings, would say that the content is both the same and different in exactly the same respect. Total symmetry with total asymmetry is therefore a contradiction. Total continuity with total discontinuity is a contradiction in the same way: the same thing, in the same setting, both smooth and jumping. Two more pairings are set aside, and four remain.

The four that remain are symmetry with continuity, asymmetry with discontinuity, continuity with asymmetry, and symmetry with discontinuity. For these four, the third consideration decides: the parts of a comparison are inseparable, which means that no part of a comparison has meaning without the others.

The Pair That Undoes Itself

Write out symmetry with continuity, each in full. Symmetry is sameness of content where there is a difference of context. Continuity is sameness of content where there is a sameness of context. Now make both perfect, together. The content is then the same whether the setting changes or does not change. So the content does not depend on the setting at all.

Then the line between content and setting does nothing: no change of setting, on one side of the line, goes with any change of content, on the other side. The line has collapsed. But content and context are two of the six intrinsics, and each is defined against the other, the way "inside" is defined against "outside". A comparison in which the two cannot be told apart is a comparison that has not been fully set up.

A small example shows why a comparison needs its content and its context kept apart. Take two ticks of a clock, and ask of them, "Are these at different moments?" The answer was settled before you looked, because being at different moments is what made them two ticks in the first place. When the respect that is being judged is the very respect that makes the two things two, the answer is fixed before the comparison is made, and an answer fixed before the comparison is made is not something the comparison found. The content and the context of a comparison cannot be one and the same respect. A description that makes symmetry and continuity perfect together asks the content and the context of the comparison to be one respect, and so the act of comparing undoes itself.

The collapse of the line between content and setting, when symmetry and continuity are both made perfect, can be seen a second way, by reading the definitions of symmetry and continuity side by side. Across the two halves of the pair, symmetry and continuity, the word "same" appears twice on the side of content, and the word "different" does not appear there at all. The count is a short way of saying what the clock ticks showed: with both halves perfect, nothing in the pair leaves any room for the content to come out different. Sameness and difference are inseparable: neither has meaning without the other, and a comparison whose

content cannot come out different finds no sameness either. So symmetry and continuity, whose definitions together say "same" twice for the content and "different" not at all, cannot both be perfect together.

Now write out asymmetry with discontinuity, and the line between content and setting collapses again, this time with difference everywhere that sameness stood before. Asymmetry is difference of content where there is a difference of context. Discontinuity is difference of content where there is a sameness of context. With asymmetry and discontinuity made perfect together, the content differs whether or not the setting differs. So again the content does not depend on the setting, and again the line between content and setting collapses. This time "different" appears twice on the side of content, and "same" not at all.

The Two Survivors

Two pairs are left, and each of them can have both its halves made perfect without the comparison undoing itself. In each of them, on the side of content and on the side of context, "same" appears once and "different" appears once. Each character is joined with its exact complement, the one character that supplies what it lacks. I call these two pairs *the two survivors*.

The first survivor is continuity with asymmetry. Continuity gives sameness of content within a sameness of setting; asymmetry gives difference of content across a difference of setting. Content and setting both appear, sameness and difference both appear, and the content depends on the setting. Continuity with asymmetry, in which the content depends on the setting, is exactly what anyone would call a smooth process that goes somewhere: an unbroken change with a direction. The poured ramp is an example.

The second survivor is symmetry with discontinuity. Symmetry gives sameness of content across a difference of setting; discontinuity gives difference of content within a sameness of setting. Again, on each side, sameness and difference each appear once. This pair describes something that stays the same and comes in separate units. The tiled floor is an example.

The four characters can be laid out as a square. Put four cards on a table in two rows of two. In the top row, put symmetry on the left and continuity on the right. In the bottom row, put asymmetry on the left and discontinuity on the right. Each column now holds a character above its opposite, symmetry above asymmetry, and continuity above discontinuity; those are the pairs of opposites, set aside by the second consideration. Each row holds one of the two pairs that undo themselves: symmetry with continuity along the top, asymmetry with discontinuity along the bottom. The two diagonals hold the survivors. One diagonal runs from continuity, top right, to asymmetry, bottom left: an unbroken order that runs one way, like the ramp, or like the experience of a single life. The other runs from symmetry, top left, to discontinuity, bottom right: the same unit repeated, with seams between, like the tiles, or like the steps of a computation.

The two survivors are closer kin than they look. Swap "same setting" for "different setting" everywhere it appears, and change nothing else. Continuity, the same content in the same setting, becomes the same content in a different setting, which is symmetry. Asymmetry, a different content in a different setting, becomes a different content in the same setting, which is discontinuity. So that one swap carries the first survivor exactly into the second. Now swap "same" for "different" everywhere, for content and for setting alike. Each survivor is carried into itself, and the two pairs that undo themselves are carried into each other. So the two survivors are one pattern of "same" and "different", written once as it is and once with "same setting" and "different setting" exchanged; and the two forbidden pairs are one pattern, written once as it is and once with every "same" and "different" exchanged. Because the two survivors turn into each other by exact swaps, matched pairs of opposite faces should be expected

wherever comparison is at work: two descriptions of one thing, one smooth and going somewhere, the other the same everywhere and countable. Before this chapter ends, it will meet the most famous pair of opposite faces in physics.

The Theorem, Stated

We can now state the theorem whole. Of the ten pairings of the four characters, applied absolutely and together, two are forbidden: symmetry with continuity, and asymmetry with discontinuity. I call each of these a forbidden pairing; my other books call it a forbidden conjunction. Two pairings survive: continuity with asymmetry, and symmetry with discontinuity. In short: symmetry and continuity cannot both be applied, at once and absolutely, to one comparison, and neither can asymmetry and discontinuity. Any comparison whose characters are made perfect, holding exactly at its foundation, is either smooth and going somewhere, or the same everywhere and countable.

In *An Immanent Metaphysics* (first published in 2001), I give the full statement of the Incommensuration Theorem, with the reason for each pairing:

> "When something is absolutely symmetric, it must also be intrinsically discontinuous. When something is absolutely asymmetric, it must also be intrinsically continuous. When something is absolutely continuous, it must also be intrinsically asymmetric. When something is absolutely discontinuous, it must also be intrinsically symmetric. Nothing anywhere at all can be both absolutely symmetric and absolutely continuous. Nothing anywhere at all can be both absolutely asymmetric and absolutely discontinuous".

In plain words: whatever is perfectly the same in every setting must come in separate pieces; whatever is perfectly different in every setting must be smooth; whatever is perfectly smooth must be going somewhere; and whatever is perfectly broken into pieces must be the same in every setting. Nothing anywhere can be both perfectly the same in every setting and perfectly smooth, and nothing anywhere can be both perfectly different in every setting and perfectly broken into pieces. In the quotation, intrinsically means by its own nature; that is a different use of the word from the six intrinsics, the parts of comparison.

The theorem's name, Incommensuration, is chosen with care. Two things are *incommensurate*, in the old sense of the word, when they have no common measure: no single scale on which both can be brought to completion together. Symmetry and continuity are incommensurate in exactly that sense. Each can be increased, and increasing one does not at once reduce the other. But perfect symmetry and perfect continuity do not share a summit. Picture two mountains that rise from the same ground and never meet in one peak: you can climb a long way up either one, and no path leads to the top of both. Two things related in this way are also called *conjugate*: each can be perfected only at the cost of the other, and neither comes before the other.

Chapter 1 used the same word for the position and the momentum of a particle, and the two uses share one meaning: two things are conjugate when neither can be brought to perfection together with the other, and neither comes before the other. The two uses of the word differ in one thing: how far the two conjugate things can both be sharpened before sharpening one begins to cost the other. For position and momentum, sharpening one begins to spread the other once the two are pinned down finely enough, near the floor that Planck's constant marks; symmetry and continuity exclude each other only when both are made perfect.

The two forbidden pairings and the two survivors are the whole of the theorem. The rest of this book follows where physics, and the four things this book asks about (signals faster than light, time travel, simulation and determinism), run into the theorem.

What "Absolutely" Means

One qualifier in the theorem is part of the theorem and is never dropped: *absolutely applied*. A character is applied absolutely when it holds exactly, at the foundation of a description, with no exception at any scale.

Within any limited region of the world, symmetry and continuity occur together, each to some limited degree. A tiled ramp is very nearly seamless when its tiles are small. A field (a quantity with a value at every point of space, like the temperature at each point of a room) can be both smooth and the same from place to place, to a very high degree of approximation, and most of working physics is done at that degree of approximation. The same holds of all four characters: in every limited region, each occurs in some degree. What symmetry and continuity cannot do is both be perfect at once.

The theorem also allows symmetry and continuity together, each one perfect, when the two are kept apart in one of three ways.

The first way of keeping symmetry and continuity apart is by level: symmetry and continuity hold of different things, so that one of the pair holds of the object and the other of the pair holds of the framework in which the object sits. A planet's path around the Sun is smooth, with no jumps, and the law of gravity that bends the path is the same everywhere and at every time. The smoothness belongs to the path, and the sameness belongs to the law.

The second way of keeping symmetry and continuity apart is by phase: symmetry and continuity hold at different times within one procedure, one after the other. Dough is rolled out into one smooth, unbroken sheet, and afterward one cutter stamps it into shapes that are all alike. The smoothness comes first, and the sameness comes after it.

The third way of keeping symmetry and continuity apart is by scope: symmetry and continuity hold over different ranges with a boundary between the ranges, so that one of the pair holds inside the boundary and the other of the pair holds across the boundary. On the tiled floor, the sur-

face of each single tile is smooth, and across the whole floor the tiles repeat. The seams are the bounds: the smoothness holds inside each tile, and the repeating holds from tile to tile.

The theorem forbids only symmetry and continuity perfected together, in one description of one thing. It is a statement about limits: it says what happens when someone asks for total symmetry and total continuity in the same description of the same thing.

The theorem can also be read in terms of parts and wholes, and *An Immanent Metaphysics* gives that reading too: “Symmetry is an assertion about the parts that is made from the perspective of the whole (the macroscopic and/or the mesoscopic). Continuity is an assertion about the whole that is made from the perspective of the parts (the microscopic and/or the mesoscopic)”. The *macroscopic* is the scale of the large, the *microscopic* the scale of the small, and the *mesoscopic* the scale in between; "and/or" means either scale, or both. Stand back at the door of the tiled room, and you can see that each tile is like every other: that is a claim about the parts, made from where the whole is in view. Walk along the ramp a step at a time, and you find no jump anywhere: that is a claim about the whole, made from where the parts are. The two claims are made from opposite ends of the scale.

What Kind of Claim This Is

The theorem is a claim about the structure of comparison itself. It depends only on what comparison is, and on the six intrinsics that every comparison must have, no one of which has meaning without the others. No experiment, no physical theory and no particular kind of mathematics enters it, and it was not learned from measurement.

The Incommensuration Theorem is a structural fact about the nature of knowing. Every act of knowing is an act of comparison, or is built from acts of comparison. The theorem is proved from nothing but the meanings of the parts of comparison, so the theorem, together with the intrinsics of

comparison from which it is built, is inherent in every single act of knowing. It is directly about comparison, the act in which every other result is reached, and so it is the deepest and most powerful epistemic basis possible. Epistemic means concerning knowing: how anything comes to be known, checked and held. No deeper knowledge of knowing can be had, even in principle.

Chapter 1 showed that interaction, comparison and relation are one act, named three ways: the Root Tautology. Each of the three is the other two seen from one side, and every description sits inside the comparison it describes. So whatever is true of comparison as such is true of every physical interaction as such, and the theorem holds of every interaction in physics, before any experiment is run.

The theorem also holds of anyone who tries to doubt it. To reject the theorem, one must compare its account of comparison with some other account. In making that comparison, one uses all six intrinsics: the two accounts are judged the same or different (sameness and difference), in what they say and in the respect in which they are set side by side (content and context), by someone, about something (the subjective and the objective). So the act of rejecting the theorem carries out the very structure being rejected, and no one can stand outside the theorem in order to doubt it. In 1637 René Descartes wrote, "I think, therefore I am": to doubt that you are thinking, you must think. The theorem is like Descartes' claim in one respect: the act of doubting the theorem enacts the theorem.

The Wave and the Particle

The two survivors are the two faces that any complete physical description must show. Within limited regions both are present, in degrees, so a description need not choose one survivor and abandon the other. At the limit, where a description is pushed to perfection, it must be one or the

other, and the two cannot be merged into a single perfect description. The oldest pair of opposite faces in modern physics is made of the two survivors: the wave and the particle.

Consider first what a wave is, before any physics is attached to the word. A wave is a pattern that varies smoothly through some medium (whatever the wave runs through, such as the water of a pond), like a ripple across a pond, so that two neighboring points differ by an amount that shrinks to nothing as the points come together. The smooth variation of a wave from point to point is continuity: the same quantity carried smoothly from one place to the next. A wave also goes somewhere. Its *phase*, how far along its rise and fall it is at a given point, moves on. At any one moment, no two points along the wave have reached the same phase, counting how far along its rise and fall the wave has gone since it started, and at any one point, no two moments have. The difference of the wave's phase from point to point and from moment to moment is asymmetry. A wave is continuity with asymmetry, the first survivor, and it is nothing else.

Try taking the asymmetry away, and see what is left. A wave whose every point was in the same state as every other point would be a constant. It would not travel, it would not overlap with another wave and add to it, and it would not carry anything. A wave with its asymmetry taken away would not be a wave. Now try taking the continuity away instead. A pattern that jumped from value to value, with nothing in between, would have no phase, because phase measures how far along the smooth change one is. It could not be laid over another wave and added to it, because that adding needs a value at every point. It would not be a wave either.

So the wave, as an idea, is the first survivor applied to something physical. The path of a moving body through space and time is another picture of the same survivor: smooth, and directed. The mathematics that describes the first survivor most directly is the mathematics of smooth change, the kind that physicists write as equations of change from moment to moment and place to place. Everything wavelike in physics

(waves adding and cancelling, waves bending around an edge, waves passing through one another, phase itself) expresses continuity with asymmetry.

Now consider what a particle is, again before any physics is attached. A particle is a unit. A particle is here or it is not; there is one particle, or two, or none, and between one and two there is nothing. A particle's coming as one, or two, or none is discontinuity: a definite change within a fixed frame, with nothing in between. A particle is also something that can be exchanged for another of its kind with no change at all. One electron is as good as another. Swap two of them, and no measurement can tell that the swap was made. An exchange of two particles of one kind that no measurement can detect is symmetry. A particle is symmetry with discontinuity, the second survivor, and it is nothing else.

Physics bears out that a particle is symmetry with discontinuity. That particles of a kind can be swapped with no change is one of the deepest facts of quantum theory, the theory of matter and light whose findings come in countable steps. Two particles of one kind are not only alike; they cannot be told apart even in principle. The two rules that govern how crowds of identical particles behave, named for Satyendra Nath Bose and for Enrico Fermi, follow from the requirement that swapping two such particles changes nothing in the description. Bose's rule lets any number of particles of one kind share a single state, as the light in a laser beam does; Fermi's rule lets no two share a state, which is why the electrons of an atom stack up in shells. That particles can be counted is just as basic. The number of particles in a region can be measured. Electric charge comes in whole units. The labels that physics gives a particle come in separate steps, and a change from one label to the next is a jump. The mathematics that describes the second survivor most directly is the mathematics of counting: the conserved charge, the whole-number label, the definite values a measurement can return.

Try taking the symmetry away. A unit that some measurement could tell apart from every other unit would be an individual, and not a particle of a kind. Now try taking the discontinuity away. Something that could be

swapped with no change, but came in no definite number, something that could be half present, would not be a particle at all. It would be more like the strength of a field, which can take any amount. The two together make a particle, and neither alone does. Chapter 2 found that wherever physics asks for perfect symmetry, it finds jumps, and that those jumps are *quantization*, coming in whole units that cannot be split. The particle is quantization given a name.

Why No Single Picture Shows Both

Since the early 1900s physics has known that light, and matter, show both faces. In 1905 Albert Einstein explained how light knocks electrons out of a metal surface: the light arrives in separate units. In the effect named for Arthur Compton, light bouncing off electrons again behaves as a unit: it hands each electron a kick of one definite size, as one ball striking another would. When light passes through narrow openings, it spreads and adds and cancels, as a wave does. Electrons, in turn, spread and add and cancel like waves, and yet each electron is counted, one at a time, as a single dot on a screen. No single ordinary picture holds both faces, and every attempt to make one produces the puzzles of the experiment with two narrow slits. In that experiment, electrons are sent one at a time toward a screen through a wall with two narrow slits in it. Each electron arrives as a single dot, and yet the dots, added up, build the banded pattern that a wave makes when it passes through both slits and adds and cancels on the far side.

In 1928 Niels Bohr gave a name to the fact that light and matter show both faces and that no single ordinary picture holds both: complementarity. Each of the two descriptions is complete where it applies; the two exclude each other; and both are needed. Bohr did not say why the two descriptions exclude each other and why both are needed.

The reason is the theorem. A description that had both faces at once, each one perfect, would be at the same time fully smooth, with a phase defined everywhere, and fully symmetric, swappable and countable. Such a description would contain symmetry with continuity, one of the two pairs that the theorem forbids when both are applied absolutely and together. The wave and the particle cannot be merged into one perfect picture, for the same reason that the two survivors cannot be merged: the wave and the particle are the two survivors.

Both faces are present, in degrees, in every quantum system. A short burst of wave has a roughly definite place and a roughly definite count of units. A particle has a wavelength, as Louis de Broglie found. Quantum theory holds the two faces on its two levels: the wave in its mathematics, and the particle in what its measurements find. So the theorem leaves every quantum system free to show both faces; it never says that a system must be either a wave or a particle.

What the theorem says is that at the limit, where one face is made exact, the other must fade. Make the number of particles exact, and the phase has no definite value. Make the phase exact, and the number has no definite value. In the physics of light this is a known rule, the relation between the number of particles of light in a beam and the beam's phase, written as an inequality: the spread in one can shrink only as the spread in the other grows. On the reading this chapter gives, that the wave and the particle are the two survivors, the rule relating number and phase is the wave and particle case of the theorem, written down in numbers.

The account of the wave and the particle as the two survivors is a reading of known physics. Physics established complementarity on its own evidence, and the reading takes that result as physics established it, without proving it again. I take complementarity as one case of the structure of comparison, a structure that was worked out without complementarity in view, and I claim that the structure of comparison is the reason complementarity holds. Since the wave is the first survivor and the particle the second, the wave and the particle cannot be merged because of what comparison is, and not because of any special feature of quantum

theory. So the same complementarity should be expected in any field rich enough to have both a description that is smooth and going somewhere and a description that is the same everywhere and countable. That expectation says what kind of relation will be found between the two descriptions, and gives no number for how large the trade-off between them will be. It is the only prediction that this reading of the wave and the particle makes.

Every Duality Is a Triple

In every pair of opposite faces there is a third thing, which relates the two. For the wave and the particle, the third thing is the act of measurement. The arrangement of the experiment, the interaction between the thing measured and the instrument, decides which face is shown. Set up the experiment to find which slit the electron went through, and the particle appears. Set up the experiment to see the pattern on the screen, and the wave appears. The two faces are not lying side by side in the electron, waiting to be looked at. They are what the electron shows under one interaction or under another; with no interaction there is no showing, and so there are no faces. Wave: smooth, and going somewhere. Particle: swappable, and countable. The two faces are what one electron shows under two interactions, not two things, and the interaction is the third member of the triple, without which there would be no pair.

The same holds outside physics. A subject and an object are related by perceiving in one direction and by expressing in the other. Every subject is the subject of some perceiving, every object is the object of some perceiving, and the perceiving is neither the subject nor the object. Theory and practice are related by making: the act in which a principle becomes a made thing, which is neither the principle nor the thing. In each case the pair of opposites cannot exist without the relation that makes them a pair. A pair of opposite faces is always three: the two ends, and the comparison that shows them.

Chapter 4 sets out three roles that run through every comparison: the pattern, laid out and the same for everyone; the act, happening here and now; and possibility, the open range within which things can happen. Here we need only where the wave, the particle and the measurement fall among them. The particle, being symmetry, belongs with the pattern. The wave, being continuity, belongs with possibility: before a dot appears on the screen, the wave is the spread of places where a dot could appear, and with what odds, a range of what could be found. The measurement, being interaction, is the act. Chapter 4 sets out three axioms that follow from the theorem, and two of them are needed here, said now in plain words. The first of the three axioms says that the act is more basic than the pattern or the possibility, and that the pattern and the possibility are conjugate: each can be perfected only at the cost of the other. So the wave and the particle are the conjugate pair, and the interaction is the act that relates the wave and the particle and makes them a pair. The third axiom adds that the three are distinct, inseparable and not interchangeable. They are distinct: a wave is not a particle, and a measurement is neither. The wave, the particle and the measurement are inseparable: none of the three occurs without the other two. They are not interchangeable: no one of them can take another's place, and in particular the measurement cannot be made into a thing of the same kind as the two it relates, one more physical thing set beside the wave and the particle, which would then need a further act, a fourth, to relate it to the wave and the particle.

The point about the measurement matters, because the puzzles of quantum theory have usually come from treating the measurement as one more physical thing beside the wave and the particle. Treat the act of measuring as one more physical thing, a wave or a particle among the others, to be described by the same mathematics that the measuring is supposed to decide, and the description runs back without end: something must then measure the measurement, and something else must measure that. The triple says why the regress arises. It comes from trying to swap the act, the third member of the triple, with one of the two faces it relates.

Duality Is Real; Dualism Cannot Stand

Two words are often used as one, and they need to be kept apart. *Duality* is the observation that one comparison has two faces, that the two faces are related by an exact swap, and that they cannot be merged into a single perfect picture. Duality does not claim that there are two kinds of stuff. Duality claims one act of comparison, and two faces that the structure of comparison gives that act. As I put it in *Unlocking Reality*, one of the earlier books of my framework, one becomes two in relation, and the two, along with the one, are three. The act of comparison and its two survivors make the most basic triple there is: the act, and its two faces.

Dualism is the claim that there are two independent things. Mind and body. Form and matter. Freedom and necessity. Wave and particle, when they are taken as two separate stuffs that a system somehow is at once. Every dualism meets the same question at once: how do the two interact? Any interaction is a relation, and a relation between two things means that the two were never independent. So two things that are wholly independent can have no answer to that question. Any answer names a third thing, the interaction, and the interaction is more basic than the two, because the interaction is the relation that makes the two a pair. So a dualism cannot stay a dualism. Either it denies the interaction, and becomes two separate accounts that never speak to each other, or it admits the interaction, and becomes a triple.

Applied to physics, the distinction says this. Wave and particle duality is real. The two faces are exactly the two survivors, the swap that relates them is exact, and the theorem says why they cannot be merged. Wave and particle dualism cannot stand. Any account that takes the wave and the particle as two things a system independently is, and then has to explain how the two are reconciled, has already left out the interaction that relates them and will have to bring the interaction back under another name.

The triple can be found inside each way of interpreting quantum theory that starts from the faces, whether from both of them together or from the wave alone. In 1952 David Bohm described a wave and a particle together: a real particle that always has a definite place, and a real wave that guides the particle. In Bohm's account, the guiding is the third member of the triple. Other accounts keep the wave alone, and take the smooth wave to be the whole of what there is; they must then explain why we see units, and that explanation is an account of measurement, and measurement is the third. The triple reads each kind of account as confirming the triple, and the triple refutes neither kind of account: whichever face one starts from, the interaction has to be supplied.

Where Physics Meets the Theorem

Physics meets the theorem in many places, and three of them can be named now, as three faces of one theorem.

The first face is the leftover between a count and a curve. Whenever a count, a whole number of units, is read against a smooth whole, something is left over at every such reading, and the leftover can never be zero, because of the Incommensuration Theorem. The tiles laid along the ramp show that leftover: however small the tiles, there is always a seam where tile and curve do not match.

The second face is the floor under paired descriptions. Where a particle is, and how much motion it carries (its *momentum*, which grows with both its mass and its speed), are two conjugate descriptions of one state. A state placed sharply in position is spread out in momentum, and a state sharp in momentum is spread out in position, and no description makes both sharp. That no description makes position and momentum both sharp is the uncertainty principle of quantum theory. I state the uncertainty principle as an instance of the Incommensuration Theorem.

The third face is the price of being inside empty space. Picture someone trying to weigh the air of a sealed room while standing inside it. The scale sits in the same air, and that air presses on it from every side at once, so part of what the air does never shows as a reading. Anyone who measures the energy of empty space is inside the empty space being measured, in the same way. I identify the fact that an observer inside empty space cannot reach part of its energy as an instance of the theorem at the scale of the whole universe, and I identify the cosmological constant (the number in the theory of gravity that measures how strongly empty space itself drives the expansion of the universe) with the part of that energy which the observer can reach. This face matters for the book because it is the theorem met at the largest scale there is, by an observer who can never stand outside what is measured; Chapter 14 comes back to it, and says at which horizon the price is paid.

One theorem, three faces: the leftover between count and curve, the floor under paired descriptions, and the price of being inside empty space.

The two great theories of physics are the two survivors themselves. Comparison had its structure before any measurement was ever made: six intrinsics in three pairs, no one of which has meaning without the others; four characters built from the intrinsics; and one theorem about which characters can be perfected together. Continuity with asymmetry can be perfected, and general relativity, the theory of gravity as the shape of space and time, is what that looks like. Symmetry with discontinuity can be perfected, and quantum theory is what that looks like. General relativity takes the mathematics of continuity with asymmetry, which belongs with possibility, and uses that mathematics on the world of the large, which belongs with the pattern. Quantum theory takes the mathematics of symmetry with discontinuity, which belongs with the pattern, and uses that mathematics on the world of the small, which belongs with possibility. Each theory had to reach across in this way; Chapter 17 shows why, and why the two theories can be joined side by side and never merged into one.

Two results from the history of physics will return in later chapters. Chapter 2 took up the result Emmy Noether published in 1918, a law about symmetry: each continuous symmetry of a law goes with an amount that is conserved. Chapter 10 takes up the result John Bell published in 1964: no account in which every outcome is fixed by hidden local causes can match all the predictions of quantum theory, and the measurements that confirm them. Chapter 10 reads Bell's result as one more place where physics runs into the theorem.

Two numbers of physics mark the edges of comparison. Planck's constant marks the small edge of the range in which comparisons can be made, and the speed of light marks the large edge. Each is more basically a limit of measurement than it is a size or a speed. Every basic measure of every domain has a lowest possible value and no highest one; that fact follows from the three axioms of Chapter 4 as a consequence, so it needs no fourth axiom. The speed of light is an edge of a different kind from a highest value. The measures inside the domain, a distance, a duration, an amount of energy, have a floor and no top: there is always room for more. The speed of light marks where the domain itself ends. Past that edge lies no larger value of anything, only events about which nothing can be asked from here, because no signal from them arrives. So the large edge is the boundary of where comparing can reach, and not the top of any scale measured inside it.

Those limits of measurement are limits of comparison, because a measurement is comparison read twice over: comparison described as an interaction, and that interaction described as a measurement. As Chapter 1 said, a physics that claims to reach beyond every possible measurement has stopped doing physics, without thereby doing metaphysics.

This way of reading physics says where measurements will run out, and why, and what shape physics must have at the places where they do. What to measure, and what a measurement will show, stay with physics; this way of reading physics does not compete with it. For a century physics has been finding those places one at a time, each one a surprise. The places where measurements run out are the edge of comparison, seen

from inside, so they need not have surprised anyone, and a physics that knows they are the edge of comparison loses nothing, and gains knowledge of where its own boundaries lie.

Why This Limit Matters

We now know the one limit that every later chapter uses. A description can be perfectly smooth and going somewhere, or perfectly the same everywhere and countable. No description can be perfectly the same everywhere and perfectly smooth at once, and none can be perfectly different everywhere and perfectly broken into pieces at once.

Each thing that this book shows cannot be is ruled out by this same limit, met in a different place. Look again at the top row of the four cards: symmetry with continuity. A world run by one law, the same everywhere and at every time, in which every state passes into the next with no gap at all, asks for that top row made perfect; Chapter 5 shows that the claim that everything is fixed in advance requires such a world. The bottom row, asymmetry with discontinuity, made perfect, would describe a content found different whatever its setting, which is pure chance; Chapter 5 shows why pure chance is excluded in the same way as the top row made perfect. The diagonals are the two pairs that can hold in the world when made perfect: the smooth order that runs one way, like a single life, and the repeated unit with seams between, like the steps of a computation. A signal that outruns light, a trip back to last Thursday, a world that is someone's program: each of these, in its own chapter, will turn out to be ruled out by this same limit, met in a different place.

Summary

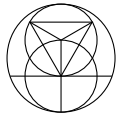
Four characters come from applying sameness and difference to content and context: symmetry, continuity, asymmetry and discontinuity. They can be paired in ten ways. Four pairings of a character with itself say noth-

ing new; two pairings of a character with its opposite contradict themselves. Of the four that remain, two undo themselves when made perfect: symmetry with continuity, and asymmetry with discontinuity. In each, the content no longer depends on its setting, and the line between content and setting, which every comparison needs, collapses. Two survive: continuity with asymmetry, smooth and going somewhere, and symmetry with discontinuity, the same everywhere and countable.

The two forbidden pairings and the two survivors are the Incommensuration Theorem. The Incommensuration Theorem is a structural fact about the nature of knowing. It depends only on what comparison is, it is inherent in every act of knowing and every physical interaction, and doubting it enacts it. Its scope is exact: it forbids symmetry with continuity, and asymmetry with discontinuity, when both halves of the pair are applied absolutely, holding exactly at the foundation of one description of one thing. In every limited region all four characters occur together in degrees, and symmetry and continuity may each be perfect when they are held apart by level, by phase or by scope. The wave is the first survivor, smooth and going somewhere; the particle is the second, swappable and countable; complementarity is the fact that the theorem forbids merging the wave and the particle into one perfect picture; and every duality is a triple, two faces and the act of comparison that shows them.

I make three kinds of claim, and each of the three kinds appears in this chapter. The theorem, its two survivors, and the fact that every duality is a triple are claims about the structure of comparison itself. The wave as the first survivor and the particle as the second, general relativity and quantum theory as the two survivors, and the two constants as the edges of comparison are readings of known physics as instances of that structure, with the structure as the reason the physics holds. The exclusions, the things that cannot be, are consequences of the structure of comparison that reach into physics, and each is argued in the chapter that takes it up.

The theorem is about comparison. Who or what is doing the comparing, and in what order do things happen?



Chapter 4

Process, Pattern and Possibility

Things or Happenings?

Is the world made of things, or of happenings? Most of us were given the first answer without ever being asked the question. There are things: stones, chairs, planets, particles. Things sit in space and last through time, and now and then something happens to them. On that picture the things come first, and a happening is something a thing undergoes.

The question has a famous sharp form. Is the moon there when no one is looking at it? If the world is made of things, the answer is yes, and the question is idle. If the world is made of happenings, the question seems to have teeth, since a moon that no one meets might seem to be a moon that nothing happens to. By the end of this chapter the answer will be yes, for a reason that may be new to you: the moon is never without happenings.

The earlier chapters have given us the ground to stand on. Chapter 1 showed that every measurement is a comparison, the act of finding whether two things are the same or different in some respect. Chapter 1 showed as well that to compare is to interact, and to interact is to compare, and that relation is both comparison and interaction. Interaction, comparison and relation are therefore one act named three ways, which I call the Root Tautology. Chapter 1 also showed that there is no such thing as observing nothing. A detector that registers no event has still registered: the

absence of a click is an outcome that the apparatus was built to be able to produce, and to be an observer at all is to be already in interaction. Which comes first: the act of comparing, or the things that are compared?

Three Ways of Looking at a Game

Watch two people playing chess at a kitchen table. There are three ways to describe what you see.

First, there is the move being made. A hand lifts a knight and sets it down, here, now, between these two players. The move being made is the game as the game happens.

Second, there is the game as it can be written down. There are the rules, which hold for every game of chess ever played; there is the position on the board, which anyone who walked into the kitchen could read; and there is the score sheet, the list of moves so far, which anyone could replay. Whoever looks, these read the same.

Third, there is what the position makes possible: the moves that could be played next, and among them the move that neither player has yet thought of. What the position makes possible is not on the board. What the position makes possible is what lets the next move be new.

No full account of the game leaves out any of the three. A board and its rules with no move being made is not yet a game in play. Moves made with no rules are not chess. A position from which only one move could ever follow would leave the players nothing to do.

The three ways of looking have names. When everything at the foundation of a domain, any field in which comparisons are made, can be gathered under three ideas, each of them needed and the three of them enough, those three ideas are the domain's *modalities*, its three basic ways of being taken.

The *immanent* is the mode of process: the relational, the interactive, the here and now, what is happening between things. The word means what is present within, and here it names what is going on. The knight being set down is immanent.

The *omniscient* is the mode of structure: the repeatable, the shareable, the third person, what holds regardless of who looks. The word means all-knowing, and it fits, because structure is what every knower would find the same. The rules, the position and the score sheet are omniscient.

The *transcendent* is the mode of the generative: the evaluative (the weighing of one possibility against another, as a player weighs the moves), the creative, the potential, what makes a domain possible rather than what is found within it. The word means going beyond, and possibility goes beyond anything already found. The range of moves still to be made, and the new move among them, are transcendent.

The same three show up again and again. Interaction is immanent; existence is omniscient; creation is transcendent. Change is immanent; causality is omniscient; choice is transcendent. At the chess table, the move changes the position; the rules tie each position to the positions that can follow it; and the player chooses which of those positions comes next. The pattern recurs at every level, and each such triple lines up with the three modalities in a way that can be checked. The third axiom, below, supplies the checks. A later chapter takes up the triple of change, causality and choice in full.

The Theorem Comes First

Chapter 3 established the Incommensuration Theorem. It says that symmetry and continuity cannot both be applied, at once and absolutely, to one comparison, and that neither can asymmetry and discontinuity. Symmetry is sameness of content across a change of setting; continuity is

sameness of content with the setting held fixed, so that nothing jumps; asymmetry and discontinuity are the matching differences. The Incommensuration Theorem is a structural fact about the nature of knowing.

The theorem comes first. It is proved from comparison alone: from the six intrinsics (sameness and difference, content and context, the subjective and the objective), from the fact that no comparison can do without any one of them, and from the four compounds that Chapter 2 built out of them. The three axioms are derived from the theorem, as its inner structure. An *axiom* is a basic statement from which other statements are reasoned; here each of the three is itself drawn from the theorem. I show two of the links, the first axiom's conjugacy and the one-way direction of the second; the full derivation is the work of the foundation volume *Incommensuration; The Theorem and the Three Axioms*, and I state each axiom that the derivation yields. I set them out in the wording of my book *An Immanent Metaphysics*, first published in 2001, because the rest of this book states its account of law and of the modalities in their terms.

One link between the theorem and the axioms can be seen at once. In the alignment this book uses, symmetry belongs to the omniscient, since what holds the same for every knower is what stays the same across a change of setting; and continuity belongs to the transcendent. The first axiom says that the omniscient and the transcendent are conjugate: each can be perfected only at the cost of the other, and neither comes before the other. The conjugacy of the omniscient and the transcendent is the theorem itself, stated in the names of the roles. The second axiom, below, sets out a cycle that runs one way, and that one-way direction is forced by the same conjugacy, as the account of the second axiom below shows.

The Three Axioms

The first axiom reads: "The immanent is more fundamental than the omniscient and/or the transcendent. The omniscient and the transcendent are in conjugate relationship to one another". Read aloud, "and/or"

means: than either one of them, and than both of them together. In plain words: process is more basic than structure, and more basic than possibility. Interaction comes before the things that interact, and before the range of what those things might become. By the Root Tautology, interaction and comparison are one act, so the first axiom also says that comparison comes before what is compared. Its second sentence says that structure and possibility are a conjugate pair: each can be made perfect only at the cost of the other.

The second axiom reads: "A class of the transcendent will precede an instance of the immanent. A class of the immanent will precede an instance of the omniscient. A class of the omniscient will precede an instance of the transcendent". A *class* is a kind, taken as a whole: possibilities as a kind, interactions as a kind. An *instance* is one particular case of a kind: this interaction, this pattern. In plain words: a possibility, taken as a kind, gives rise to a particular interaction; interactions, taken as a kind, give rise to a particular structure; and structures, taken as a kind, give rise to a particular possibility. The order in the second axiom is the *directed cycle*. It runs from structure to possibility, from possibility to interaction, and from interaction back to structure, and its direction cannot be reversed.

Picture a lawn between a bus stop and a bakery. People cross it every morning, each by a slightly different line. Their crossings, taken together, wear one particular path into the grass: interactions, as a kind, have given rise to a structure. The path, once it is there, makes certain crossings possible and easy, including a shortcut that no one has taken yet: a structure, as a kind, has given rise to a particular possibility. Tomorrow morning someone takes the shortcut: a possibility has given rise to a particular interaction. And the cycle runs one way only: the crossings wear the path, and the path never makes the crossings that wore it.

Why one way only? Because structure and possibility are conjugate, they cannot both be fully present at once, so any process that involves both must pass through them in turn, never both together. And both are made from the act of comparing, so the two meet only through that act.

From structure comes a possibility: the worn path opens the shortcut. From a possibility comes an act: someone takes the shortcut. From the act comes new structure: the shortcut, walked, becomes part of the path. To run the other way, a possibility would have to give rise to structure directly, with no act between them, which would be a stance with no engagement in anything, and so empty. The foundation volume *Incommensuration* derives this direction from the theorem in full.

A laboratory runs the same cycle. The apparatus, set up just so, is a structure, and it gives rise to a particular range of possible outcomes. Out of that range comes one particular interaction: a click. The click is laid down as a record, which is structure again; and the records, taken together, set up the range of possibilities for the next run. The record follows the click it records, always. Every process in every domain runs this cycle, and the cycle is what the word process means when the word is used exactly.

The third axiom reads: "The (classes/instances) of the immanent, omniscient, and transcendent are distinct, inseparable, and non-interchangeable". Read aloud, the part in parentheses means: the classes of the three, and likewise the instances of the three. In plain words, three things hold of the three modalities. No one of the three can be reduced to the others: they are distinct. No one of the three can occur without the others: they are inseparable. And no one of the three can be put in the place of another without the substitution being a mistake: they are not interchangeable. This third axiom is the axiom that makes a triple a triple. It supplies the three tests that every foundational triple in this book is checked against. Among those triples are relation, domain and identity; interaction, existence and creation; choice, change and causality; and reality, existence and objectivity. At the chess table, taking the score sheet for the game being played would be a substitution of exactly this forbidden kind.

The axioms describe their own domain. The first axiom is a statement about what holds regardless of position, and so it has the nature of the omniscient. The second axiom is a statement about process, and so the second axiom has the nature of the immanent. The third axiom is a statement about what makes the three modalities possible at all, and so the

third axiom has the nature of the transcendent. That the axioms fall into the pattern they describe does real work: it is what lets the *metaphysics*, the study of the structure of comparison, ground itself without appeal to anything outside itself. Describe the axioms by means of the axioms, and then describe that description in the same way: at that second step the description closes on itself, and a third step adds nothing new, because describing the second step needs no concept beyond the ones the axioms already have, so a third step would only repeat the second.

One further principle is needed before the theorem can be read exactly, and it follows from the axioms as a consequence, so the axioms stay three. Every domain is *zero-positive*: its fundamental measures have an absolute minimum and no absolute maximum. Each kind of quantity a domain measures has a floor that nothing goes below, and no ceiling. The speed of light is no ceiling of this kind: as Chapter 3 showed, it marks where the domain itself ends, and not the top of any scale measured inside it. The qualifier in the theorem, absolutely applied, has its own meaning, which Chapter 3 gave: holding exactly, at the foundation of a description; the Appendix says how that qualifier stands to the principle that every domain is zero-positive. Symmetry and continuity appear together approximately, as they do in every real system, and most of working physics lives there. What the theorem forbids is the two holding exactly in what one and the same comparison finds, and likewise asymmetry with discontinuity. The two survivors of Chapter 3 can hold exactly: continuity with asymmetry, smooth and going somewhere, the character of a wave; and symmetry with discontinuity, the same everywhere and countable, the character of a particle. The two forbidden pairs cannot.

So the three axioms are three faces of one structural result. The ground of this book rests on a theorem alone, and the axioms are its inner structure, not additions to it.

With the three axioms in hand, the completeness of the ground that Chapter 1 promised can be stated. The three axioms and the three ways of looking are closed and consistent under reason (reasoning from them never leads outside them, and never to a contradiction), in process and in

theory, in several overlapping and intersecting ways, and they cover the basic notions of comparison, description, process and interaction completely. In physics, an interaction goes by the name particle: the dot on a screen where an electron is found, one at a time, is the place where an interaction happened. The basic notions of comparison, description, process and interaction also appear in physics under the names intersection (the place where two lines of happening cross, as two roads cross at a junction), field and measurement; all of these words are different words for the same patterns. All of mathematics, computer science and physics, and all of reason and argument itself, lies within the scope of this ground, its terms and its process, and nothing the account needs lies outside it.

Interaction Before Existence

How do you know that there is a chair in the next room? You could walk in and sit on the chair, or knock against the chair in the dark and hear the chair scrape. A friend could trip over it and tell you so.

Ask what has to be assumed before anyone can say that something exists, or that it is real, or that it is objective. In each case the answer is the same: some possibility of engaging with it, in a way that can be confirmed. If I can interact with it, if I can measure it, if there is some way in which it could make a difference to me and I to it, then I can say that it exists, or that it is real, or that it is objective. In every one of these cases the idea of interaction is taken first, in order to establish the idea in question. Taking the idea of interaction first, in each of these cases, is the first axiom, met in practice.

The three words sort by the modalities. Interaction is immanent; existence is omniscient; objectivity is transcendent; and reality, the plural of what is real, is immanent, because what is real is what is met in interaction. Reality here means the many real encounters, each one real, and not one total made by adding them up.

To exist is to have a place in a pattern: to stand as structure, the same for every comparer; to be repeatable, locatable and shareable, the kind of thing a third person could be told about. *To be real* is to be encountered, here and now, in this interaction: to be in an actual act of comparison. *To be objective* is to hold across perspectives, to be what any observer would find. A claim that something is objective is a claim about the whole range of possible observers, and so a claim about possibility, as the score sheet and the judgment of a lost position showed in Chapter 2. Objectivity also comes in degrees: it is the agreement of many findings across many comparers, and no single act of comparison has it on its own; a single act is real, and only real. At the chess table, the game being played tonight is real; the rules and the score sheet exist; and the judgment that the position is lost for white, which any strong player would reach, is objective.

The three are distinct, inseparable and not interchangeable, as the third axiom requires. Because they are distinct, one of the three words can be said of a thing where another cannot. There are things that are real which do not exist: encountered, but with no place in any pattern, as the click of a detector is real in the moment it happens, before any record gives it a place. And there are things said to exist which have a place in a *formalism*, a theory's mathematics written out as rules, but which have never been met, such as the many worlds that no one could ever observe, which I name in a later section of this chapter.

It follows that the process of perception is more basic than either the perceiver or the perceived. This can be seen from the inside. I can perceive what is the object of my perceiving. I cannot perceive the perceiver: when I try, what I find is another object. I cannot perceive the perceiving: it is not among the things perceived. Of the perceiver, the perceived and the perceiving, perception reaches only the perceived.

Light does not perceive light. Two beams of light cross in a region of space, and their waves pass through one another with nothing registered. The interaction is not itself one of the things interacting. The interaction is

its own kind of being, just as the relation between "is" and "is not" is neither something that is nor something that is not, and the relation between "is" and "is not" must be of a deeper order than either.

A perception can be pictured as a channel. There is a world that produces a signal, the signal travels through a medium, and a self is on the receiving end. What flows through the channel is taken to stand for what is perceived. In *An Immanent Metaphysics* I distinguish two directions of interaction. In *perception* the change flows from the objective toward the subjective; the work done is abstraction, and overall, details are removed. In *expression* the change flows from the subjective toward the objective; the work done is instruction, which means putting into structure, and overall, details are added.

A measurement is a perception in this sense: it abstracts. A thermometer reads twenty degrees, and that one number is what is left of the motions of countless molecules of air, each with its own speed and direction, all of it set aside by the instrument on the way to the reading. The pointer of any instrument works this way: its reading is what remains of an enormous quantity of detail that the apparatus has discarded. Discarding detail belongs to what measurement is, so no better instrument will stop measurement from discarding it: measuring is taking a content out of a context, and leaving the context behind.

Setting up an experiment is the other direction: it is an expression. The experimenter puts structure into the world, and only then reads back what the world returns. Every experiment is an expression followed by a perception, and the pairing of the two is what makes it a test.

A Law Exists; a Measurement Is Real

Chapter 1 noted that the word relation has two uses, the relating itself and a relation written down as structure, and I now show how the two uses fit together. A relation in mathematics exists as structure. It is a sameness or a difference of content held across contexts: timeless, repeatable, readable

again and again, defined and laid out. It has the nature of the omniscient. An interaction is an event. It is here and now, it cannot be undone, and in its particulars it cannot be repeated. It has the nature of the immanent. So a relation in mathematics is not an interaction.

A relation in the full sense is real. It is given in an interaction between the things it relates, and, as the Root Tautology holds, it is one act with that interaction and with the comparison. A relation in the full sense is the relation that the Root Tautology names. A relation in mathematics is the relating of a relation in the full sense, laid out as structure, as a score sheet lays out a game. The structure laid out and the act that is real are distinct, by the third axiom, and the first axiom says which of the two is more basic: the act. A law exists; a measurement is real. The score sheet of a chess game can be replayed by anyone, at any time, forever. The game itself happened once, on one evening, between two people.

To take the structure that exists for the act that is real is to take mathematics for physics. Mathematics is, and cannot fail to be, founded on patterns of relation laid out as structure. Physics is, and cannot fail to be, founded on patterns of interaction. That mathematics is useful in physics does not make the two the same, either in what we can know or in what there is.

Two bonds are tight. There is no physics without measurement, and there is no mathematics without relations among its objects. The bond between a structure laid out and the interaction it describes is strictly weaker than either of these. Structures model interactions; they do not become them. A model of an interaction is a relation among the interaction's abstractions, and the abstracting has removed the very thing that made the interaction an event.

Chapter 1 drew the consequence for how far physics reaches, and it can now be seen from the axioms. Since physics is founded on measurement, physics cannot be taken to go beyond the ideas that measurement needs. As Chapter 1 said, a theory that claims to describe something bey-

ond all possible measurement has stopped doing physics, without thereby doing metaphysics, and what no comparison could ever reach can be the ground of no assertion.

A claim about what lies beyond measurement in principle is judged by standards of its own, and being judged by them is no rebuke to it. Those standards are needed later in this book, because hidden variables, the third example below, return in Chapters 10 and 11. A claim about what lies beyond measurement in principle is judged by whether it is consistent; by whether it respects the categories it uses; by what it makes understandable; and by whether it has quietly borrowed the authority of physics for a position that physics cannot reach. Three examples show the kind. One is a multitude of worlds that no one could ever observe, taken as literally existing. Another is the wavefunction, the smooth mathematical wave by which quantum theory describes a system and sets the odds of each outcome, taken as a thing in itself, prior to and independent of every interaction. The third example is a set of *hidden variables*, values held to fix every outcome in advance, taken to be out of reach of any measurement in principle. Each of these is a claim about what no comparison could reach. None of the three claims (the many worlds, the wavefunction taken as a thing in itself, and hidden variables out of reach of every measurement) is physics, and existence with no possible interaction cannot be told apart from not existing at all.

Two Theories of Access

I read the two great theories of physics as they stand, leaving each of them unchanged, and I say what kind of theory each one is, in the terms that this book uses.

General relativity, the theory of gravity as the shape of space and time, is a theory of the limits of the flow of signals. It is a model of who can signal whom: from which *where-when* to which where-when, a where-when being an event, one place at one moment. Its first object is the light cone:

at each event, the shape that marks out every event that event can reach and every event that can reach that event. Chapter 12 takes the light cone up in full. The theory's measure of the shape of space and time, which physicists call the metric, is, before it is anything else, a statement of which events can be connected by an interaction and which cannot. General relativity is by its nature a theory of large-scale interactions, and it views them from the outside in: the whole of space and time is laid out, and the observer surveys it as a structure of possible connections.

Quantum theory, the theory of matter and light whose findings come in countable steps, is a theory of the limits of what can be known. It is a model of who can know what, and when. Its basic rules, which physicists call commutation relations, say which pairs of quantities can be definite together; where a particle is and how it is moving, for example, are a pair that no description makes perfectly definite at once. Its state is a catalog of what an interaction could yield, and with what weight each outcome comes. Quantum theory is by its nature a theory of small-scale interactions, and it views them from the inside out: the observer is within the domain, receiving what the interaction delivers, and the theory says what can be delivered.

In the alignment this book uses, symmetry is omniscient and continuity is transcendent, and each of the two theories applies one of these to the domain that belongs to the other. The flow of signals is a matter of connectedness: whether one event can be reached from another without crossing a boundary. Connectedness, the reaching of one event from another without crossing a boundary, is continuity. So general relativity is continuity with asymmetry, applied to the large scale, which is the omniscient domain, the whole laid out as if seen from a balcony, as Chapter 2 pictured it. What can be known is a matter of what stays the same, and can therefore be counted: quantities that every observer finds the same, the whole-number labels that mark a particle's state, and particles such as two electrons, which can be swapped with no difference at all. What stays the same, and can therefore be counted, is symmetry. So quantum theory is symmetry with discontinuity, applied to the small scale, which is the

transcendent domain, the single event that is only itself. The two theories are the two survivors of the Incommensuration Theorem, the two pairs that can hold exactly, and each reaches across into the domain of the other. The reach of each of the two theories into the domain of the other is required, and no accident, for a reason that I give in Chapter 17.

Both theories are *access controls*: rules that say who can reach what. Think of the permissions on a shared computer, a list that says who may open which folder. Each of the two theories states a limit of that kind on what can be perceived or measured: general relativity on what can be signaled between where-whens, quantum theory on what can be known of a what. And both are stated in a formal, deterministic way, deterministic meaning fixed in every detail by what came before. The equations of general relativity carry the shape of space and time forward with no freedom in it; the equation of quantum theory carries its wave forward in the same way. Yet what these theories set limits on need not itself be fixed in advance. A deterministic description of the limits on what can be known does not show that what is known was fixed beforehand. The list of permissions can be exact and unchanging, and it says nothing about whether the letters inside the folders were written in advance. The map of the access controls says nothing about whether what is accessed is determined.

There is an irony here, which the next chapter develops. Purely deterministic methods, abstract and symbolic, can prove formally that in actual concrete reality there is no consistent way of knowing the world as fully determined. The formal apparatus by which we describe the world guarantees that the world is not fully describable by that apparatus. That the world is not fully describable by the formal apparatus is a structural feature of how a formalism stands to reality, and of how a relation stands to an interaction. So the limit on describing the world is no failure of physics. The limit on describing the world is exactly what we were led to expect in the previous section, since structures model interactions and never become interactions.

The Screen and What Lies Behind It

The first axiom says what comes first: interaction. A family of theories puts something else first, and the contradiction those theories fall into shows the first axiom at work.

Think of the screen of a phone. The icons you tap are shaped for your convenience; they do not show what is going on inside the machine. The interface theory of perception, which the cognitive scientist Donald Hoffman set out in his 2019 book *The Case Against Reality*, is the clearest member of the family. The theory holds three things together. Everything we perceive comes to us through an interface, like the icons. No interface is privileged, since each is shaped for the needs of the perceiver rather than for the truth of the perceived. And yet there is a fundamental reality which is privileged, and the interfaces conceal it. The same structure appears in physics whenever someone says that the true reality is some formal object, a piece of mathematics, and that measurement is only our species' limited view of it.

The structure holds a contradiction, and it can be stated exactly. No interface is privileged, and a level of reality is privileged; so somewhere along the way there must be a translation from the privileged level to the interface. Ask of that translation: is it real, or is it not?

Take the first answer: the translation is real. Then it is an interaction, something that actually happens between the privileged level and the interface, and it is then privileged in exactly the sense the theory denied to every interface: it is what actually occurs. The theory's own account of how the interface comes to be has contradicted its claim that nothing at the interface is privileged.

Take the second answer: the translation is not real. Then the privileged level has no connection to any interface, and the claim that it lies beneath the interfaces cannot even be stated, because lying beneath is exactly the connection that has been denied. A privileged reality with no real relation to what is perceived is not a reality being perceived through an interface. It is a supposition with no route to the one supposing it.

Either way the theory holds an inconsistency within itself, which must be resolved before it can be taken seriously as a foundation for physics or for *epistemology*, the study of how we know. The account of this book resolves it by rejecting the premise that both answers share. What is privileged is the interaction, and no level is. There is no fundamental reality beneath the interface, and there is no merely unprivileged interface. There is interaction, which is real, and existence and objectivity are abstractions from it. The first axiom says this directly: "The immanent is more fundamental than the omniscient and/or the transcendent". A theory that makes some laid-out structure fundamental and makes interaction secondary has turned the first axiom upside down, and the contradiction over whether the translation is real is where that upside-down order shows itself.

The Cat in the Box

In 1935 Erwin Schrödinger, one of the founders of quantum theory, proposed a thought experiment. That year was the year in which Niels Bohr published his answer to the question whether the quantum description of physical reality can be considered complete, and the two belong together: both ask whether the mathematics of quantum theory says everything there is to say about a system. Schrödinger placed a cat in a sealed chamber, together with a small quantity of radioactive substance, a counter, and a flask of poison that the counter will release if it registers the decay of an atom. After an hour, the mathematics of quantum theory assigns the whole system a *superposition*: a combination of alternatives held together in one description, here the atom decayed and not decayed, and so the cat dead and alive. Schrödinger meant the picture as an absurdity: a reason not to take the superposition literally at the scale of a cat.

Two answers have become standard. The first standard answer is the disturbance reading: the cat is either dead or alive, and we do not know which. On the disturbance reading, looking at the cat either finds out

whether the cat is dead or alive, or forces the cat to be one or the other. The second standard answer is *decoherence*: the cat is already in constant interaction with its surroundings, and the surroundings act as a measurer. On the decoherence answer, the superposition is destroyed long before anyone opens the box. I accept decoherence as a description of the mechanism, and I reject the disturbance reading as a description of what is at stake. I also give a reason for both answers, which neither answer gives for itself.

The reason begins with the axioms, which forbid perfect closure, that is, a system with no interaction at all with anything outside it. The first axiom puts interaction before existence, and so existence with no possible interaction cannot be told apart from not existing. A domain with no relations to anything outside it would therefore have no existence relative to anything outside it. What the science of heat calls a closed or isolated system is an idealization: a system whose exchanges with its surroundings are small enough to neglect. That no system is perfectly closed shows the role of observation in physical theory. *Observation* is any interaction that relates a system to what is outside it. Observation is not something added from outside to a system that would be complete without it. It is the condition for the system to have a definite state relative to what is outside it.

Being definite is a relation. To have a definite state is to be a content within a context, and a content within a context is established by comparison, which is interaction. With no interaction relating the cat to the outside, there is no fact of the matter for the outside. The reason is not that a fact is hidden, since that would assume the fact. The reason is not that a fact would be disturbed by looking, since that would assume it too. The reason is that a fact for the outside names a relation between the cat and the outside, and that relation has not yet been made. Observation is required by the structure of things for anything to be definite, in the way that a difference is required for a comparison. Where there has been no interaction, there is nothing that the word definite could be applied to.

The disturbance reading assumes a definite value already there, which the measurement reveals or knocks about. The assumption of a definite value already there is exactly what I reject. It is also the assumption that John Bell showed in 1964 cannot be held together with locality, the rule that what happens at one place is affected only through what is near it. The account here rejects the assumption from above, from the idea of comparison; Bell rejects it from below, from the correlations measured between distant particles (the matching patterns in their results). The two routes meet. Chapter 10 tells Bell's result in full.

Inside the chamber, the cat is not in doubt. The atom, the counter, the flask and the cat are in interaction with one another the whole time, and relative to those interactions, things are definite in there. Being definite is relative to the interaction in which it is established, and the interaction inside the box is not the interaction across its wall. None of this makes facts a matter of opinion. Interaction is real, and the first axiom makes it the most basic thing there is. A fact is a content in a context, and the context is supplied by an interaction, never by a view from nowhere.

On this reading, decoherence is the observation that was always already happening. The surroundings are nothing special: they are a partner in interaction, and any interaction is a comparison. So the account expects large superpositions to be found fragile, for a structural reason. A large system is one that is in a great many interactions, and each of them makes relations that a superposition requires to be absent. That fragility is an expectation of the reading, about the kind of thing that happens; the calculation of how fragile a given superposition is belongs to physics.

The quantum state, on this account, exists as the description of possibility: the range of outcomes that a given setup could yield, each with its weight. So a quantum state before measurement is not a particular outcome. Measurement does not reveal a particular result that was already there; it makes one outcome actual, out of the range of possible outcomes. On this account, what is called the *collapse of the wavefunction* is the measurement making one outcome actual, out of the range of possible outcomes. The wavefunction, met in the last section, is the smooth wave by

which quantum theory describes a system. Collapse is not a physical process that destroys a superposition. It is the passage from the possible to the actual in the act of comparison, after which the outcome is a new finding, and the earlier state is replaced by it. A measurement in a laboratory is the second axiom at work: a range of possibilities gives rise to one particular interaction, and that interaction is laid down as a record. The reading says what being definite is, and it leaves physics to judge between the versions of quantum theory that write a collapse into their mathematics and the versions that do not.

Now the moon. Does this mean that things are there only when we look at them? Is the moon there only because someone is looking at it? No. The moon interacts all the time, with sunlight and with everything near it, and relative to those interactions the moon is definite. No person needs to look at the moon for the moon to be definite. Consciousness is not what makes anything definite; any interaction will do, and a Geiger counter, the instrument that clicks when a particle of radiation passes through it, is as good a measurer as a physicist. The reading says what being definite is, and it says that the puzzle of the cat comes from asking for definiteness where the relation that would make it has not yet been made.

Where a Stance Stands

Every comparison has two sides that Chapter 1 named: the subjective, the stance from which the comparison is made, and the objective, what the comparison is about, whoever makes it. The two are inseparable in every comparison, and they are distinct, and the theorem gives each its character.

The objective, what is pointed to and shared, is symmetric: the same content holds whoever compares, wherever and whenever. And it is discontinuous, since what can be counted and shared comes in separate units, as electric charges, the whole-number labels of particles, and the kinds of particle do. The subjective, the stance, is asymmetric: it sets this

moment apart from all others, this place apart from all others, and this outcome apart from all those that did not occur. And it is continuous, since a stance is an unbroken engagement, which cannot jump without ceasing to be a stance. You, reading this sentence, are in one evening, in one chair, in the one way your day actually went, and your reading runs on from the sentence before without a gap.

The two characters divide the four compounds between them, with no overlap, and they are the two survivors of the theorem: the objective is symmetry with discontinuity, and the subjective is continuity with asymmetry. So the distinction between the subjective and the objective is there to stay, and no better idea will close it. A single description of both would perfect symmetry together with continuity, which the theorem forbids.

The asymmetry and continuity of the subjective give it a definite consequence: it is *situated*. I use the word situated for the stance, to keep it apart from the way a content is *localized*, which is to occupy a place. The stance is situated in time: bound to one moment, the present, set apart from the past that is remembered and the future that is anticipated, and joined without a break to the moment just before. It is situated in space: bound to one place, here, where what is found is found directly, and what is there is inferred. And it is situated in possibility: bound to one actuality, the outcome that occurred, set apart from all those that could have occurred and did not. The three ways of being situated, in time, in space and in possibility, are the three distinctions that I found in every measurement in Chapter 1, past from future, here from there, and the actual from the possible, seen now from the side of the one who measures.

The objective is not situated. A law holds at all times, at all places and for every possible arrangement; shift it in time or in space, and it stays the same. Each content, as found, is localized: it is found at a place, and it keeps any other content out of that place; and the law that holds of that content holds everywhere. So the stance is situated, the content as found is localized, and the law is neither.

Being situated in possibility matters most of the three. It is the structural basis of actuality: the fact that this outcome occurred, and not the others. The mathematics of quantum theory gives every possible outcome its *amplitude*, the weight that sets its odds, and singles out none of them as the one that occurred. The singling out is the act of a comparer, situated in possibility. Because the singling out of one finding is the act of a comparer, the part of quantum theory that the lawful description of quantum theory leaves unfixed is the *measurement problem*. The measurement problem is the question of how one definite finding comes out of a description that singles out none. In a later chapter I call that unfixed part the residual. The mathematics is continuous, the definite finding is discontinuous, and the singling out of one finding is made from the comparer's stance.

The immanent is the event in which structure and potential meet. Among the instances of the immanent are the middle scale, between the very large and the very small; the boundary; and the present. Each is a meeting place. A boundary, like a shoreline, is where two sides meet, as land meets water. The middle scale is where the large and the small meet. The present is where what has been laid down meets what can still come to be. I identify the present, where a stance is situated, as where the subjective and the objective characters meet. There both are present, and neither does away with the other. The theorem guarantees that the two cannot be unified, and the present is where they live side by side.

What We Now Know

We now know which comes first. Interaction comes first; what exists is laid out from what is real. Process is more basic than structure and than possibility; structure and possibility each pull against the other; the three run in one cycle, from possibility to act to record, and the cycle never runs backward; and none of the three can be reduced to the others, occur

without them, or stand in their place. A law exists; a measurement is real; and mathematics models the world's happenings without becoming them.

So the moon is there when no one looks, because the moon is never out of interaction. And the cat in its box is definite inside the box, and is not yet definite for us, because being definite for us is a relation that our looking makes.

The distinction between what exists and what is real matters for the question of this book. Three of the requests the book examines ask something that exists to do the work of something real. Determinism asks a laid-out pattern of law to be the world's happenings. Time travel asks a laid-out shape of time to be walked like a place. The simulated world asks a program, which is structure, to be a world of real acts. The chapters ahead take each of these requests in turn.

Summary

The immanent, the omniscient and the transcendent are the three modalities: process, what is happening here and now; structure, what holds the same for every comparer; and possibility, what makes things possible and new. The three axioms come from the Incommensuration Theorem as its inner structure: the first puts process before structure and possibility, and makes those two conjugate; the second gives the one-way cycle by which possibility gives rise to interaction, interaction to structure, and structure to possibility; the third makes the three distinct, inseparable and not interchangeable. The ground rests on the theorem alone.

Three kinds of claim were made in this chapter. That measurement is interaction, described as measurement; that interaction comes before existence, reality and objectivity; and that a relation laid out as structure exists while an interaction is real: these are claims of the metaphysics itself, the study of the structure of comparison. The reading of general relativity as a theory of the limits of signaling, the reading of quantum theory as a

theory of the limits of knowing, and the reading of the cat as a case of being definite by relation: these are readings of known physics, and each gives the structural reason that a known result or puzzle has the shape it has. The expectation that large superpositions are fragile is an expectation about the kind of thing that happens, and not about how fragile any given superposition is.

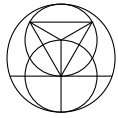
From these the rest follows. Every measurement has a subjective side that no formalism removes, and that side is situated: in one present, at one here, in one actual outcome. Mathematics models physics without becoming it, and a physics that claims to reach beyond every measurement has stopped being physics without becoming metaphysics. The two great theories are theories of access, one of signaling and one of knowing, and the determinism of their formalisms is the determinism of maps, and not of territories. A theory that puts a privileged reality beneath unprivileged interfaces cannot say whether the translation between that reality and the interfaces is real. And the cat is not definite for us, not because we have not looked, and not because looking would disturb it, but because being definite for us is a relation that our looking makes.

Physics has always known that there is no measurement without interaction. What the metaphysics adds is that no measurement can be made without interaction because of what measuring is, by definition, and not because of any practical limit of instruments; and that the whole of physics is the study of what interaction, as such, permits and forbids.

If the world is happenings, is every happening fixed in advance by the ones before it?

PART TWO

CAUSES WITHOUT CLOCKWORK



Chapter 5

Is Everything Already Fixed?

The Demon and the Piano Roll

Is everything that will happen already fixed by what has happened? Many people carry a picture in which it is. In that picture the world is a vast machine: every part pushes the next, and once the parts are set going, every later moment follows from the one before, down to the last detail. In that picture, the state of the world at any moment is completely and uniquely fixed by its state at any earlier moment. In that picture, given complete knowledge of the present, the whole of the future follows with certainty. In that picture, too, given complete knowledge of the past, every detail of the present was already settled.

The clearest form of the picture is an old one. In 1814 the mathematician Pierre-Simon Laplace imagined an intellect that knew, at one instant, every force acting in nature and the position of every thing in it. For such an intellect, he wrote, nothing would be uncertain: the future, like the past, would be present before its eyes. That imagined knower has come to be called Laplace's demon.

A player piano gives the same picture in a single room. A paper roll punched with holes passes over the instrument, and each hole sounds a note. Every note of the piece is fixed on the roll before the first note is

struck, and nothing in the playing adds anything to the roll. The clockwork picture says that the world is like the roll: the whole piece is punched in already, and the passing of time only plays it through.

The picture has real appeal, and the appeal is not foolish. It grows out of the experience of lawfulness: similar conditions produce similar outcomes, physical law is reliable, and causes bring their effects in a regular way. Drop a cup, and it falls, every time. Now push that reliability to its limit. Declare the regularity total, where experience shows it high; exact, where experience shows it close; certain, where experience shows it probable. What you arrive at is *determinism*: the claim that everything is fully and exactly fixed, every value of every content for every context, with no observer and no time needed. (The content of a comparison, as Chapter 1 put it, is what the comparison finds; its context is the respect in which the two things compared are two, such as two places or two moments.)

The question of this chapter is whether the reliability can be pushed that far. It cannot. The reason is a fact about the structure of comparison: the limit that determinism names is a place that no comparison can reach. The reason is not some irregularity that the world happens to contain. In this chapter I show why no world can be purely and completely deterministic, and what science establishes instead.

What Determinism Claims

What, exactly, does determinism claim? To see why it cannot be had, we need the claim stated exactly, and the philosophy of physics supplies it. A world is deterministic if the laws of physics, applied to a complete specification of the state at one time, yield a unique state at every other time. In plain words: describe everything about the world at one moment, completely, apply the laws, and you get one single state of the world at every other moment, earlier or later. The philosopher John Earman gave this definition in his book *A Primer on Determinism* in 1986, and the philosoph-

er Carl Hoefer adopts it in his survey of determinism for the *Stanford Encyclopedia of Philosophy* (in its edition of 2016). It is a good definition, and I take it exactly as it stands.

Said as fully as it can be said, determinism asserts that there can be a complete and absolute specification of all values of all content for all values of all context, independent of any observer. Content comes in three kinds: pattern, force and probability, which is to say a mass, a push, a likelihood. Context comes in three kinds as well: space, time and possibility, which is to say where, when, and among which possibilities a thing is measured. Determinism says that every content is fixed at every context, with nothing left over. Every variable is fixed, every setting of every quantity is specified, and every state of affairs is fully determined by all the states that came before it.

Determinism is therefore a claim about what a complete description of the world could fix, in principle: every content at every context. The claim holds only where nothing remains: no *degree of freedom* left unspecified (a degree of freedom is a way in which something could still vary), and no gap between the description and what is described.

Two features of the claim matter for everything that follows. The first feature is that determinism needs nobody. The complete specification does not require anyone to observe it; it does not change with the place or the scale of any observer; and it is indifferent to whether any finite point of view exists at all. That independence from every standpoint is what makes it determinism rather than prediction. A prediction is made by someone, from where that someone stands. Determinism is the specification that would be there with no one standing anywhere.

The second feature is that determinism needs no time. A fully deterministic account is already complete. The next state is already contained in the present state and in the rule that carries the present state forward, and so the next state is not waited for. The passing of time adds nothing to the specification that was not already there.

A finished novel is a fair picture. Every page is already printed. Reading the novel in order adds no event to the story, and the page numbers only say where each event stands. One may still tell the story in time words: “on page forty the letter arrives, and on page forty-one she reads it”. In the same way one may still describe a deterministic system in time words: “given this state now, that state follows a moment later”. The time words are allowed, and they add nothing. The rule and the first state already settle the next state; the clock is a convenience, and never an ingredient.

So two kinds of time are in play, and they must be kept apart. The first kind of time is time as one more line of measurement on a picture laid out all at once, like the page numbers of the novel. The second kind of time is the order of actual acts of comparison, each made here and now, one after another, in one direction. A deterministic specification has the first kind of time and lacks the second, and causes need the second kind. In a process that runs in the second kind of time, something happens that was not already contained in what came before; on a deterministic path laid out in advance, nothing does. Chapter 14 gives the two kinds of time their names and follows them further.

Two Perfections at Once

Put in the fewest words, determinism is the claim that one law carries every state into exactly one next state, without a break and without exception. What does that claim ask of a description? It asks for two things at once, and Chapter 2 gave us the words for both.

It asks first that the laws hold the same in every setting: the same rule carries one state into the next here and there, now and later, for this system and for that one. Holding the same in every setting is sameness of content across difference of context, which I called symmetry in Chapter 2: change the setting, and the content stays the same. Determinism asks for symmetry without exception.

It asks second that the state be fully specified at every point, with no gaps, no missing values, no ambiguities, and no place where the description gives out or the state is left open, so that each state passes into the next with no break. A state fully specified at every point, with no break, has sameness of content within sameness of context, carried all the way down, which I called continuity in Chapter 2: hold the setting fixed, and nothing jumps. Determinism asks for continuity without remainder.

Picture the two demands as two floors. One floor is laid with identical square tiles: the same tile in every place, which is symmetry. The other is a smooth ramp poured in one piece: move a little along it and its slope changes only a little, with no seam anywhere, which is continuity. Laplace's demon needs one law that carries every state into the next, the same everywhere and at every time: the repeated tile, made perfect. And the demon needs every state specified at every point, with no gap, so that each state passes into the next without a break: the seamless ramp, made perfect. The standard definition of determinism, then, is a demand for perfect symmetry and perfect continuity in one and the same description.

That demand meets the theorem of Chapter 3. The Incommensuration Theorem says that symmetry and continuity cannot both be applied, at once and absolutely, to one comparison, and that neither can asymmetry and discontinuity. The Incommensuration Theorem is a structural fact about the nature of knowing.

The reason for the theorem, in one sentence, is this: a content that is the same whether its context changes or does not change has no dependence on its context at all, and a comparison in which content and context cannot be told apart is no comparison. Content and context are each defined against the other. Perfect symmetry together with perfect continuity asks the two of them to do no work, and then the act that was to be described is no longer there to describe.

Now set the two results side by side. Determinism is the assertion of total symmetry and total continuity, applied at once and without remainder. Total symmetry with total continuity is, word for word, the first

pairing the theorem forbids. So no world is purely and completely deterministic. Total specification, which is continuity, of lawful structure, which is symmetry, is ruled out by the structure of comparison itself. The demon asks for exactly the pair that the theorem forbids.

Determinism is therefore unreachable as a matter of structure. Determinism is not a claim that happens, as a matter of fact, to come out false of this world. Carried to completion, the concept of determinism erases the difference between content and context that comparison needs in order to be comparison. And a world that admits no comparison is not a world in which anything is determined, because it is not a world in which anything is anything.

What the Conclusion Rests On

What does the conclusion rest on? It rests on the structure of comparison alone. The theorem itself is no claim learned from experiment. It leans on what comparison is. So the conclusion stands before any experiment is run, and it stands apart from the places where a reader might expect an argument about determinism to lean: it uses no feature of quantum theory; no theory of chaos (the study of systems, such as the weather, in which tiny differences grow into large ones); no uncertainty principle (the rule of quantum theory that certain pairs of quantities cannot both be pinned down at once); and no particular formalism (a formalism is the system of mathematics and symbols in which a theory is written). And by the Root Tautology of Chapter 1, which says that interaction, comparison and relation are one act named three ways, comparison is interaction, and interaction is what physics is about. So a limit on comparison is a limit on every physical interaction, and determinism is excluded for the world itself, before any experiment is run.

What the argument establishes is exact, and its scope is exact too. The argument leaves the regularity of the world standing: the laws of physics are reliable, experiments repeat, and the future is very largely fixed by the

past. What fails is the claim of totality: the claim that the regularity can be complete, that the specification can be without a gap, and that both can be perfected in one description. Within any limited range (a limited region, a limited span of time, a limited range of scale) a description can be very symmetric and very continuous, and the whole success of physics lives there. The theorem concerns the limit, and determinism is a claim about the limit. So the Incommensuration Theorem reaches determinism.

The argument has one more feature: whoever questions it uses the structure of comparison that it describes. Any attempt to reject the Incommensuration Theorem uses the structure of comparison that the theorem describes. To argue that comparison has no built-in structure, one must compare the structured account with an unstructured one, using a content, a context, sameness, difference, a standpoint, and something the standpoint is about. A content, a context, sameness, difference, a standpoint and what the standpoint is about are the six parts of every comparison that Chapter 1 named. So a reader who doubts this chapter makes a comparison in order to doubt it, and in making it supplies everything the chapter needs.

The Two Edges

Where, then, does determinism live? Determinism has a place, and its place is an edge. Comparison has a domain, as Chapter 1 said: the *domain of comparison* is the range within which comparisons can be made. The domain has two edges. At its large edge stands the speed of light: from beyond the reach of light, no signal arrives. At its small edge stands Planck's constant, the very small fixed amount that sets the size of the steps in which quantum findings come. Below the amount that Planck's constant marks, no measurement can pin down together what a conjugate pair of quantities would need fixed (a conjugate pair, as Chapter 3 put it, is two quantities each of which can be made exact only at the cost of the other, such as where a particle is and how it is moving). What I mean by a

boundary is exactly this: no coherent signal. From beyond either edge there is no meaningful measurement, so there is nothing that can be compared, and so there is nothing that can be known.

Imagine seeing a city from every window of every building and from every street corner, all in one look, from no particular place at all. Determinism is that look, taken of the whole world. It represents what lies beyond the large edge. It is the view from everywhere at once: the whole of the state across the whole of the context, the world seen as a completed object with no standpoint inside it. In the terms of Chapter 4, it is the omniscient mode (structure, what holds the same for everyone, laid out as a pattern) pushed past the edge of what can be measured.

The bill for that view comes due at the other edge. If determinism could hold, determinism would have to be made good all the way down to the smallest scale there could be, which I call the absolute microscopic: the scale of exactly zero size. Determinism would have to reach the absolute microscopic because the claim of determinism is that nothing is unaccounted for at any scale, however small. Every detail of every state, at every size, down to no size at all, would have to be fixed.

Beyond the small edge lies the opposite of determinism. *Pure randomness* is a content found different whatever its context: complete indeterminism, in which nothing at all is fixed. Pure randomness is an abstraction, as determinism is. Whatever crosses into the domain from below the small edge cannot appear as anything other than random, because nothing about it could have been compared before it arrived. What arrives in that way, unable to appear as anything other than random, is hard randomness: the remainder that every comparison leaves, real, and inside the domain. Chapter 8 takes hard randomness up in full. Here it is enough to see that hard randomness is met inside the world, and that it is no edge.

Between the two edges runs a line, and four ideas sit along it. At one end is determinism: everything can be predicted, always, and nothing is unaccounted for. Next comes causality: most things can be predicted, most of the time. Next comes choice, between causality and pure random-

ness: some of what happens follows from what went before, and some of it does not, and both parts are real. At the far end is pure randomness: nothing can be predicted at all, and everything is unaccounted for.

Determinism and pure randomness are both limits, and neither one is inside the domain. The world we can compare is neither of them. It is the middle: causal, and open. Hard randomness is real, and it is met inside the domain; it is not a limit. The determinist takes one edge for the whole, and the believer in pure randomness takes the other. Both mistake a boundary for a territory, like a traveler who walks the border of a country and believes he has seen the country.

The theorem excludes the two limits that lie past the edges, determinism and pure randomness, one with each half. Its first half excludes determinism, as we have seen. Its second half excludes pure randomness: a content found different whatever its context would be asymmetry and discontinuity perfected together, and nothing anywhere at all can be both absolutely asymmetric and absolutely discontinuous. So the world is neither the player piano's roll nor a heap of unconnected noise. Determinism and pure randomness are the two forbidden pairings, and every actual world lies between them.

Three Scales

At what size of things would determinism have to hold, and at what size are causes found? The way a whole hangs together, its coherence as a whole, is called *teleology*, and a description of a thing by its purposes, ends or final states tries to capture it. Teleology is named here because it is the third of three descriptions, and each of the three has its own scale; seeing all three in their places shows where determinism's demand falls. Causality, determinism and teleology are not rival descriptions of one level of reality. Each belongs to a different one of the three scales that Chapter 1 named.

The first scale is the microscopic, the very small, the scale of the most basic physics: particles and fields, below which the old idea of a definite path, a track through space, no longer applies. The second scale is the mesoscopic, the middle, the observer's own scale: living things, instruments, and the readings on the instruments' dials, the "moderate-sized specimens of dry goods" that fill ordinary life, in the phrase of the philosopher J. L. Austin (1962). (Physicists use the word mesoscopic for the scale between single atoms and bulk matter; in this book I use the word mesoscopic for the observer's own scale.) The third scale is the macroscopic, the very large: the scale of the cosmos as a whole, the total system taken without reference to any observer inside it.

Determinism's demand falls on the microscopic. Its claim reads the same at every scale, but it requires the complete specification at every scale, and the microscopic is where that specification has to be made good. The microscopic is where the hope of complete specification has its natural home, and it is where that hope fails.

Causality belongs to the mesoscopic. The mesoscopic is the scale where science works, where experiments are run, and where a finite observer can see that what came after is consistent with what came before.

Teleology belongs to the macroscopic. A description by purposes makes sense only when the whole path of a thing is available to be described, and that needs a vantage point that no observer inside the world can occupy.

Laid out as one picture, the three scales stand one above another. At the bottom is the very small, where determinism would have to be paid in full and cannot be. In the middle is the observer, reading a dial, where causes are found. At the top is the whole, where purposes would be read from a place no one inside can stand. So each scale sets a limit on a finite observer: at the bottom, a specification no one can complete; in the middle, the causes one can see from where one stands; at the top, a view of the whole that no one inside can take.

A Picture With No Time in It

Why does determinism seem so natural to so many physicists? One further difference answers that question, and it is the difference most often lost. Determinism is *atemporal*, which means that it has no time in it. The complete specification it imagines is not a specification at a time, in a place, from a standpoint. It is the whole laid out at once, in the way a mathematical object is laid out. Causality is temporal. Causality is a pattern seen by an observer who is somewhere and somewhen, looking from a past that is known toward a future that is not yet.

Most physicists are entirely comfortable with timelessness, because the models they work in are timeless and the models work. I regard that comfort as a mistake, and the next two sections say why. It is a mistake one can make while remaining perfectly consistent within the model. The fault is not one of logic. It is a fault of philosophy: the argument is valid, and it is not sound. An argument is *valid* when its conclusion follows from its starting points; it is *sound* when it is valid and its starting points are also true of what it is about. The physicist who treats the world as deterministic because the model is deterministic is right about the model, and wrong about what the model is a model of.

A Ball on a Dome

The mathematics most often assumed to guarantee determinism already shows determinism giving way, and the philosophy of physics has found several places where it does. These results come after the argument, which is already complete: they are what the argument leads one to expect, and reading them in its light shows why they are there.

The cleanest case is Norton's dome, set out by the philosopher of physics John Norton in a paper of 2008. Picture a smooth dome, like an upturned bowl, with a ball resting exactly on its top. The surface of the dome is shaped in a special way, so that the force on the ball grows as a fraction-

al power of its distance from the top (as the square root of that distance, for instance, rather than in step with it). Newton's laws allow the ball to sit at rest on the top forever. They also allow the ball to sit at rest for any length of time one likes, and then, without any cause, to begin to slide down in any direction one likes. Newton's equations are satisfied throughout and the starting state is fully specified, and yet the future of the ball is not unique.

In the language of mathematics, there is a theorem that guarantees one single future for equations like Newton's, and that theorem needs the force to change smoothly in a strict sense: over any small step, the change in the force must stay within some fixed multiple of the length of the step (mathematicians call this condition Lipschitz continuity). At the very top of the dome the slope of the force has no limit, and the condition fails at that one point.

Read the dome in the terms of this chapter. The laws are perfectly symmetric: the same equations hold everywhere, at every time. The failure comes exactly where the continuity of the specification is not perfect: at one point, in the slope of the force. Withdraw perfect continuity by that much, and perfect determinism goes with it. Norton's dome is a place where the two ideals were being asked for together, and could not both be had.

John Earman catalogs more cases. In Newton's mechanics with no limit on speed, a particle may arrive from infinitely far away in a finite time; Earman calls such arrivals space invaders. Then the state in any bounded region at one moment does not fix the state there later, because something can come in from outside. Where the state in a bounded region at one moment does not fix the state there later, the specification has a gap. A gap in the specification is a failure of continuity, in the sense in which I use the word 'continuity', under laws that remain perfectly symmetric.

There are also descriptions of space and time in which a single history can branch into two histories that cannot be told apart at the point of branching. Picture a road that splits in two so smoothly that, at the fork it-

self, no traveler could say which of the two roads he stands on. One past then leads on to two futures, and the place where they part belongs to both of them at once. In the two earlier cases the gap was in the specification of the state; in this case the gap is in the setting, the space and time in which every state is placed. There the context itself fails to be continuous.

And there is the hole argument, which John Earman and John Norton published in 1987. It concerns general relativity, Einstein's theory of gravity as the shape of space and time. General relativity is written so that its laws take the same form however one labels the points of space and time; physicists call that rule general covariance, and it is symmetry carried to its limit. The hole argument shows that one physical situation then admits many mathematically different descriptions, which agree everywhere outside some region (the "hole") and differ inside it. Taken as a description, the mathematics then fixes what happens inside the region only if one gives up something about what its points refer to: one must stop taking each mathematical point to name a place that is there in its own right, apart from what happens at it. Here symmetry, carried to its limit in the form of general covariance, undoes the uniqueness of the specification: the description no longer gives one single answer. So Einstein's theory already shows the limits of determinism within its own mathematics.

These results belong to Norton and Earman, and to the physics and mathematics they draw on, and I add nothing to them. What I claim is that the theorem says why there had to be results like these. Determinism requires a pairing that cannot be perfected, so any formalism pressed toward determinism will give way at some point, and the point where it gives way will be a point where symmetry and continuity have been asked to hold together at the limit. Each of the cases above is a breakdown at just such a point. The literature has found the cases one at a time and treated each as a curiosity. Read through the theorem, they are one conflict, between perfect symmetry and perfect continuity asked for together, showing through in different places.

The dome also shows why determinism cannot be had in part. One point, in one slope of one force, is enough to undo it. Determinism is impossible by the Incommensuration Theorem, and it cannot even be moderately held, because any crack in the specification is the end of the specification.

The Fence and the Field

What should we make of an argument against determinism that is itself made with exact, deterministic tools? The whole situation has an irony in it, and the irony is what defines the situation: the argument against determinism is made with deterministic tools. The tools by which the argument against determinism is made are themselves deterministic. Mathematics is deterministic. Logic is deterministic. The theorem that guarantees one future for well-behaved equations is a piece of exact reasoning, and so is Norton's display of its failure at the dome. The Incommensuration Theorem itself is an argument from definitions, valid at every step. And what these purely deterministic methods formally establish, in symbols and in the abstract, is that there is no internally consistent epistemology of determinism in actual, concrete reality: no consistent account of how determinism could be known to hold of anything real. The tools with which we describe the world are what guarantee that the world is not fully describable by those tools.

The irony follows from what a formalism is and what reality is, and it is no failure of physics. Any metaphysics adequate to physics (any account of what physics is about at its foundation) has to account for the irony rather than explain it away. On the account I give, a formalism is a system of relations. Relations are the mathematics of comparison, and they are timeless, symmetric and exact. Reality is a system of interactions. Interactions are comparisons that occur, somewhere, somewhen, from a standpoint, and they are in time, one-way, and finite. The Root Tautology says that relation and interaction are one act seen from two sides. Two

sides of one act remain two sides. A formalism can describe the limits on interaction with perfect exactness and still not be an interaction, and the thing it describes need not have the exactness of the description.

The two great theories of physics are both formalisms of that kind. General relativity is a theory of the limits of signaling: who can send a signal to whom, from where and when, to where and when. Quantum theory, the theory of matter and light whose findings come in countable steps, is a theory of the limits of information: who can know what, and when. Both are theories of access, of what can be measured and what cannot, and both are written in a formal, deterministic way. Yet the information whose limits they describe need not itself be deterministic, even though the theories that describe the limits are.

A surveyor can set the fence posts around a field to the width of a hair, and the grass inside still grows as it grows. A perfectly exact fence does not make the field inside it exact. The irony of determinism is that the fence was mistaken for the field.

The comfort with timelessness is a mistake one can make consistently, and here is why. Within the formalism, determinism is a platonic truth: a timeless truth, true in the way a theorem of geometry is true, and it is what the relations, taken as relations, say. The mistake is to carry the truth of the relations across to the interactions without noticing the crossing: from the timeless to what happens in time, from the symmetric to the one-way, from what needs no observer to what cannot occur without one. (An observer, here as everywhere in this book, is a comparer: anything that makes a comparison, a person, an instrument, or an atom that absorbs light.) Nothing in the formalism licenses the crossing, and it is the crossing, not the formalism, that the theorem forbids.

So the conclusion of this chapter concerns findings, not formalisms. An equation that carries each state into exactly one next state is deterministic in the formalism, which exists as structure, and not in the findings, which are real. (Chapter 4 drew the line: to exist is to stand as structure, the same for every comparer; to be real is to be in an actual act of compar-

ison, here and now.) A deterministic equation is therefore no counterexample to the conclusion. Its determinism is a property of the formalism, which exists. What a comparison finds is real, and no world is deterministic in what comparisons find. The completeness that determinism adds to the regularities of science is nothing science has ever found; it belongs to the mathematics in which science writes its findings down.

What Remains

What is left when determinism goes? Causation remains. The failure of determinism leaves causation, and not nothing, because the pairing that causation has is never forbidden. Causation is continuity with asymmetry: a connection without a break, along an order that runs one way. That pairing is one of the two survivors of Chapter 3, the one that is smooth and going somewhere. (Chapter 6 keeps the word causation for that one-way order itself, and the word causality for the order as an observer perceives it.) So causation survives where determinism does not. Causation is established, and determinism never.

Causation rests on continuity, and it needs a comparer for whom one finding follows another. Every world has causation to some positive degree. What science establishes is causal law, not deterministic law, and the difference is the difference between a pattern seen by an observer in time and a specification that needs no observer and no time.

One stronger claim falls with determinism, and I name it now because it returns later in the book. Some have proposed that even the setting an experimenter chooses for a measurement, and the state of what is measured, are both fixed in advance by one law from one history. The proposal is called superdeterminism, and Chapter 11 gives it a full hearing. It asks for everything determinism asks for, and more. So every form of it that asks for complete determinism is excluded a fortiori, that is, with all the more force. And every form of it, complete or not, that makes the choice of a setting only a consequence of earlier conditions makes a different mis-

take: it reduces the choice of a setting to the causal history of the particles. Cause, change and choice are three distinct parts of what is real, one for each of the three modes of Chapter 4: cause for structure, change for process, choice for possibility. (Chapters 6 and 7 walk through all three; a choice, there, is the bringing about of something new that nothing before it fixed.) By the third axiom, which says that the three modes are distinct, inseparable and not interchangeable, choice and cause are distinct, and neither can stand in for the other.

So the chapter ends on the three, cause, change and choice, and on the remainder that no law fixes, which the theorem requires in any world that admits comparison at all. The next chapter gives that remainder its name, the residual.

Why It Matters

The clockwork picture is excluded at its root, before any experiment. It asks for a world in which the structure of everything, which exists, stands in for every finding, which is real, and it asks the description to be perfect in two ways that no one description can be perfect together. The clockwork picture, excluded at its root, is the first answer that I give in this book to the question of what cannot be, and why. Determinism asks for a pairing that comparison cannot support, and so, taken all the way, it describes a world in which nothing is anything. The same request, that what exists stand in for what is real, returns later in the book in the ideas of time travel and of a world that is a program.

Summary

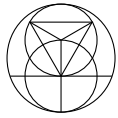
The central claim of this chapter, that no world is purely and completely deterministic, is a structural exclusion: it says what physics will not find, and it says so before looking. That science establishes causal law and not

deterministic law, and the readings of Norton's dome, Earman's space invaders and the hole argument, are readings of known results, each giving the structural reason the result holds.

Determinism is the claim that everything is fully and exactly fixed: every value of every content for every context, independent of any observer and needing no time, so that one law carries every state into exactly one next state, without a break and without exception. We now know that the claim asks for perfect symmetry, the same law in every setting, and perfect continuity, every state specified without a gap, together in one description; that the Incommensuration Theorem forbids exactly that pair; and that therefore no world is purely and completely deterministic. The conclusion depends on no experiment, on no feature of quantum theory and on no theory of chaos; it depends on the structure of comparison alone.

Determinism lives at an edge, never inside the world. It is the view from everywhere at once, beyond the large edge of the domain of comparison, and its claim would have to be made good down to the scale of exactly zero. Pure randomness lies beyond the small edge, and the theorem excludes it as well. The world we can compare is the middle, causal and open, and hard randomness, the remainder every comparison leaves, is real inside it. The known breakdowns of determinism inside its own mathematics (Norton's dome, Earman's space invaders, the branching histories, the hole argument) are one conflict, between perfect symmetry and perfect continuity asked for together, showing through in different places. Determinism is a timeless truth of the formalism, carried to the world only by a crossing the theorem forbids: the fence mistaken for the field. Determinism cannot be had, and it cannot even be moderately held, because any crack in the specification is the end of the specification. What science establishes instead is causation, and every world has it to some positive degree.

Then what is a cause, if it is not a clockwork?



Chapter 6

What Is a Cause, Then?

A Click That Was Caused

If nothing is fixed in advance, are causes real at all? The last chapter showed that determinism cannot be had. Determinism is the claim that everything is fully and exactly fixed: every value of every content (what a comparison finds: a mass, a force, a likelihood), for every context (the setting in which it is found: where, when, and among which possibilities), with no observer and no time needed. Determinism needs one law that is the same everywhere, which is symmetry made perfect, and every state laid down without a gap, which is continuity made perfect; and the Incommensuration Theorem forbids those two to be perfected together. A reader may fairly worry that causes go down with determinism. Causes stand, and in this chapter I show why causes stand.

Start with a picture. A Geiger counter sits on a bench beside a lump of radioactive ore. A *Geiger counter* is an instrument with a gas-filled tube that clicks each time a fast particle passes through the tube. *Radioactive* ore is rock that holds atoms which break apart on their own, each sending out a particle as it goes; physicists call that breaking apart *decay*. Every so often the counter clicks. Each click has a cause: an atom in the ore decayed, sent out a particle, and the particle passed through the counter's tube. Nobody doubts that the decay caused the click. Yet nothing fixed when that atom would decay. Physics gives the chance that the atom decays

within the next minute, and that chance is all there is to know in advance; no record of the atom's past, however complete, names the moment. The click of the counter is an event that was caused and was not determined. The click followed from the decay, and nothing made the decay happen at that moment rather than at another.

The counter is a small picture of the whole of this chapter. Every event in a world can have a cause without anything in that world being fixed in advance: a world can be thoroughly causal without being deterministic at all. The words "cause" and "determinism" are used as though they were interchangeable, across the sciences, in law and in ordinary speech, as though they named one relation looked at from two angles. They differ in structure, in scope, and in what they require of an observer. Treating them as one changes the answer to every question that depends on the idea of a cause: whether anyone ever chooses, whether the future is open, whether a mind could be a program, and whether the universe could be a simulation. That determinism is impossible does not make causes impossible, and the whole practical weight of the argument rests on keeping the two apart.

What a Cause Is

Watch someone strike a match. The match flares; a moment later the paper beside it catches. The watcher registers that the second followed the first, and that the same thing has followed before, every time a lit match met dry paper. That registering is causality, seen from where it is actually found.

In An Immanent Metaphysics I defined causality as "the subjective perception (context) of a consistency between two objective contents". In plain words: two things happen out in the world, the flare and the catching, and these are the objective contents, what the comparison is about, whoever makes it; and someone registers a consistency between them, and that registering is the subjective perception, which is the context in

which the two are compared. So *causality*, or *cause* in ordinary speech, is a consistency that an observer finds between two things that happen in the world, one following the other, again and again.

Where determinism needs no observer, causality is always relative to one. Every causal relation belongs to a perceiver who stands at a particular scale and at a particular moment, and who registers a consistency between what came before and what follows. The perceiver here is an observing self. Anything that makes a comparison can register a reading; an instrument can, and so can a thermostat. Registering one reading and then another is not yet finding a cause. The perception of a consistency, as a cause, belongs to a self that holds both events in view, the earlier and the later together, and expects the one from the other. An observing self, in this book, is a comparer that holds an earlier and a later event in view together and expects the one from the other; Chapter 7 says what a self is. Causality requires a finite observer set inside a domain (a range of the world within which comparisons are made), working at the scale at which perceiving, knowing and acting take their natural form.

Where determinism is timeless, causality runs one way in time. The causal relation runs from earlier to later. It presupposes a direction of time, and a process of becoming in which the observer's registering of the consistency happens in a sequence. Picture a frozen block universe, all of time laid out at once like the pages of the finished novel in the last chapter. Causality appears in that block only where someone is there to perceive a cause, and perceiving is itself an act in time, which cannot be reduced to anything timeless.

Causality has a place in a set of three concepts that belong together: change, causality and choice, which I call the *triple of the real*. Each of the three plays one of the three roles that Chapter 4 called the modalities: the immanent (what is happening, from within, here and now), the omniscient (structure held in view as a whole, outside time, the same for everyone), and the transcendent (what could be: the open, possibility, the interior). In that triple, causality has the nature of the omniscient. It is concerned with what is actual and what is conserved: with what is kept

the same across a change. The pattern that causality registers is omniscient, the pattern conserved across the change; the act of registering the pattern is in time, and runs from earlier to later. A modality is given to a term only relative to the triple in which the term is considered, and here the triple is change, causality and choice. Chapter 7 walks through the triple of the real in full.

The Roll and the Pianist

A player piano and a pianist can play the same piece. In the player piano, a paper roll fixes every note before the first is struck; nothing in the playing adds anything to the roll. The pianist plays each note in the light of the last, hears that it followed, and could have played otherwise. The roll is a picture of determinism, and the pianist a picture of causality, and the three differences between them can be stated exactly.

The first difference is about the observer. Causality requires an observer; determinism requires none. A deterministic specification describes everything from no point of view. The philosopher Thomas Nagel, in a book of 1986, named that kind of description the view from nowhere: the description that leaves nothing out because it takes no particular position. A causal relation is always someone's causal relation: a consistency perceived from inside a domain, by a subject who stands at a scale.

The second difference is about the direction of time. Causality requires a direction of time; determinism does not. A complete deterministic specification fixes the past as fully as it fixes the future. The content of a deterministic specification has no direction of time in it: the arrow of time (the one-way running of time from past to future) belongs to the way the description is used, and not to what the description says. Causality runs one way: from what is perceived first to what is perceived to follow, from earlier to later in the experience of the observer.

The third difference is about the smallest details. Causality pays no attention to the microscopic details of a process; determinism must attend to every one of them. (The *microscopic* is the scale of particles and fields, the scale of fundamental physics.) "The match caused the fire" is a judgment made at the scale of the observer, and it is indifferent to the quantum states of the atoms in the match. A deterministic specification, if one existed, would have to fix those states as well, because any detail left unfixed, any way in which anything could still vary, breaks determinism's claim to be complete. This third difference is structural, and it has nothing to do with how fine an instrument is: causality, by its nature, does not reach into the microscopic, and determinism, by its nature, must.

In Incommensuration, my book on the theorem itself, I put the three differences in one sentence: "Causality asserts a consistent pattern of observations of relations between events, as perceived by an observing self. It requires reference to an observer. It is time-asymmetric. It makes no claim about the complete microscopic specification of all events". Time-asymmetric means that causality runs one way in time, with causes before their effects. The difference between causality and determinism can be put as two sentences side by side. Causality says: similar conditions reliably produce similar outcomes. It does not say: all conditions produce unique, fully specified outcomes for all time. The second sentence is determinism, and it is a different claim.

What Science Can Establish

Science can establish causality. It cannot establish determinism, and I mean cannot in the strict sense. The method of science carries a quiet assumption of symmetry: repeat an experiment and it gives the same result. Every check of that symmetry is itself a finite set of comparisons, so no method can check that the symmetry holds everywhere. And every meas-

urement has a resolution, a finest grain below which it cannot tell two values apart, so no method can check that the content is continuous everywhere.

So the claim science can make is causal law, and never deterministic law. The limit of science to causal law belongs to what science is, and no better science would remove that limit. Science is a system of comparisons made from standpoints, and the theorem says that such a system cannot reach the limit that determinism names.

Every scientific finding passes through observation. Every scientific finding is shaped by the limits on what a finite observer can reach of a domain, and it is made at the observer's own scale. Science can narrow what is possible: that is engineering, which narrows what a device can do to what it is meant to do. Science can also widen what is possible: that is discovery, which finds possibilities no one had seen. The complete specification that determinism requires is out of the reach of science. The aim of science is causal; its achievement is causal; its limits are the limits of causality itself.

Science is, by its nature, built entirely on observation, and observation can establish only causal relations. To demand that science deliver deterministic certainty is to demand that science become something other than itself. What science delivers, knowledge of causes shaped by the conditions of observation, is what a finite observer set inside a domain can, in principle, achieve. Science finds regularity: the same link from what comes before to what comes after, found again across changes of circumstance. The completeness that determinism adds to that regularity is something science has never found; it is a property of the mathematics in which science writes its findings down.

Why the Theorem Leaves Causes Standing

Chapter 3 pictured the theorem with a floor of identical square tiles and a smooth ramp poured in one piece: tiles can follow the ramp as closely as anyone likes, yet a tiled floor always has seams, and a curve never repeats itself exactly. A floor cannot be perfectly repeating and perfectly seamless at once.

That everyday fact is one instance of the Incommensuration Theorem, which Chapter 3 set out: symmetry (the same content across a change of setting, like the tiles) and continuity (the same content within the same setting, so that nothing jumps, like the ramp) cannot both be applied, at once and absolutely, to one comparison; and neither can asymmetry and discontinuity. The Incommensuration Theorem is a structural fact about the nature of knowing. The theorem also gives the structural reason why causality can be had where determinism cannot, and it gives that reason twice.

The first reason is that causality involves symmetry: the pattern repeats across different contexts, or it would not be a pattern. Causality does not require total continuity: it does not need the specification to be complete at every point, only good enough to see the pattern by. A finite degree of symmetry together with a finite degree of continuity is exactly what the theorem permits, and it is exactly what causality is. The theorem forbids the two perfected together; the two together in some degree, short of perfect, are found in every real system.

The second reason lies at the level where the theorem rules. It concerns *causation*, the one-way, unbroken order of what follows what at each event, which causality is the perception of; a later section of this chapter keeps the two words apart. Causation's own pairing is continuity with asymmetry: a connection without a break, along an order that runs one way. Causation's pairing, continuity with asymmetry, is one of the two pairings that the Incommensuration Theorem permits, the two survivors

of Chapter 3: continuity with asymmetry, smooth and going somewhere; and symmetry with discontinuity, the same everywhere and countable. So the failure of determinism leaves causation standing, and not nothing.

One more principle joins the theorem here. Every domain is *zero-positive*: its basic measures have a floor at zero and no ceiling. Zero-positivity follows from the three axioms of Chapter 4; it is a consequence of them, and not a fourth axiom. Interaction is one such measure. There is no negative degree of contact between a comparer and the world: contact can be slight, and it can come close to the floor, but there is no degree beneath the floor. A world that admits comparison at all has comparison in it, and so interaction, and its measure of interaction stands above the floor; and every interaction carries causation's own pairing, which the theorem keeps. So every world that admits comparison at all is causal to some positive degree, some degree above zero; that is zero-positivity at work. And by the Incommensuration Theorem, no world is purely and completely deterministic. The two statements sit together without strain. Picture the degree of causality as a reading on a gauge with a floor and no top. From inside the domain, one sees the needle off the floor: every world is causal to some degree. From the edge of the domain, one sees the place where the needle can never stand: total regularity together with total specification, which the theorem forbids, so that no world is purely deterministic. The two statements describe one structure, seen from inside the domain and from its edge.

The theorem's second clause excludes the opposite extreme. *Pure randomness* is a content found different whatever its context: complete indeterminism, with no part of what is found fixed by anything at all. Pure randomness would be asymmetry and discontinuity perfected together, and nothing anywhere at all can be both absolutely asymmetric and absolutely discontinuous. Pure randomness is the limit opposite to determinism. It is an idealization, a description perfected on one side, and the theorem excludes it just as it excludes determinism.

The domain has two edges. The far edge, the *macroscopic boundary*, is the edge past which no signal arrives. The near edge, the *microscopic boundary*, is the edge below which no measurement tells things apart. Total regularity and total specification cannot both be had. Past total regularity lies the macroscopic boundary and the view from nowhere, the whole seen at once from no place inside it, which Chapter 5 called the view from everywhere at once. Past total specification lies the microscopic boundary and pure randomness. The reason is the one Chapter 5 gave: a specification pushed to completion would have to reach the scale of exactly zero size, and so it would have to pass the small edge, below which no measurement tells things apart; and what lies past the small edge is the opposite of a specification, pure randomness, since nothing there could have been compared before it arrived. Between them is the world, which is causal, and open.

Something real lies on the near side of the microscopic boundary, and it must be kept apart from pure randomness. *Hard randomness* is an outcome that no prior specification fixes, however complete, for any comparer whatsoever, even one with unlimited access to the state. Hard randomness is the remainder that every comparison leaves, an outcome not fully determined by what came before. Hard randomness is what arrives through the microscopic boundary, where nothing could have been compared before the arrival. Hard randomness is real and inside the domain. The decaying atom of the Geiger counter is an instance of hard randomness. So pure randomness is an idealization, excluded by the theorem, and hard randomness is real, and the theorem guarantees it. Chapter 8 walks through hard randomness in full.

The floor and the ramp are a picture. The theorem is about every comparison, because it is proved from the meanings of the parts of comparison alone. In physics it appears as the characters of the two great theories. Quantum theory, the theory of matter and light, whose findings come in exact, repeated steps, is symmetric and discontinuous. General relativity, the theory of gravity as the shape of space and time, whose space and

time are smooth and one-way, is continuous and asymmetric. In this chapter the theorem appears as the impossibility of determinism and of pure randomness.

How the Two Get Run Together

The confusion of cause with determinism follows a definite path, and the path has three steps. First, change is collapsed into causation: the process of changing is identified with the lawful pattern of the change. Then choice is collapsed into causation: whatever the pattern does not explain is declared to be an instance of the pattern not yet understood. With all three collapsed into one, causation is identified with determinism, and the two are treated as exact copies of each other. Each step feels harmless, and the last is the one that produces the mistake. At the end of the path, "this caused that" and "the whole state of the world was fixed in advance by one law" are taken to say the same thing.

A philosopher knows that these have to be different. A physicist working inside a model in which they are not different is not wrong about the model. He is wrong about what the model is a model of. A feature built into a model is easily taken for a feature of what the model describes, and careful people make this mistake for exactly that reason: the premise is built into the models they work with every day.

Chapter 5's two words make the point exact: an argument is valid when its conclusion follows from its premises, and sound when its premises are also true of what it is about. An argument that treats causation and determinism as one thing can be valid: each step can follow from the one before. It is still not sound, because one of its premises, that causation and determinism are the same thing, is true of the model and false of the world.

What Is Left Over

Determinism is impossible, so there is always a residual. Chapter 5 met it as the remainder that no law fixes; here it gets its name. The *residual* is what no lawful description fixes: a part of what is real that lawful structure, however completely specified, does not capture. The theorem guarantees the residual in every description. Because symmetry and continuity cannot both be perfected, there will always be aspects of what happens that no structural description, however complete, can render in full.

I want to be exact about what kind of residual this is. The residual is a gap that the theorem requires, and it is a different thing from ignorance. It is not a gap in what we happen to know, of the kind that a more powerful theory or a finer instrument could in principle fill. The residual is there in any world that admits comparison, and it is there for the same reason that comparison is possible at all. There is, in every such world, real newness: something not fixed by prior conditions, not obtainable from structural law, and not reducible to mechanism.

I name this residual choice, and I place choice in the triple of the real with change and causality. Choice is the residual as it is met from inside, by the one in whom the new arises; Chapter 7 shows that the same residual, described from outside, is what Chapter 8 calls hard randomness. The three are distinct, inseparable and not interchangeable, in the same way as the three modalities. Causality is the lawful, repeatable pattern; it has the nature of the omniscient. *Choice*, in this sense, is the emergence of the new, the single act that is only itself; it has the nature of the transcendent. Change is the actual process of passing from one state to another, the interaction in which pattern and newness meet; it has the nature of the immanent, and it is the wider kind that holds the other two within it.

Choice and causality are conjugate with respect to change: each can be perfected only at the cost of the other, neither comes first, and neither can be reduced to the other or defined away. “With respect to change” says where the two meet: both are ways of accounting for change, the larger kind that holds them, and inside change the one is never present without

its partner. Causality therefore implies *potentiality*, the presence of what could happen. A lawful order says what follows from what, and so it is an order among possibilities: where there is lawful order, there are other possibilities besides the one that actually occurs, and potentiality is a concept of choice. Choice is as much a way of accounting for change as causation is. Choices cannot be predicted; a choice that could be predicted would be causal, and not a choice.

To collapse choice into causality is to claim that newness is an illusion: that the appearance of the new is always and only the unfolding of what was already hidden in the old. The theorem says that this collapse cannot be carried through. The three remain three. That the three remain three is what makes superdeterminism malformed, which is worse than false. Superdeterminism, which Chapter 5 met, asks for everything determinism asks for and then for more; the more that it asks for, that the experimenter's choice itself be causation, is the collapse of one member of the triple into another, which the structure identifies as a *category error*, taking a thing of one kind for a thing of another kind. Chapter 11 takes superdeterminism up in full.

One way of seeing why the triple must be a triple comes from the side of the self rather than the world. Suppose a person has choice, but the choices have no effect: nothing that is chosen makes any difference to anything else. Where choices have no effect, there is no causation, and there is also no choice, because a choice that changes nothing was not a choice of anything. Now suppose the opposite: everything the person does affects everything, and every act spreads its effects without limit. Then there is no free choice either, because any move at all carries too much, and the person is bound.

The first case has no symmetry: the act does not carry across contexts. The second has too much continuity: the act carries without a break into everything. Choice lives where there is enough of each and not all of either, which is exactly where the theorem says anything at all can live.

Where Causes Live

There are three scales of description. The microscopic is the scale of particles and fields. The *mesoscopic* is the scale of observers: organisms, instruments and their readings, the scale of ordinary experience. The macroscopic is the scale of the cosmos as a whole, the total system considered without any particular observer inside it. Causality is perceived by observers set inside a domain, at their own scale; that is why causality is native to the mesoscopic.

A reader may ask whether causes stop at the edges of the observer's scale. They do not stop there. At the two edges of the observer's scale, causality is joined to its partner, choice, and the two stand to each other differently at each edge; and below and above the observer's scale physics finds causal order all the same. The answer has these two parts.

The first part is choice. Choice, causality's partner in the triple of the real, is joined to causality at both edges of a domain, its smallest scale and its largest. Choice and causality are conjugate: each can be held more fully only as the other is held less fully, and neither comes before the other. At the microscopic limit the two stand in strict opposition: in a single quantum event there is a real alternative, and no consistency that anyone could perceive in that one event (Chapter 7 shows this). At the macroscopic limit, in the words of Incommensuration, "the degree of realness of either aspect is directly proportional to the degree of realness of the other". In plain words: across the whole, the more fully lawful pattern is part of what is real, the more fully openness is part of it too. Conjugacy is a trade in how fully each can be held; what becomes proportionate at the macroscopic limit is how real each is.

The second part is the order that physics finds. Physics finds a causal order at every scale. The light cone of an event is the region of space and time that a signal sent from the event can reach, traveling no faster than light; the light cone orders what can affect what, and Chapter 12 walks through it. In Measurement I named that order: "The causal order is the

cone of precedence at each event, continuous and one-way: that is causation, in the character of continuity with asymmetry". That order is causation at each event: local, one-way and without a break.

So I keep two words apart. *Causality* is the perception of a consistency between two events, found by an observer at its own scale. *Causation* is the one-way, unbroken order at each event that physics finds at every scale: continuity with asymmetry. Causality is causation perceived; causation is what causality perceives, wherever causation holds.

Keeping the two words apart settles an old question at once, the question of what caused the universe as a whole. A consequence follows: the universe has no cause. To ask for the cause of the whole is to ask for something outside the whole with which the whole could be compared, and the whole has no outside. In Measurement I put it this way: "A description of the whole world from a vantage outside it has no comparison behind it, because the whole has no outside to compare it with". Claims about "the cause of the universe" are therefore category errors: they apply a concept found within the world, between events, to the world as a whole, where nothing stands outside to be compared. The order of the light cone is a different kind of claim: it is causation at each event, and never a causation of the whole. What applies at the macroscopic scale is the proportion of choice and causality in the whole, and that proportion has a name, *teleology*: the coherence of a whole, which descriptions in terms of purposes and ends try to capture.

Causes Before Anyone Saw Them

One further distinction, inside causality itself, completes the picture. *Subjective causality* is causality as this chapter defined it, rooted in an observer's perception of a consistency, and it is the foundation. Subjective causality is the form of causality on which everything else rests: the structure of domains, the possibility of free will, and the workings of learning, of evolution and of feedback.

Objective causality is derived from subjective causality. "The earthquake caused the tsunami", said as a fact about the world rather than about anyone's perception of it, extends a concept whose first home is the subjective into a domain where the perceiver has been set aside in thought. That extension is legitimate and useful, and it is derived. Its limits show at the microscopic and macroscopic extremes. There, the observer who was set aside in thought comes back, because that observer's limits of access shape what any objective statement can contain. There are three such limits. The uncertainty relation limits how sharply two paired quantities can be found together. The light cone limits what can reach what. And the horizon is the edge beyond which nothing can be seen. Later chapters set out all three. Chapter 1 called the light cone and the horizon two faces of one limit, the absolute elsewhere, and Chapter 9 counts them that way, as one of three limits on knowing, with the uncertainty relation and the limit in time.

So a reader may ask: if causality needs a perceiver, did nothing cause anything before there were observers? Events before any observer were related as they are now; what did not yet exist was the perception of those relations as causes. Objective causality is the extension of that perception to times and places where no one stood, and the extension is legitimate: it is true to say that an impact caused the crater on the Moon. What the extension cannot do is stand on its own, as though "cause" named a relation that needs no comparison to be found. Every finding about the early universe is made now, by observers comparing records, and the causal order it reports is the order those comparisons reveal.

Does the perceiver have to be human? No. Anything that makes a comparison, such as an instrument, a thermostat or a sense organ, fills the role of a comparer; in the words of Measurement, "An observer is a comparer: the word names a position inside the world, and it never names a mind". What a comparer registers is a finding, the record of one act of comparison. So a thermostat registers findings, readings of warm and

cold, and the perception of a consistency between findings, as a cause, belongs to an observing self, which holds both events in view and expects the one from the other.

The Non-Polarization Principle

The result so far can be stated as a principle, which the rest of this book applies again and again. I call it the *non-polarization principle*. All worlds are causal to some degree above zero. No world is purely and completely deterministic, and no world is purely and completely logical; a purely logical world would be one made of nothing but the timeless relations of mathematics. The principle reaches the self as well as the world: neither experience nor creativity is determined. Experience comes to a self by perception, and creativity goes out from a self by expression: these are the two directions in which interaction crosses the boundary of the self. In Incommensuration I derived the principle's first instance from the theorem: "No world is purely and completely deterministic. But every world is to some positive degree causal".

Deterministic descriptions remain useful and powerful. Every world has *determinate structure*: lawful pattern, conserved across changes, which deterministic descriptions capture and on which engineering relies. Determinate structure is present to a degree above zero in every world, and complete in none. Determinism is the claim that determinate structure is complete, and that claim has no degrees: it is total or it fails. So the principle denies only one thing about deterministic descriptions: that any of them can be total, that the specification can, even in principle, exhaust the domain it describes. There is always a remainder that determinism, by its own aim of completeness, cannot take in: the observer's point of view, the direction of time, and the mesoscopic grain at which causality works.

A natural question can now be answered. Can there be a world that is purely deterministic and not at all causal, or one that is purely causal and has no determinate structure at all? No. The very conditions that make a

world a world (interaction, comparison, and distinctions between content and context) already carry a degree of causality above zero, and with it a departure from determinism; and they carry a degree of determinate structure above zero, and with it a departure from pure randomness.

The Roll, the Hiss and the Music

Are cause and determinism the only pair of ideas that have been run together in this way? They are one case of a habit: taking one end of a pair of ideas, perfected, for the whole of what is real.

Set the player piano beside the pianist again, and put a third thing beside them: a radio tuned between stations, hissing. The paper roll is order with no openness: every note fixed in advance. The hiss is openness with no order: no sound related to the one before. Music lies between them, and so does every pianist who has ever played: each note follows from what came before, and none was fixed before it was played.

The roll is determinism, symmetry and continuity perfected together. The hiss is pure randomness, asymmetry and discontinuity perfected together. Neither can be found in the world. What can be found is always between them, and the region between them is no compromise and no average of two real things. It is the only real thing of the three; the two ends are what thought makes by perfecting one side of the real thing.

The separation of cause from determinism is the first case of a pattern that recurs, and I call the cases together the chain of distinctions. In each case a pair of ideas names two ends, or *poles*. Each pole, taken absolutely, is an *idealization*: a description perfected on one side, which no domain contains. What is real lies between the poles: a condition in which each side depends on the other, which a finite observer set inside a domain occupies, and from which both poles are abstractions. The non-polarization principle is the general statement of the pattern. For each pair, no domain is purely one pole or purely the other, and every domain involves both to some degree above zero. For the first pair the principle takes the form

already given: every domain is causal to some degree above zero, and has determinate structure to some degree above zero, and no domain is deterministic.

The chain has four pairs. The first pair is determinism and pure randomness, with causality between the two poles. Determinism is the idealization of perfect order; pure randomness is the idealization of perfect openness, a content found different whatever its context. The theorem excludes each pole, the first by its first clause and the second by its second. Causality, together with change and choice, is what is real between them: the triple that An Immanent Metaphysics calls "the composition and basis of all-that-is". The other three pairs are science and mathematics, the physical and the non-physical, and realism and idealism. All four pairs fail at their poles for one reason: to be real is to be interdependent, and each pole asserts either a complete independence or a complete dependence.

Science and Mathematics

Science is the practice of establishing causes through observation and experiment. Science is immanent through and through: every scientific result passes through the conditions of an experiment, made at a scale, by an observer, at a time. Science also works in the omniscient modality: it produces theories that generalize beyond any one observation and claim to hold for cases not yet examined. As a method, science joins the immanent (the experiment) and the omniscient (the generalization); as a practice it also holds the transcendent, the belief that its models correspond to reality.

Mathematics works in the omniscient modality alone. It describes pure relations: structures that hold apart from any particular instance, apart from whether anyone observes them, and apart from time. Its content says what must be the case, given certain definitions and rules, and says it from no particular point of view. What makes a piece of mathematics

count as established, though, is a set of acts: a proof counts as a proof because someone has checked it, step by step, comparing each line with the rules, and every check is an act of comparison carried out in the world. That fact places mathematics in an order, and the order says where the sureness of mathematics comes from. A companion book, Physics Is Not Math, sets out the order in which four fields depend on one another, each presupposing those before it: metaphysics, physics, computation, mathematics. Mathematics remains the finest instrument there is for what is laid out as structure, and it stands after the acts that check it.

Mathematics keeps its exactness. Within its definitions mathematics is exact, and a checked proof is as sure as anything can be. What the order of dependence adds is where that sureness comes from: from acts of checking, which are comparisons made in the world. Mathematics keeps every power it has, in its place.

The pair follows the pattern. The purely mathematical pole would be formal structure with no contact with anything actual; the purely scientific pole would be observation with no formal structure at all. Neither is found. Every science involves mathematics, and every mathematical system, to be about anything at all, must connect to some domain in which comparisons are made.

The Physical and the Non-Physical

In this pair the word physical carries a special sense, narrower than its everyday sense. The *physical*, here, is what is both actual and, in its description, deterministic: every state specified, every change governed by a rule, nothing left over. A clock pictured as nothing but gears, each position fixed by the one before, is the physical pole taken alone. The *non-physical* is what involves potentiality and a process that is not deterministic: the domain in which more than one outcome is available, and in which the complete specification that the physical pole aims at cannot be had. The open moment before a choice, when more than one outcome is still

available, is that side taken alone. The non-physical in this sense is the side of possibility, and it has nothing to do with the supernatural or the ghostly. In Incommensuration I stated the pair exactly: "The physical and the non-physical are both idealizations. The physical is existence, both actual and deterministic; the non-physical is creation, both potential and non-deterministic; and neither can be fully instantiated".

Physics itself measures possibility. Every physical measurement measures one of six quantities: three kinds of content (pattern, force and probability), each in one of three kinds of context (space, time and possibility), which Chapter 19 calls the six measures. Probability and possibility stand among the six beside the others. Planck's constant, which sets the size of the smallest quantum step (the size of the countable steps in which quantum findings come), gives possibility its grain. So the non-physical of this pair is not what physics leaves out. It is the side of possibility taken alone, as the physical pole is the side of fixed structure taken alone.

Every actual system has both sides: some degree of openness, which the free will theorem makes precise for quantum systems (a result about particles measured in two distant laboratories, which Chapter 9 explains), and some degree of actuality, through which its possibilities take effect. The domain of choice, and of real alternatives, is joined to the physical: it is where a system's openness is met from within.

Realism, Idealism and the Four Errors

Two words carry this pair. *Perception* is the direction of interaction that runs from the world to the self; *expression* is the direction that runs from the self to the world. In Incommensuration I wrote: "Perception and expression are the two directions of interaction across the boundary of self: perception runs from reality to the self, and expression from the self to reality".

Realism, taken absolutely, is the claim that perception is independent of expression: the world is as it is, whether or not and however anyone engages it, and the observer can be wholly separated from the observed without loss. That separation is exactly what determinism requires and causality denies. All knowledge passes through observation, and observation is always from a point of view, so the complete independence of reality from perception is an idealization that cannot be kept up absolutely.

Idealism, taken absolutely, is the claim that perception is dependent on expression: what is perceived is made entirely by what the perceiver brings to it, so that the act of perceiving makes its object. Idealism makes the observer the source that makes reality what it is. It gives the finite perceiver a creative power that belongs to creation, the infinite counterpart of choice: where a choice selects among what is open, creation makes the range of what is open, without bound, as Chapter 7 explains. Realism and idealism are the same kind of error. In Incommensuration I said of the two that "each is a limit case taken as if it were a perspective".

Perception and expression can each be taken to two limits, absolute independence and absolute dependence, and that gives exactly four errors. In Comparison I called them the *errors of extrema*, "the four errors that arise when the relationship between self and world is taken to an extreme". Two are errors about the world, and two are matching errors about the self.

Take perception first, in relation to expression; these are the errors about the world. Taken as absolutely independent, perception gives realism: the observer sees without being shaped by anything the observer brings. Taken as absolutely dependent, perception gives idealism: what is perceived is made entirely by the perceiver. Now take expression, in relation to perception; these are the errors about the self. Taken as absolutely independent, expression gives *prejudice*: what is expressed (stated, asserted, created) bears no relation to what is observed. Taken as absolutely dependent, expression gives *reaction*: what is expressed is entirely fixed by what is perceived, leaving no room for a creative response.

Reaction deserves a second look, because reaction is determinism in a person. In Comparison I called it the claim that "expression is absolutely dependent upon perception, which is to say mechanistic and deterministic". A person who believed it would believe that every word they speak is fixed by what they have taken in: the player piano's roll, written into a life.

All four errors fall for one reason. Every action has at least some reaction. In Comparison I put it this way: "There are no absolutely asymmetric interactions. No interaction, no relation, can ever be fully and completely asymmetric. Every action has at least some reaction". Each of the four errors makes the relation between perception and expression run one way only: realism and prejudice by making one side wholly independent of the other, idealism and reaction by making one side wholly dependent on the other. So the non-polarization principle ranks positions: it finds the extremes structurally incoherent, because each asks a relation to run wholly one way, in a domain where no relation can.

Interdependence

The idea beneath all four pairs, and beneath the non-polarization principle itself, is interdependence. An Immanent Metaphysics, aphorism 1.76-6, sets out three terms: "Existence is characterized in the terms of independence. Reality is characterized in the terms of interdependence. Objectivity is characterized in the terms of dependence". To be real is to be interdependent, and interdependence is more basic than both existence (independence) and objectivity (dependence). Aphorism 1.84-3: "There is neither absolute dependence nor absolute independence; there is only interdependence (interaction)".

Each of the two extremes undoes itself. Suppose one term depended absolutely on another. Then everything in the first would be fixed by the second, and nothing in the first could be told apart from what the second fixes. The principle of identity says that "that which is indistinguishable must be the same", so the dependent term would then be no second term

at all, and no relation would be left. Suppose instead that the two terms were absolutely independent. Then no interaction would pass between them, and what enters no interaction is not real, since to be real is to be interdependent.

Independence and dependence are the extremes. Independence is what determinism aims at: a world that is what it is without relation to anything else, every state fully specified without relation to any observer. Dependence is what pure idealism claims: a world that is nothing apart from the perception that makes it. *Interdependence* is the condition in which each term of a relation is what it is only in and through the relation, without being reducible to the relation. Picture two dancers who hold hands and lean back, each against the other's weight. Each stands as she does only because the other does, and yet each is still a dancer of her own, and not a part of the other.

Interdependence is why the non-polarization principle holds. A world that is real is interdependent: neither purely independent nor purely dependent. That interdependence guarantees that the world is always to some degree causal, always to some degree physical and non-physical, always involves both observation and structure, and is never reducible to either pole of any of these pairs.

Consciousness Between the Poles

One more instance of the condition between the poles is the reader's own consciousness, and it prepares Chapter 8, which finds three things bound together wherever there is a first-person point of view. Consciousness is itself an instance of the condition between the poles. An Immanent Metaphysics defines consciousness in three parts. Consciousness refers to the total volume of one's perceptions across all available channels; it is the summation of choice within the volume of the being of the self; and the concept of consciousness combines only the concepts of interaction and creation, without the concept or assumption of existence. In plain words:

consciousness is everything one perceives, taken together; it is where one's choices gather; and it is made of interacting and creating, and not of a thing that exists apart.

The last part places consciousness exactly. Consciousness belongs to the interdependent region between the purely physical, where existence would place it, and the purely non-physical, where dualists would place it. (*Dualists* hold that mind is a second substance, apart from matter.) Consciousness is no kind of existence: it is not a substance, not a property of matter, and not a feature that could be located the way physical objects are located. Nor is consciousness nothing: it is not an illusion, not an *epiphenomenon* (a by-product that does nothing), and not a word with nothing to refer to.

Consciousness is neither reality nor self but their interaction. It is neither the world perceived nor the perceiver, but the act of perceiving. That act is, by the Root Tautology of Chapter 1, one with comparison and with interaction: "Comparison is isomorphic with interaction and relation" (aphorism 1.42-4), which is to say that comparison, interaction and relation are one act under three names. Chapter 8 shows what this act is wherever it happens as a first-person point of view: three things bound together and taken as one act, the one-way direction of time, hard randomness, and the subjective.

The Middle Is No Compromise

The four pairs can now be said side by side. Determinism and pure randomness: one pole is perfect order, symmetry with continuity; the other is perfect openness, asymmetry with discontinuity; what is real between them is causality, with change and choice. Science and mathematics: one pole is the purely mathematical, structure with no contact with anything actual; the other is the purely scientific, observation with no formal structure; what is real between them is actual inquiry, which always has both. The physical and the non-physical: one pole is the purely physical, actual

and deterministic; the other is the purely non-physical, potential and not deterministic; what is real between them is actual systems, with openness and actuality both. Realism and idealism: one pole is realism, perception independent of expression; the other is idealism, perception dependent on expression; what is real between them is interdependence, which is interaction.

Is the middle just a compromise between two good ideas? No. A compromise splits the difference between two things that each exist. Here the poles do not exist; they are made by perfecting one side of what does. The pianist is not half player piano and half radio hiss. The pianist is the real thing, and the roll and the hiss are what one gets by imagining the pianist's order without openness, or openness without order.

Is this account itself a kind of idealism, since it starts from perception? No. It starts from interaction, of which perception is one direction and expression the other, and An Immanent Metaphysics holds that perception is more basic than the perceiver and the perceived. Idealism makes the perceiver the source of what is perceived; this account makes neither the source of the other.

What Kind of Impossibility This Is

Chapter 5 took the standard definition of determinism, the one John Earman gave in 1986 in his book A Primer on Determinism and Carl Hoefer adopted in 2016 in his survey of the subject, and found that it asks for symmetry perfected together with continuity perfected, a pairing that no description of a comparison can support. So determinism is structurally unreachable.

The impossibility of determinism is structural. The description that determinism asks for cannot be the description of a comparison, and a world is what comparisons are of. The impossibility is not a matter of logic: a fully deterministic mathematical object is perfectly consistent as an object. Nor is it a matter of experiment: no experiment is cited, and none is

needed. It is a consequence of what comparison is; comparison is interaction, and interaction is the act of every physical event. The conclusion rests on the theorem and on the Root Tautology, and, as Chapter 5 showed, whoever disputes them uses them, since to dispute them one must compare, and comparison has the structure they state. The world is causal; it is not deterministic; and the argument depends on keeping those two claims distinct.

Can physics show determinism failing in the world itself, and not only in thought? It can, and the Geiger counter on the bench is the first sign: a single atom decays when it does, and no record of its past, however complete, names the moment. Chapter 8 follows that atom into the laboratory, and Chapter 9 follows the experimenter who chooses what to measure.

The Distinction in Its Tradition

The distinction has a history of partial statements, and placing it among them shows what this chapter adds.

David Hume, in A Treatise of Human Nature of 1739, found in everything observed only things that go together again and again (he called it constant conjunction), and the mind's expectation that they will go on going together; he found no necessary connection between a cause and its effect. Hume's finding is the nearest forerunner of the definition given here, causality as the perception of a consistency. Hume took the subjective side of causality to be a lack; I take it to be the foundation, from which objective causality is derived.

Bertrand Russell, in "On the Notion of Cause" of 1913, observed that advanced physics states its laws as relations between quantities, equations that tie one measured amount to others, and does not use the word "cause" at all. That observation fits the third difference: causality, by its nature, does not reach into the microscopic, and the physics of the microscopic has no need of it.

G. E. M. Anscombe, in Causality and Determination of 1971, argued that causality consists in an effect's coming from its cause, and that this does not require the effect to be forced. Her example was a bomb wired to a Geiger counter, so that it goes off if the counter clicks: the explosion is caused when it happens, and it was not determined to happen. Her thesis is the thesis of this chapter, and I give her thesis a structure and a ground. The structure is the three differences, which say what kind of thing each concept is: relative to an observer or not, in time or not, and at which scale. The ground is the Incommensuration Theorem, which shows that causality does not carry determinism with it, and more: that determinism cannot be had at all, and why causality remains when determinism fails.

The long debate about free will has mostly assumed that cause and determinism are one, which is the confusion this chapter takes apart. *Compatibilists* hold that free will is compatible with determinism; *libertarians*, in this old sense of the word, hold that free will requires determinism to be false. Both sides take determinism to be at least a possible state of the world. Determinism is impossible, and so the question changes. The question becomes what choice is in a world that is causal and open; whether freedom can survive in a fixed world no longer arises, since there is no fixed world.

What We Now Know

We now know what a cause is. A cause is a consistency that an observer finds from inside the world, between two things that happen, one following the other, again and again. Causes are real, and causation, the one-way order at each event, runs through every scale. Determinism holds in no world. The Geiger counter's click was caused, and nothing determined it. For the question of this book, what cannot be and why, this is the first settled ground. The clockwork world cannot be, because the clockwork needs perfect symmetry and perfect continuity together, and the theorem forbids the two together. The theorem leaves whole the causes that science

finds, because their own pairing, continuity with asymmetry, is one the theorem keeps. The limits that Part Four of this book draws, on signals and on time, are limits on causation of exactly this kind: on what can reach what, and on what can follow what.

Summary

The claims of this chapter are of three kinds. That no world is purely and completely deterministic is a structural exclusion: it says what physics will not find, and it says so before looking. That every world that admits comparison is causal to some degree above zero is a claim of the structure of comparison itself: zero-positivity at work, a consequence of the axioms. That science establishes causal law and not deterministic law is a reading of known results, giving the structural reason the result holds.

Causality and determinism are two concepts. Determinism is the complete specification of every content for every context: it needs no observer, it has no time in it, and it would have to be made good all the way down, to the scale of exactly zero. Causality is the perception of a consistency between two things that happen in the world: it belongs to an observer, it runs one way in time, and it is native to the observer's own scale. Causation is the one-way, unbroken order at each event that physics finds at every scale, and causality is causation perceived. Each scale of description has its own concept. Determinism would have to be made good at the microscopic, all the way to the scale of exactly zero. Causality is perceived at the mesoscopic, the observer's own scale, and teleology names the whole, at the macroscopic.

The Incommensuration Theorem, a structural fact about the nature of knowing, forbids perfecting symmetry with continuity, and so excludes determinism; it forbids perfecting asymmetry with discontinuity, and so excludes pure randomness; and it keeps continuity with asymmetry, which is causation. Every world is therefore causal to a degree above zero and deterministic to none, and a world can be thoroughly causal without

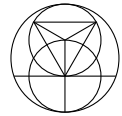
being deterministic at all. Science establishes causal law; determinism is out of the reach of every science, in principle, and therefore in all future practice. Pure randomness, an idealization, is excluded with determinism; hard randomness, an outcome no prior specification fixes, is real, and inside the world.

Because determinism is impossible, there is always a residual: real newness, the emergence of the new, which no lawful description captures and no finer instrument could fill. Its name is choice, and choice stands with change and causality in the triple of the real, the three distinct, inseparable and not interchangeable.

The separation of cause from determinism is the first link of a chain of distinctions. In four pairs (determinism and pure randomness, with causality between; the purely mathematical and the purely scientific; the physical and the non-physical; realism and idealism) each pole taken absolutely is an idealization, and what is real lies between. The four errors of extrema (realism, idealism, prejudice and reaction) are the poles of the relation between self and world, and none can stand, because every action has at least some reaction. The common reason is interdependence, the condition of being real, which is more basic than independence or dependence. Consciousness is an instance of the same condition: neither world nor perceiver, but the act of perceiving.

The universe is not a machine. It is a process: lawful, causal, and irreducibly open. There is always a residual of real newness, and the residual is no fault in the world and no fault in physics. It belongs to the structure of reality, and it is the same feature of that structure that makes physics possible at all.

If causes are real but never total, and every world keeps a residual that no law fixes, where does the new come from, and what is a choice?



Chapter 7

Choice, Change and Cause

Two Ways to Spend an Evening

Where does anything new come from, if causes do not fix everything? Chapter 5 showed that determinism, the claim that everything is fully and exactly fixed, cannot be had: it would need one law perfectly the same everywhere and every state laid out without a gap, both at once, and that pair is the one the Incommensuration Theorem forbids. Chapter 6 showed that causes are real all the same: a cause is a consistency that a finite comparer finds between two things in the world, one following the other, again and again. So the world is causal, and it is not determined. This chapter asks what it is to choose in such a world.

Someone has a free evening and a list: a film, a dinner with a friend, an early night. Weighing the three, giving up two to have one, is a *decision*. Then the friend calls with an idea that was on no list, a walk to watch the sun go down over the river, and the evening becomes something none of the options was. The friend's call, bringing something that was on no list, is a choice.

A decision picks from what is already laid out. A choice brings about something that was open but not yet laid out. In a world that is causal and not determined, the difference between the two is real, and the rest of this chapter says what it is. A deterministic world, the world Chapter 5 excluded, would have no such difference, since the list would be all there is.

The question is about you, the reader. Consciousness, in this chapter, is the condition from which you are reading these words, and never a theoretical category held at arm's length. In this chapter I tell decision from choice, and choice from creation. I show how choice and cause trade against each other at each scale. I also find in the self the same condition between two extremes that I found in the world in the earlier chapters.

Three Members of the Real

Go back to the evening. The friend's call, bringing something that was on no list, is the choice. The walk to the river, unfolding step by step, is the change. Seeing afterward that the call led to the walk is the cause.

Change, causality and choice together are what I call *the triple of the real*: the three members that everything real has. Chapter 4 named the three modalities, the three ways in which anything can be taken: the immanent (process, what is happening here and now), the omniscient (structure, what holds the same for everyone, laid out as a pattern), and the transcendent (possibility, what makes things possible and new). In the triple of the real, change is immanent, causality is omniscient, and choice is transcendent. Change is process: what is happening, the unfolding of experience from within. Causality is structure: it registers what is kept, what persists from one state to the next. Choice is possibility, and it belongs to the inside of the one who chooses: it reaches beyond what is given, and brings the new into the world from within the chooser's own standpoint.

In my book *An Immanent Metaphysics*, first published in 2001, I defined the three together, in the terms of comparison that Chapters 1 and 2 built. Change is "the continuity of content and asymmetry of context of the interaction and definition of the subjective and the objective". In plain words: a change is the meeting in which a self and its world are defined, going on without a break along an order that runs one way; it is the actual transition, the happening itself. Causality is "the subjective perception

(context) of a consistency between two objective contents". In plain words: a cause is a pattern relating events in the world, found again and again by a comparer from where it stands, and laid out as structure that others can share. Choice is "the objective expression (content) of a consistency between two subjective contexts". In plain words: a choice is something new coming out into the world, an outcome not drawn from earlier conditions alone, given from a stance.

Causality and choice mirror each other. In causality, a self perceives a consistency between two things in the world. In choice, a self expresses into the world a consistency between two of its own contexts. A self's own contexts are its own settings, the standpoints from which it compares, such as the person before the friend's call and the same person after it. The walk to the river is the new outcome that carries a consistency between those two settings out into the world.

Change is the most basic of the three. Chapter 4 gave the first of the three axioms: process, the immanent, is more basic than structure or possibility. Change is the immanent member of the triple, and without something actually happening, neither causality nor choice has anything to act in.

The three together are what all that is consists of, and its basis: distinct (each is its own thing), inseparable (none occurs without the other two), and not interchangeable (none can do the work of another). Whatever is real, being in an act, takes part in all three at once: it takes part in lawful pattern, it is in actual transition, and it holds something that what came before did not fix. Something that lacks one of the three is less than real. Without change, it is an abstraction. Without choice, it is a mechanism. Without causality, it is an event with no consistency that any comparer could find again.

The triple also sorts three words that everyday speech runs together, which Chapter 4 first set apart. To exist is to be describable as structure, with a place in a pattern; that structure is the same for every comparer, which makes existing an omniscient matter; and to exist involves no inter-

action. To be real is to be in an actual act of comparison, here and now, which is an immanent matter. To be objective is a matter of degree: objectivity is measured over many related acts, as the agreement of their findings across comparers, and it belongs to no single act, since a single act is real outright.

Why the Three Hold Together

Causality, change and choice hold together, and every attempt to break the three apart shows why the three hold together. The third axiom, as Chapter 4 gave it, says that the three modalities are distinct, inseparable and not interchangeable. So there are three ways to try to break the triple: collapse one member into another, separate two members, or exchange two members. Each of the three attempts fails.

Take collapse first. Choice cannot be collapsed into causality: a world in which every choice is only a cause is a world fixed in advance, which is determinism, and the theorem excludes determinism. Causality cannot be collapsed into choice either: a world with choice and no causality has no regularities, nothing to base a prediction on, and no way to communicate. Even a choice needs a causal setting in which the choice makes a difference. The friend's call would change nothing if no walk could follow from it.

Take separation next. Causality cannot be separated from change: a law that never governs an actual transition is a mathematical relation with no physical content. Change cannot be separated from choice: a change in which every step is fully fixed is only the unrolling of what was already contained, and the difference between a state and the state after it becomes a label, never a process.

Take exchange last. Choice and causality cannot trade places. To treat a creative act as a mechanism is to try to predict and control what is new by its nature. To treat a lawful regularity as a decision is to deny that nature can be relied on. Collapse, separation and exchange all fail, and so the triple stands.

One Order Only

The three run in one order, and in no other. The second axiom, as Chapter 4 gave it, is a one-way cycle, from possibility, to act, to record; the Incommensuration Theorem derives that axiom. In the triple of the real the cycle runs from choice, to change, to cause, and back to choice: choices start changes; changes can be understood by their causes; and patterns of cause set up the conditions for new choices.

Two everyday cases show the cycle turning. To pick up a cup, a person moves a great many muscles in exact time together, like the steps of a dance, and those many small selections give one result: the cup has moved. And for a film to be on the evening's list at all, there must already be roads, cars, theaters, films, actors, money, lights, lenses and chairs: a vast number of results of earlier choices, all in place, so that one possibility can be open. Many possibilities come before one selection, as the evening's options came before the walk: that is the step of choice. Many selections, the muscles, come before one result, the moved cup: that is the step of change, the act itself. Many results, the roads and the theaters, come before one new possibility, the film: that is the step of cause, whose settled patterns set up the conditions for new choices. The cycle moves from many to one, three times over.

This order is why causes run from earlier to later, as Chapter 6 found them doing. Causality runs from earlier to later because it is the understanding of changes that choices began. Reading the cycle the other way round, with cause first and choice last, leaves choice out of the cycle and

makes it only an effect of what came before. That backward reading is *reductionism*, the theory that explains every whole entirely by its parts and its earlier states; and reductionism is determinism in disguise.

Decision, Choice and Creation

With the order in place, the evening's two words, decision and choice, can be made exact, and a third word, creation, joins them. A choice is a process with three parts: what is possible, a selection, and a consequence. A choice is personal, the act of one self, and finite, bounded by the conditions of the one who chooses. A choice involves both the random and the dynamic: something that nothing before it fixes, and something that moves on by law. A choice sits between the completely random, which would be no choice at all, only noise, and the completely determined, which would be no choice at all either, only the carrying out of a program. Choice needs the gap between those two extremes, and the gap is guaranteed: Chapter 6 gave the guarantee its name, the non-polarization principle, which says that every world lies between the two poles and no world sits at either one.

A decision is a narrower act. A decision is a selection from a set of possibilities already defined, made by cutting away the others. Anything that involves trade-offs is a decision. A choice may realize a possibility that was open but was never in the set. Decision works with causality: it weighs what is known to follow from what. Choice works with creation.

Creation is choice's counterpart at the infinite scale. Choice selects among what is open; creation generates the range of what is open, the possibility that never runs out, from which new structure comes. Choice is bounded by the conditions of the chooser; creation has no bound. Choice is personal, the act of one self; creation is cooperative, many agents, processes and judgments of worth coming together into one shared outcome. At its limit, where creation would be perfected, it has the same form as the self taken whole: all of self at once, and not one self among others. And

there is a second triple, existence, interaction and creation, in which creation is the transcendent member: it holds there the place that choice holds in the triple of change, causality and choice. The second triple matters later in this chapter, because the four results near its end are drawn from two of its members, existence and creation.

A reader may ask, at this point: if I choose, am I responsible for what follows? Yes. Consequence is one of the three parts of a choice, and a choice without its consequence is not yet a choice. Picture a choice undone so completely that the clock rolls back, as if it had never been made. A choice undone that completely could not be told apart from no choice at all, and what cannot be told apart in any way at all is the same; so the undone choice would be no choice. A choice is not made until it is done. Because the world is causal, what is chosen goes on to have effects; because the world is not determined, the one who chose could have chosen otherwise.

Choice and Cause at Every Scale

Choice and causality are conjugate, in the sense Chapter 3 gave the word: the more fully one is held, the less fully the other is, and neither comes before the other. How the two trade against each other changes with scale.

At the smallest scale, the scale of single quantum events, the trade is strict opposition. John Conway and Simon Kochen proved a theorem in 2006, and strengthened it in 2009, which Chapter 9 sets out in full: if the experimenter's choice of what to measure is free, then the particle's response is not fixed by any information that came before it. In the theorem's own word, the particle is free. Chapter 9 also shows why the experimenter's choice is free. At this scale there is a real alternative, and no consistency that anyone could perceive in the single event.

At our own scale, the scale of a person in a room, both are present, and their proportion changes from act to act. A habit is a choice made so often that it runs almost wholly as causality: the hand reaches for the light

switch as you come through the door. A turning point in a life is causality bent almost wholly by choice. Most acts lie between, and the life of a person is the movement of that proportion.

At the scale of the whole cosmos, choice and causality become partners in proportion: how real either one is goes directly with how real the other is. Conjugacy is a trade in how fully each can be held; what becomes proportionate at the largest scale is how real each is. The whole, taken as a whole, shows both structure and openness, and the two fit together instead of opposing each other. I give the name *teleology* (from a Greek word for an end or purpose) to that proportion in the whole: the coherence that makes the cosmos a world, and not a collection of events. The coherence is real. A description of it in terms of purposes and ends would need the whole course of the cosmos in view, and so it is never complete for an observer inside the cosmos. At this scale the non-polarization principle is plainest: the cosmos is neither purely determined nor purely random, and the play of structure and choice is what gives it its character as a world.

Chapter 6 found that the universe has no cause, since a cause is a relation found by comparing events within a whole, and the whole has nothing outside it with which to be compared. That finding answers the old argument for a first cause at its premise, the premise that the universe, taken as a whole, needs a cause. What stands at the scale of the whole is teleology, the proportion of choice and causality in the whole, and creation, the cooperative making of what is possible.

Felt Cause Comes First

Chapter 6 told apart two kinds of causality: subjective causality, the consistency a comparer registers from within its own domain, and objective causality, the same consistency extended to the world as it is apart from any one perceiver. In a person the two take a familiar shape. Felt causality, the experience of one thing following another and of the world holding together from moment to moment, is subjective causality. Impersonal

causality, the statement that this caused that as a fact about the world, made as if from no standpoint at all, is objective causality. The felt kind is the foundation; the impersonal kind is drawn from it.

That the felt kind comes first matters for choice, because choosing is met first as an experience. A person can experience choosing only because subjective causality is real: the world, as perceived from within, has a causal structure that is not deterministic, and so it leaves room for the new. Both halves of that are needed: that the world as perceived from within has a causal structure, and that the structure is not deterministic. Without a felt causal structure there would be nothing for a choice to act on and no effect to follow from the choice; with a felt structure that was deterministic, there would be nothing open for a choice to meet.

The Self

Picture two people. One is on a crowded trading floor, answering every call. The other is on a long silent retreat, answering none. Both are selves: one is turned almost wholly outward, into the world's traffic, and the other almost wholly inward.

In *An Immanent Metaphysics* the self is "the appearance of the aggregation of that aspect of interaction which is choice". In plain words: every interaction has a side that is choice, and the self is what appears when that side of many interactions is gathered together. The self is one of the parts no choice can do without, in the way that content and context are parts no comparison can do without.

One principle shapes every self: the more it interacts, the less it resembles the non-physical, and the less it interacts, the more it resembles the non-physical. The non-physical, here, is the word in Chapter 6's sense: the side of possibility, where more than one outcome is available, and nothing ghostly. The self takes part in interaction: it perceives, it chooses, it is set within a domain. How much it interacts is never zero, since it is always engaged with something, and never infinite, since it is never the

same as the whole of what it interacts with. The more the self interacts, the more deeply it is set in the physical, and the more fully it engages the causal structure of the world. The less it interacts, the further it withdraws from engagement, and the nearer it comes to pure possibility. Neither end is ever reached. The trader still has an inward side, and the one on retreat is still engaged with something.

The principle that shapes every self is the non-polarization principle, applied to the self. The self is neither purely physical nor purely non-physical. The self is the condition between those poles, in which each side depends on the other, and that condition is what it is to be real, since to be real is to be interdependent. Consciousness is real; what is illusory is the demand to explain it by one of the impossible extremes, pure physicality or pure non-physicality.

The Four Theorems of Being in the Self

Two everyday facts, set side by side, show the theorem at work in the self. The laws of physics hold the same in every place and at every time, while every place and every moment can be told apart from every other. And the "I" who reads this sentence is the same "I" who read the last one, while what that "I" perceives and thinks changes all the time.

Chapter 2 met these as the four theorems of being, which I stated in *An Immanent Metaphysics*: "Subjective context is continuous: the persistence of self-identity. Subjective content is discontinuous: the diversity of thoughts and perceptions. Objective content is symmetric: the conservation of matter and energy. Objective context is asymmetric: the arrow of time". In plain words: the setting of the self is unbroken, so the self persists; what the self holds jumps, from one thought or perception to the next; what the world holds stays the same across settings, so matter and energy are conserved; and the setting of the world runs one way, which is the arrow of

time. In the world, content is lawful across changing context: laws that hold everywhere, applied in particular circumstances. In the self it is the mirror image: one lasting context, the self, across changing content.

The lasting of the self is given; nothing builds it. The lasting of the self is the ground of all experience, as experience is lived. Because the self persists without a break, one and the same self stands at the end of its past and at the start of its future. From within, the past shows itself as settled, held now only in memory, and the future shows itself as open, held in intention and deliberation. The present, where the two meet, is where choice is made, bringing about something of what is open.

The four theorems of being are one form of the limit the Incommensuration Theorem names. Experience is continuous in its context and discontinuous in its content: the self goes on without a break while what it holds jumps. Determinism would need something else from experience: content that stays the same across every change of setting, which is perfect symmetry, together with perfect continuity. No arrangement could make experience both perfectly symmetric and perfectly continuous. Pure indeterminism, which Chapter 6 called pure randomness, would require perfect asymmetry together with perfect discontinuity, and the theorem's second clause excludes it equally. Both are illusions, made by pushing one pole to its extreme; what is real is the condition between them.

One Rule, Four Doors

Can the theorem, bound directly to the ideas of existence and being, creation and doing, give results of its own? It can, by a method I call *domain binding*: the theorem's terms, symmetry and continuity, are bound to particular ideas of a field, and the theorem's limit is read in that field's own language. Four bindings give four results: that I am and I think; that I think, therefore I am; that measurement is the meeting of creation and doing; and that the scientific method is the necessary form of any process that relates what exists to what is done.

Chapter 3 pictured the theorem as a floor. A floor laid with identical square tiles repeats perfectly and always has seams; a ramp poured in one piece has no seams and never repeats itself exactly; no floor can be both perfectly repeating and perfectly seamless. The same rule can be read in any material: wood, stone, sound, thought. Each reading puts particular things in the places of "repeating" and "seamless", and the rule then says something about those things that was not obvious before. The four readings below are four doors opened by one rule. They differ from the four theorems of being, which state the theorem in the self directly; these four come from binding the theorem to four chosen ideas.

The theorem speaks in general terms: symmetry and continuity cannot both be perfected in any comparison in any field, and neither can asymmetry and discontinuity. Its power comes from that generality, and applying it takes three steps. First, name precisely which ideas of the field are its symmetry and its continuity. Second, show that the naming is correct, so that each idea really is an instance of the term it is matched with. Third, let the theorem's limit give the field's result, with no further reasoning special to the field. Each of the four arguments below takes those three steps.

Domain binding is therefore a derivation. A reader may ask whether a binding is only an analogy. An analogy notices a likeness; a binding identifies the field's ideas with the theorem's terms, shows that the identification is correct, and then lets the theorem do the work.

Four ideas carry the bindings, in the field of the self and the world. *Existence*, here, is what holds the same across instances, apart from the self; it always takes the part of symmetry, because a sameness of content across differences of context is exactly what symmetry is. *Creation* is the making of what is open, unbroken and without seams; it always takes the part of continuity, because a sameness without seams is exactly what continuity is. Being, what the self is, and doing, what the self does, are the partners of existence and creation in the bindings: each takes whichever part the other idea in the same binding leaves free. So being takes the part of con-

tinuity in one binding and the part of symmetry in another. The part an idea takes, symmetry or continuity, belongs to its binding, and not to the word that names the idea.

Here are the four pairings, said one at a time, with the part each idea takes. Existence with being: existence takes symmetry, so being takes continuity; the result is "I am and I think". Creation with being: creation takes continuity, so being takes symmetry; the result is "I think, therefore I am". Creation with doing: creation takes continuity, so doing takes symmetry; the result is measurement. Existence with doing: existence takes symmetry, so doing takes continuity; the result is the scientific method. A listener needs to hold only one rule to keep the four straight: existence is always the symmetry and creation always the continuity, and the other idea in each pair takes the part that is left.

I Am and I Think

An argument. This binding pairs existence, which takes symmetry, with being, which takes continuity. Take the self as a comparison in which the one comparing is also what is compared. "I am" is the self's existence: the same self across different moments and places, a sameness of content across a difference of context, which is symmetry. "I think" is the self's being, taken as what the self is while it goes on, its activity: the unbroken going-on of thought, each moment passing into the next with no break, which is continuity. Every comparison has content and context and finds sameness in each to some degree, so both characters are present in any self. By the theorem, the two cannot both be perfected. Perfected on the side of existence, the self would be all sameness and all seams: a fixed structure with no activity, like a statue of a person. Perfected on the side of being, it would be all flow and never the same: a flow with nothing that stays the same, an activity that is no one's. Neither is a self. A self is therefore both, each to a degree and neither absolutely, and the two cannot be

separated absolutely. The fact that I exist and the fact that I think are two sides of one condition that cannot be taken apart; they were never two independent facts.

The binding of existence and being states that existence and thought go together. The famous saying of Descartes runs in one direction, from thought to existence, and the next binding gives that direction.

I Think, Therefore I Am

An argument. This binding pairs creation, which takes continuity, with being, which now takes symmetry. Thinking still sits on the side of continuity, as it did in the last binding; what has changed is the part that being takes. In this binding, thinking is creation: the continuing making of thought after thought, unbroken, which is continuity. The thinker's being is symmetry: the sameness of the "I" across those thoughts. A symmetry is a sameness of content across a difference of context, and here the differences are the passing thoughts themselves; so the sameness of the "I" is found only by comparing thoughts, which is to say, it is read from the thinking. The thinking needs no "I" already in place in order to occur: perceiving is more basic than the one who perceives and the thing perceived. By the theorem the two are never perfected together, so the "I" is never a structure complete in itself, apart from the thinking; it is always the sameness found across an activity. Hence the direction: from thinking, a thinker. "I think, therefore I am".

The conclusion is the one René Descartes reached in 1637 and again in 1641, and the ground on which I reach the conclusion is new. Descartes derived existence from doubt: whatever he doubted, the doubting itself was thinking, and so a thinker was there. This derivation reaches the same conclusion from the theorem's limit on comparison. Descartes sought what could not be doubted even in principle, and stopped at the thinking

self. He stopped too early. Doubt is a process, and a process has the parts of comparison; as Chapter 3 found, doubting comparison is already an act of comparing. One cannot doubt comparison without comparing.

Measurement: A Caused World and a Chosen Event

In the binding of creation and doing, the theorem yields measurement, and its result fits in one line: a caused world and a chosen event yield a chosen world and a caused event.

In plain words: before a measurement, the world is caused, since its earlier states fix, by law, how likely each possible outcome is; and the event of measuring is chosen, since the experimenter chooses what to measure, for instance which way to set the detector. After the measurement, the world is chosen, since one actual outcome has been singled out from among the possible ones; and the event is caused, since the reading follows from the interaction. Choice and causality exchange their roles across the act.

The exchange yields something in the world and something in the one who chose. In the world, it yields one actual outcome, with the information that outcome carries. In the one who chose, it yields knowledge of the chooser: once made, the choice stands among the actual things the chooser can perceive, and it is a fact about the one who made it. We learn about ourselves through the choices we make.

An argument. Doing, here, is symmetry: the lawful procedure of measuring, the same operation against the same standard, whatever the occasion. Creation, here, is continuity: the field of what is possible, smooth and unbroken before anything is found. By the theorem the procedure cannot be perfected together with the field: a perfectly lawful procedure cannot fix completely what an unbroken field of possibility will yield. So every measurement ends in one actual outcome that the procedure did not fix: what was possible becomes actual. Measurement is the act in which creation, the bringing in of the new, and doing, the carrying out

of a lawful process, meet. The act is creative, since something new is brought into being, and causal, since it follows from the structure of the interaction. The roles can exchange because choice and causality are conjugate with respect to change: choice goes with creation, possibility and continuity; causality goes with existence, actuality and symmetry.

A reader who remembers Chapter 6 may ask how causality can stand with symmetry here. Chapter 6 kept two words apart. Causality is the pattern a comparer perceives, and in this binding it stands with existence and symmetry, as what is kept across a transition from one state to the next. Causation is the one-way order at each event, and it has one of the two pairings the theorem keeps, which Chapter 3 named: continuity with asymmetry, an unbroken connection along an order that runs one way. The two statements are about the two things the two words name, and each statement holds in its own place, because a part belongs to its binding and not to the word.

The Scientific Method

In the binding of existence and doing, the theorem yields the scientific method: the way an observer comes to know a domain, by observation, hypothesis, test and revision.

An argument. Existence, bound as symmetry, is what holds the same across instances: a law. Doing, bound as continuity, is the unbroken course of an inquirer's action. To join the two completely would be to know the law in one continuous act, finished; the theorem forbids a perfect symmetry together with a perfect continuity, so that act is impossible. What remains possible is a cycle of separate acts, each compared with the last: an observation, which is a finding; a hypothesis, which is a sameness proposed across findings; a test, which is a new finding compared with the proposed sameness; and a revision. The cycle never closes, because closing it would be the forbidden perfection. The scientific method is there-

fore a structural necessity, and no convention: it is what the theorem requires of any process that seeks to establish how what exists relates to what is done. The scientific method is made of comparison.

What the Four Show

The four results are familiar: existing and thinking go together; thinking implies a thinker; measurement turns the possible into the actual; science is the method of knowing. Derived from one principle, they are consequences of one structural fact, that symmetry and continuity cannot both be perfected together, and they are four faces of that fact, never four separate observations.

Each binding works because the world it binds is causal without being deterministic: the self is never a finished structure, the measured world is never fixed before it is measured, and knowledge of a law is never completed in one act. In a deterministic world all four results would fail together, and in a purely random world they would never begin.

From Outside, Chance; From Inside, Choice

Suppose I choose to raise my left hand and not my right, and suppose the choice is a real one, fixed by nothing before it. To someone watching from outside, the raised hand looks exactly like a random event. No outside measurement can tell the two apart.

Chapter 5 showed that no world is purely and completely deterministic: whatever lawful description is given of a world, something happens in that world which the description does not fix. Chapter 6 named what is left over, what no lawful description fixes, the residual; the theorem guarantees a residual in every description, and Chapter 8 follows it into the laboratory. From outside, in the third person, as one describes someone else, all that can be said of the residual is that certain outcomes cannot be predicted. An observer outside, with the best theory and the most com-

plete data, cannot tell whether an outcome no one predicted is random in the fullest sense, without cause or significance, or a choice, specific and significant to the process it came from. From outside, the two look the same.

Watching people gives the same result. People are fairly predictable; a good deal of what a person will do can be foreseen. Watch someone for a long time and write down everything about them that can be predicted, the way a program that shrinks computer files writes down the patterns in a stream of data, so that the stream fits in less space. However good the program, some part of the person's behavior will not shrink, and that leftover cannot be told apart from randomness. That leftover is the part the person regards as their own choosing.

The two look the same because they are the same event, described from two positions. I call the passage from one position's description of an event to another's a *modal rotation*: the event stays as it is, and the position from which it is described changes. A modal rotation rests on one fact: a role belongs to a position of description, and never to a thing. (The Appendix sets it beside a second relation that rests on the same fact.)

From the omniscient position, the position of structural law and of description in the third person, the residual shows up as a gap without law, which cannot be told apart from ignorance: it has no pattern, and it is the gap in the specification. From the omniscient position the residual is without meaning, because the omniscient position describes what is the case and does not weigh meaning. From the transcendent position, the position of weighing worth and of engagement in the first person, the same residual is choice: the specific outcome that was not fixed in advance, carrying the character of the process it came from. From that position it is meaningful, and the meaning is the transcendent description of the same event, never a further property laid on top of the randomness.

Each of the two descriptions is correct in its place, and neither can replace the other. The outside description works in the terms of symmetry: it describes the residual as a place where symmetry is missing, where con-

tent fails to stay the same across a change of context. The inside description works in the terms of continuity: it describes the residual as the unbroken engagement, in the first person, with one actual outcome. The two are conjugate, and the theorem forbids making both perfect in one description, so they cannot be joined into one.

So an event can be unpredictable and meaningful at once. A process can be unpredictable to every other process, because nothing outside it fixes the outcome, and not because the others are ignorant, and still be meaningful to itself. Its meaning is one with its being that choice, and the meaning cannot be extracted, passed on in its particulars, or reproduced by any process other than the one that is the choice. From outside, randomness; from inside, choice: one event, two descriptions, both correct, neither complete.

Meaning Cannot Be Removed

Meaning belongs to the structure of the residual, and nothing can remove it. Meaningful, here, means significant to the process an outcome came from: the outcome carries that process's character, and it has a share in value. The argument has three steps.

The first step asks where meaning could stop. Suppose meaning is present in one event and absent from another, and the two are joined by a chain of interactions, as all the events of one connected world are. Then either the chain has a sharp edge somewhere, where meaning stops, or meaning carries along the chain. A sharp edge would need a rule that says where meaning stops, and drawing that line is itself an act of comparison, since every distinction is made by comparison. A comparison is made from a stance, and making it is a meaningful act, so no edge can be drawn from a position outside meaning. The same holds of meaninglessness, taken as a definite condition and not as a mere absence. So either every comparison in the connected world is meaningful, or none is.

Second, meaning has one case that cannot be removed. The idea that choice is meaningful is itself an idea with definite content: it identifies the residual, seen in the first person, as a specific outcome carrying the character of its process. That idea is meaningful, so it is itself a case of meaning. It applies to itself. A thing that applies to itself in this way is called a *fixed point*: it is itself a case of the very thing it is about, and no description can take that case away.

Third, every attempt to define meaninglessness is itself meaningful. A definition of meaninglessness that had content would be meaningful, and so it would not be a case of what it defines. To say what it is for something to lack meaning is to compare, from a stance, and every comparison made from a stance is a meaningful act; so the definition is itself meaningful. So meaninglessness has no case of its own.

So the either-or of the first step comes out one way: meaning is present wherever comparison occurs, and it cannot be removed. The claim is that the residual has a character of value that no description can strip from it; it is not the claim that every event has a significance someone could write down. *Nihilism*, the view that nothing is meaningful, would need the opposite fixed point, a meaninglessness that is a case of itself, and that fixed point does not exist. The act of declaring that nothing matters is an act of weighing worth, and so it is itself a case of mattering.

Three things the theorem guarantees now join. The theorem gives the residual, and with it the impossibility of determinism. The one-way order of comparison, act before record, gives the direction in which the residual unfolds; Chapter 13 shows that this direction is built into every single comparison, and calls it the structural arrow. And meaning, which cannot be removed, gives the residual its character of value. Together they give an account of the real in which law, process and significance are all present, none can be removed, and all are grounded in one theorem. The three, joined, answer this chapter's question of where the new comes from: the residual is where the new comes in, the arrow is the direction in which it unfolds, and meaning is the character of value it carries for the one from whom it came.

Where the New Comes From

We now know that the real is the triple of causality, change and choice: pattern that can be found again, the happening itself, and something new that nothing before it fixed. The three are distinct, they never occur apart, and none can do the work of another; and they run in one order, from choice, to change, to cause, to new choice.

The triple answers the chapter's question. The new comes from choice, the member of the triple that reaches beyond what is given. A world without choice would be a world in which choice had been collapsed into causality, and that world is determinism, which Chapter 5 excluded: determinism would need one law perfectly the same everywhere and every state laid out without a gap, the pair the theorem forbids. So choice belongs to everything real, as fully as cause does.

Summary

The claims of this chapter are claims of the structure of comparison itself, argued from the Incommensuration Theorem and the three axioms and from nothing else: the triple of the real, which holds because the third axiom makes its members distinct, inseparable and not interchangeable; the four results of domain binding, each a derivation from the theorem and not an analogy; and the two consequences of the theorem with which the chapter closes, that choice and randomness are one event described from two positions, and that meaning cannot be removed from that event.

The real is the triple of causality, change and choice. Causality is a pattern relating events, found again and again by a comparer and laid out as structure that can be shared; choice is the coming of something new, an outcome not drawn from earlier conditions alone, given from a stance; change is the actual transition, the happening itself, and the most basic of the three. Causality is omniscient, change immanent and choice transcendent; the three are distinct, inseparable and not interchangeable, and

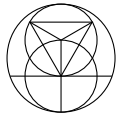
every attempt to collapse, separate or exchange them fails. Whatever is real takes part in all three of causality, change and choice. Without change, what is real would be an abstraction, and without choice, a mechanism. Without causality, what is real would be an event that no comparer could find again.

The three run in one order, choice to change to cause to choice, and that order is why causes run from earlier to later. A decision selects within a set already laid out; a choice realizes what was open but not in the set; creation, cooperative and without bound, makes the range of what is open. Choice and causality are conjugate: strict opposites in a single quantum event, mixed in changing proportion at the scale of a person, and partners in proportion at the scale of the cosmos, where teleology names their proportion in the whole. Felt causality is the foundation from which impersonal causality is drawn, and it leaves room for the new.

The self is the condition between the purely physical and the purely non-physical, resembling the non-physical less the more it interacts; consciousness is real. The four theorems of being show the theorem in the self: one unbroken self across changing content, mirroring lawful content across a one-way setting in the world. Bound to existence and creation, being and doing, the theorem gives four results: I am and I think; I think, therefore I am; measurement, in which a caused world and a chosen event become a chosen world and a caused event; and the scientific method, as the necessary form of any process that relates what exists to what is done. All four rest on a world that is causal without being deterministic.

What no lawful description fixes is, from outside, randomness, and from inside, choice: one event under a modal rotation, two descriptions, both correct, neither complete. Meaning is present wherever comparison occurs, and it cannot be removed; nihilism has no fixed point to stand on. Law, process and significance are all present in the real, none can be removed, and all are grounded in one theorem. The scope of these results is every act of comparison, at every scale where comparers are.

From outside, a choice and an outcome that nothing fixed look exactly alike. So what does an outcome that nothing fixed look like when it happens in a laboratory: is it only the name for what we do not yet know, or is it fixed by nothing at all?



Chapter 8

A Click No One Could Predict

Two Meanings of Random

When we call something random, are we only saying that we do not know what it will do? Or can an outcome be open in itself, with nothing anywhere that settles it before it happens? The word ‘random’ carries two meanings, and everything I argue in this chapter turns on keeping the two meanings apart.

Start with four ordinary things we call random. The toss of a coin. The weather next month. The next digit of a long decimal that no one has yet worked out. The card at the top of a deck that has just been shuffled. In each case there is, or might be, a fact of the matter that fixes the outcome in advance; we do not have access to it. Better instruments, more computation and a fuller theory would in principle reduce the randomness, perhaps to nothing. I call this *soft randomness*: unpredictability that comes from a limit of what we know, and that better knowledge could shrink. Soft randomness is a statement about our knowledge.

The second meaning is stronger. A thing is random in the second sense when there is nothing from which it could have been predicted. No knowledge of causes, however complete, would have yielded the outcome, because the outcome was not contained in anything that came before it. Better instruments do not help, because there is nothing there for them to find. A fuller theory does not help, because the part of the outcome that

no knowledge predicts is a feature of what the theory is about, and the theory has no gap that a better theory would fill. I call this *hard randomness*: an outcome that no description of what came before fixes, however complete that description is, for any comparer whatsoever, even one with unlimited access to the state of things. (A comparer, as Chapter 1 said, is anything that makes a comparison: a person, an instrument, or an atom that absorbs light.) Hard randomness is a statement about the world.

The word hard has one specific meaning. It says that there is no possibility, not now and not ever, of turning what is left over into knowledge of causes, or into any other kind of prediction that logic could make. Random, in this sense, is what remains when no account by causes can account for a thing at all, even in principle. Any account of causes that could be brought to bear on it would make it not random. Picture a program that shortens a long file by finding every pattern that repeats in it and writing each pattern down only once. When the program has found every pattern there is, what remains has no pattern left to find, and nothing can make it shorter. Hard randomness is like that remainder: what is left after every regularity (every pattern that repeats) has been taken out, and what no further taking out can shrink. I call it the *incompressible residue*. The name sounds close to a name from Chapter 6, the residual, and the two are kept apart this way. The residual is the part of what is found that no lawful description fixes, and the Incommensuration Theorem guarantees that every description leaves such a part. The incompressible residue is that same residual as it is met in the findings, the part that is left when every pattern in them has been taken out. So the residue is the residual, found.

Hard randomness is real, and the single outcome of a quantum measurement is its plainest instance. Quantum theory is the theory of matter and light whose findings come in countable steps. In a quantum measurement, a count arrives in one detector and not in the other. The state of the particle before the measurement, known as fully as quantum theory allows, gives only the probability of each outcome: the share of many

identical runs in which that outcome will come up. It gives nothing that fixes which outcome comes up this time. That count, the click in one detector and not the other, is the click no one could predict.

A third idea goes further than hard randomness, and it is kept apart from both soft and hard randomness. Chapter 6 named it pure randomness: a content found different whatever its context, an outcome with nothing at all accounted for. Pure randomness is an idealization, as determinism is. Recall the four ways of being the same or different from Chapter 2: *symmetry* is sameness of content across a change of setting, *continuity* is sameness of content within one setting, with nothing jumping, *asymmetry* is difference of content across a change of setting, and *discontinuity* is difference of content within one setting: a jump, a click, a step. Determinism would be symmetry and continuity, both perfected together. Pure randomness would be asymmetry and discontinuity, both perfected together. The Incommensuration Theorem forbids both: symmetry and continuity cannot both be applied, at once and absolutely, to one comparison, and neither can asymmetry and discontinuity. The theorem is a structural fact about the nature of knowing. So of pure randomness and hard randomness, only one is real. Pure randomness is excluded, as determinism is. Hard randomness is real. Every actual world lies between the two poles of determinism and pure randomness.

A Line With Two Ends

It helps to see hard randomness against a range. Picture a line drawn across a page. At the left end stands determinism: nothing unaccounted for, every outcome fixed by what came before, always and completely. At the right end stands pure randomness: nothing accounted for, no outcome fixed by anything before it. Both ends are abstractions. Between them lie the two ideas that describe the world we live in.

The first of the two ideas that describe the world we live in is causality. Causality, as Chapter 6 defined it, is the condition in which most things can be predicted most of the time: a consistent pattern of relations between events, perceived by an observer, holding to a degree. The second of the two ideas that lie between the ends is choice. Choice is the condition in which some part of what happens can be predicted from what came before and some part cannot; the part that cannot be predicted is the contribution of an agent, and it is something other than noise. Change, the larger class that holds both, lies between them on the line. Change holds the two as a whole holds its parts: some change is accounted for by causes and some by choices, so on the line it sits between the two parts it holds. Causation is mostly predictable, choice is mostly unpredictable, and change, which holds both, is about half predictable. "About half" names a place on the line, between "mostly" and "mostly not"; it is no measured amount.

Chapter 5 showed that the left end, determinism perfected, lies beyond what comparison can reach. This chapter is about what comes into the world from the side of the line where pure randomness stands: hard randomness, which is real.

Causation Is Not Determinism

Hard randomness is hard to see clearly because of a habit of thought that is nearly universal among people who work with mathematical models, and the habit needs a name before going further.

Chapter 6 traced the habit in three steps: change is collapsed into causation, then choice is collapsed into causation, and then causation is identified with determinism. As the three steps bear on randomness, the second step makes choice a name for a causal process we happen not to have traced, and the third treats the pattern of causes and effects as an exact copy of the mathematical rule that takes a complete state of the world to the next complete state. At the end of the three steps only one idea is left,

determinism, and hard randomness has no place to stand. A residue that no cause accounts for has been defined out of existence before any experiment is consulted.

Each step is a mistake, and the third is the one that matters here. Determinism is perfectly applied. It is a mathematical idea, an ideal form of the kind Plato described: a rule that holds without exception, without time, and without an observer. Causation holds to a degree. Causation is a pattern seen from inside time, from a standpoint, by a comparison. Causation holds as far as the pattern repeats, and cannot be perfectly applied. Determinism is atemporal: it stands outside time. Causation is temporal: it happens in time.

One can be entirely comfortable with a rule that stands outside time inside a model, and most physicists are; the mistake comes when the model's timelessness is carried across to the world. The result is a position that can be logically consistent and philosophically inconsistent at once. In Chapter 5's two words, the position the habit produces is valid but not sound: its conclusion follows from its premises, and its premises are not all true. The premises of the position include an identity that does not hold: the claim that causation is determinism.

John Bell showed in 1964 that no account in which every outcome is fixed in advance by hidden local causes (causes that reach each place no faster than light) can match all the predictions of quantum theory, and the measurements made since then side with quantum theory. Chapter 10 walks through his result step by step; here the reader can take it as physics has established it, and we need one thing from it. Bell's result shows that causation and determinism are two different things. It shows that causation cannot be perfectly applied to the correlations of separated particles. A *correlation* is an agreement between two sets of results that shows up again and again. There is more correlation in nature than any coherent account of causes can explain. Determinism is perfectly applied by its definition, and causation, by Bell's result, cannot be; so causation is

not determinism, whatever the models say. With a strict difference between the two, there is room for a residue that is causal in no sense, and that residue is what hard randomness names.

Three Results Read Together

Three results of physics are each well understood when taken separately. Taken together, they establish that hard randomness is a real idea: one that cannot be removed from a theory without the theory ceasing to agree with experiment. The three are Bell's theorem, the uncertainty principle that Werner Heisenberg published in 1927, and Planck's constant.

Bell's theorem shows that there are correlations in nature that no cause accounts for. Separated particles agree in ways that neither a shared past nor an influence passing between them can account for, and the agreement cannot be used to send a signal. The correlations are real; they are observed after the fact, when the two sets of records are brought together; and no causal process accounts for them. So the first result, from Bell's theorem, is that there is something in nature that causality does not cover.

Heisenberg's uncertainty principle shows that there is a limit of access. It says that certain quantities come in pairs whose two members no measurement can make sharp at once. Take a property of a particle called its *spin*, which behaves in some ways like a small arrow fixed to the particle, pointing in some direction. Picture the arrow at some angle to the measuring apparatus. The angle, if it had a definite value, would be perfect information about the particle. No measurement returns the angle. A measurement asks a single question, vertical or horizontal in the frame the apparatus sets, and receives a single answer, and the act of asking turns the arrow it asked about. One cannot measure the angle both ways on the same particle. One cannot get around this by preparing a pair of particles whose results agree, and measuring one way on each: the pair does not route around the limit, because the measurements taken at right angles to each other come out with no correlation at all. The dodge would need the

partner's answer, asked at a right angle, to say how the first particle would have answered at that right angle; with no correlation between the two answers, the partner's answer says nothing about the first particle, and the angle stays unfound. The angle is unknowable, and it stays unknowable however good the technology becomes. The limit is categorical: beyond it lies only the unknowable, it is written into the mathematics, and no improvement in engineering moves it. So the second result, from Heisenberg's uncertainty principle, is that there is a boundary past which measurement does not go.

Physicists call the mathematical description of a particle, from which they compute the chance of each outcome of a measurement, its *quantum state*. I will add here a reading of my own. The quantum state is a record: the record of what has been found, and of what a given setting of the detector makes possible, and it pictures no physical thing. That record is the same for everyone who has the outcome. It is nobody's private degree of belief, and it is no knowledge of a hidden angle, since there is no such angle. The act of measuring is an interaction: it makes one outcome actual, and the record of what can be known is updated to that outcome. The update is not a physical change of the system measured. The mathematics tracks what can be known by an interaction of a given kind, and where the mathematics says knowledge stops, the mathematics is telling us about interaction, and not only about our ignorance. This account of the quantum state is my interpretation, the frame in which I find the argument most natural; the argument of this chapter stands on the three results without it.

Planck's constant measures the size of the limit that the second result names. It is the number of physics that marks the microscopic limit, the scale below which measurement gives out. Its numerical value expresses the definition of a unit: it has been fixed by the definition of the kilogram since 2019, and it is quoted from physics, not obtained from anything in this book. What it measures is a bound on knowing two things at once. The pairs Heisenberg named, such as where a particle is and how it is moving, are *conjugate*, in Chapter 3's word: each of the two can be made

sharp only at the cost of the other. Planck's constant is an amount of *action*, a quantity physics uses that combines energy with time, or motion with distance. Below the action it marks, no measurement can pin down both members of a conjugate pair together as finely as a full description of them would need. A reader may picture that limit as a smallest grain of space and of time. Planck's constant sets no least length and no least duration: how fine a comparison can be made depends on the energy brought to it. Heisenberg says that there is a boundary; Planck's constant says how wide the uncertainty is at the boundary. So the third result, from Planck's constant, is that the boundary has a scale.

Now put the three together. Bell says there are correlations that no causal process accounts for. Heisenberg says there is a limit past which no measurement obtains information. Planck's constant says how much is unobtainable. Each of the three is an observation, supported by experiment and confirmed again and again. Together they mean that there is no coherent account of causes that can explain all the results of measurement. Some results are not explained by causes. A result that no cause explains is random, and since the failure of causation here is categorical, built into what measuring is, rather than practical, the randomness is hard. Once hard randomness is asserted on the basis of actual experiment, it can no longer be thrown away. It cannot be taken out of any theory that means to agree with what is measured.

The Tests and the Razor

Physics accepted hard randomness because experiment forced physics to accept it. The history of the response to Bell's result shows the idea being forced on a community that did not want it.

Bell's result sets a ceiling, called *Bell's inequality*, on how strongly the results at two separated detectors can agree if each result is fixed by causes local to it. When the first experiments found agreement above that ceiling, the response was that something must be wrong with the experi-

ment. So the experiment was done again, more tightly. Objections were raised. The detectors might be too inefficient, so that the pairs they caught were an unrepresentative sample. The settings might be chosen too slowly, so that a signal moving at the speed of light could have carried one setting to the other side. The settings might not be independent of the source of the particles. Each objection named a *loophole*, a gap through which an account by causes might after all have slipped. Experiments were then designed to close the loopholes, and in 2015 three teams of experimenters, led in their published reports by Hensen, by Giustina and by Shalm, closed them together. Their detectors were far enough apart, their settings were chosen fast enough and randomly enough, and their detection was efficient enough, that no account by causes through any of the named channels remained. The inequality was still violated.

At that point only two options remained. One option was superdeterminism, which Chapter 5 met and Chapter 11 gives in full: the position that everything is fixed at literally every scale, including the settings of the detectors, the random-number generators that chose them, and the experimenters who built the generators, with the large arrangement of the laboratory and the small state of the particle source fixed together from one common past. The other option was to accept hard randomness as a basic feature of reality.

Chapter 11 gives a structural argument against superdeterminism: superdeterminism is determinism made larger, and the theorem excludes determinism, and superdeterminism with it. That argument is not needed to choose between superdeterminism and hard randomness. What superdeterminism has to assume, once the tests reached the point they reached, is so extreme that an old rule settles the matter on its own. *Occam's razor* says: of two accounts of the same facts, prefer the one that assumes less. Weigh the two. One hypothesis says: the whole of the past, at every scale, has been arranged so that the free choices of experimenters happen to line up with the states of particles, to produce exactly the correlations that an account with no causes predicts, while leaving no trace of the arrange-

ment anywhere else. The other hypothesis says: some of what happens is not caused. Of the two hypotheses, the second is simpler by a very large margin. We accept hard randomness. We live with it.

One more attempt to escape the idea deserves a direct reply, since it gave the field its vocabulary. The hypothesis of *hidden variables* was the idea that something in the source of the particle pairs, something we do not know but that has a definite state, might account for the correlations. If both particles carry the answer from the start, their agreement is no mystery. The account in which both particles carry the answer from the start is the account by causes through a shared origin. Bell's inequality is the ceiling on the account by causes through a shared origin, and nature exceeds the ceiling.

The point I want to make here comes before any arithmetic. Hidden variables were proposed as the ground of causation: the definite state from which each outcome follows. When one uses the hypothesis as a model to predict the outcomes, one finds that causation has still been violated: the correlations still exceed what any definite state at the source can supply. And once causation has been violated, it does not matter what the variables are. A variable is a placeholder for a cause, and a placeholder cannot restore causation after causation has been shown to be insufficient. The hidden variables program was a temporary argument; it was tried, and it failed. The failure of the hidden variables program is what physicists mean by "no local hidden variables": no variable could have accounted for the correlations, which is a stronger finding than a search that came up empty.

One Atom, One Avalanche, Three Balls

Arguing that hard randomness is required is one thing, and pointing at hard randomness is another. Two examples, one direct and one indirect, show where hard randomness enters the world of ordinary scale.

The direct example is atomic decay. The *nucleus* of an atom is its small, heavy center, and some nuclei are unstable: sooner or later each one breaks apart, or *decays*. Take a large batch of such nuclei. The rate at which the batch decays is among the best-measured regularities in physics; physicists state it as a *half life*, the time in which half of a large batch will have decayed. The moment at which one single nucleus decays is predicted by nothing. There is no internal clock we have failed to read, and no hidden state winding down. The batch keeps its half life; the single atom decays when it does. The statistics, the pattern over many cases, are lawful, while the individual event is not. Atomic decay is hard randomness in its purest observed form: an event that could not be predicted on the basis of any past knowledge, because there is nothing from which a prediction of it could be made. Atomic decay is the clearer of the two examples because there is no hidden cause to appeal to.

An event at the scale of atoms produces a result at the scale of laboratory instruments by amplification. A detector is built so that one microscopic event sets off a cascade. In a Geiger counter, a single charged particle knocks an electron loose from an atom of gas, an event called an *ionization*. The freed electron knocks others loose, and they knock loose more, in an avalanche of further ionizations that ends as a pulse of current, which the counter registers as a click.

The indirect example is the *three-body problem*, the old problem of predicting how three bodies move when each acts on the other two, and through it the weather. Take three balls on a table that come together and touch at the same instant, and bounce apart without losing energy. Which way does each go? The outcome depends with extreme sensitivity on exactly how they came together: on differences in position and timing that can be made as small as one likes while still changing the result entirely. This sensitivity is no special case. A great deal of what happens in nature magnifies very small differences, and the weather is the familiar instance: it is so sensitive to where it starts that long-range prediction fails, however good the model.

Now follow the amplification downward. The differences that decide the outcome can be made smaller and smaller, and at some point they fall below the limits of measurement, below the microscopic limit. Picture a list of every state the system could be in, the *space of states*. At the microscopic limit that list is divided into cells, and differences inside one cell are not distinct states. The reason is the limit itself: two states that differ by less than the limit allows cannot be told apart by any interaction at all, and, by the principle of identity of Chapter 6, what cannot be told apart in any way is the same; so the two are one state, one cell of the list. At that point the *microstate* of the system, which is the exact state of the system down to the smallest parts of the system, cannot be known. The reason is not that our instruments are poor, but that no instrument for knowing the microstate can exist. And yet the system still produces an outcome. What produces it? Whatever is happening at the scale below the boundary of measurement; and that, from the standpoint of anything inside the range where measuring works, is hard random.

Atomic decay below the limit, amplified by sensitive processes above it, is happening all the time and everywhere. Add the correlations Bell found, which no cause accounts for and which appear in ways no one could have expected in advance, and the picture is of a world in which the incompressible residue is not confined to the laboratory. It seeps upward, the way groundwater rises wherever it finds a path, through every process that amplifies what it cannot resolve.

An Edge Built Into Measuring

I have been calling the microscopic limit a boundary, and whether hard randomness is hard, and not soft, depends on what kind of boundary the limit is.

Physics works inside the *domain of comparison*: the range within which comparisons can be made. The domain is specified by measuring. To be in the domain is to be measurable: to be able to enter an interaction that

yields a signal an observer can compare with another. A *signal* here is something that passes from one event to another and lets a record at the second depend on what happened at the first. The microscopic limit, marked by Planck's constant, is where no interaction yields such a signal: below the action it marks, no measurement can pin down both members of a conjugate pair together as finely as a full description of them would need. What lies there is out of reach of every measuring tool: of the tools we have, and of any tool that could ever be built, because a tool is itself an interaction, and the limit is a limit on interaction. The boundary is therefore *categorical*: it marks where measurement itself gives out, and with it the domain; it is not a mark of where our technique gives out.

There is a second boundary of the same kind at the other end, and seeing the two together makes the idea of hard randomness fall into place. At the large-scale edge of the domain there is a *horizon* past which no signal can reach us: whatever is farther away than light has had time to travel is unseeable, in principle. That horizon is the edge of what one observer can ever compare with, and not an edge of the universe. The present expansion of the universe began thirteen point eight billion years ago. That age is measured: it comes from the faint glow left over from the early universe, the *cosmic microwave background*, as the Planck Collaboration reported its measurements in 2020. Light from beyond the *observable universe* has not had time to reach us in those thirteen point eight billion years, and light sent out now from beyond another edge, the *cosmological event horizon*, never will. Zero measurable signal comes from too far away. In exactly the same sense, zero measurable signal comes from too small. A boundary, in this sense, means no coherent signal, and from beyond either edge there is no meaningful measurement. Chapters 12, 18 and 19 walk both edges in full; here we need only the parallel.

Now place the two ends of the line from the start of this chapter at the two edges. Determinism describes what lies beyond the large-scale boundary: the completed, the fully specified, the whole seen at once. The whole seen at once is the ideal of the view from everywhere, which we met in Chapter 5, and no signal from inside the domain can reach that

ideal. Determinism, if it could hold, would have to be made good all the way down to the absolute microscopic, the scale of exactly zero. Pure randomness describes what lies beyond the microscopic boundary, that from which no signal reaches inside. There randomness would be, in a strict sense, perfected, and so pure randomness is an abstraction, as determinism is. "Beyond", said of determinism and of pure randomness, means where thought places an idealization, by pushing one side of the real past the edge; it names no place that anyone could reach. Hard randomness is what comes through the microscopic boundary. Anything that comes through the microscopic boundary of the domain can appear only as random from our side, because no signal could come up through it carrying the information that would make it appear otherwise. The word hard in hard randomness says that the microscopic boundary is categorical: measurement itself gives out there. There is no possibility, ever, of turning what arrives through it into knowledge of causes, because knowledge of causes is made of signals, and the boundary is defined by their absence.

That the boundary is categorical has a consequence for what we can assume. Can we assume that there is something beyond the boundaries? No. For some boundaries "beyond" is not defined at all. The limit that the Incommensuration Theorem draws is such a boundary. The horizon of one observer is not: other observers stand past it, with horizons of their own. The limit of comparison itself has no far side for anyone. It is like the curving of a field itself, rather than a wall standing at the edge of a field that would otherwise go on: there is no far side of it to stand on. Where a beyond is defined, as past the horizon of one observer, we still cannot assume anything about what is there, because there is no way to tell. To tell, one would have to interact with it, and an interaction would produce an outcome we could register, a signal. There is no signal. So there is no hypothesis to be made about what lies beyond, and a hypothesis offered anyway is a pretense of knowing what the boundary itself makes unknowable. What is outside the limits of the domain is unknowable. That unknowability is the whole of what can be said about what lies outside the limits of the domain. I say no more, since saying more would claim

knowledge of what cannot be known, and a claim to know what cannot be known is an inconsistency in knowing. Chapter 20 takes up what physics cannot reach as a subject in its own right.

I lean on Planck's constant, and one distinction says exactly what in the constant I use. Planck's constant and the microscopic limit differ as a number differs from a reading of it. The number of Planck's constant has been fixed by the definition of the kilogram since 2019, and the number of the speed of light by the definition of the meter since 1983. Their numerical values express the definitions of units. What physics measures are the *dimensionless ratios* the units leave untouched: pure numbers, one quantity divided by another of the same kind, so that every unit cancels. The limit is the reading of the constant as a boundary of the domain of measurement. I make claims about the limit, as a boundary, and none about the number. The same holds for the speed of light: its numerical value expresses the definition of a unit, and the boundary is what is read.

Chance From Outside, Choice From Inside

Hard randomness also has a place in the structure of the whole account. Chapter 4 set out *the three modalities*, three aspects of everything there is: the *immanent*, process, what is happening here and now; the *omniscient*, structure, what holds the same for everyone, laid out as a pattern; and the *transcendent*, possibility, what makes things possible and new. The alignment this book holds fixed throughout places the microscopic domain under the transcendent: the potential, the creative, the generative, possibility and choice. It places the large-scale domain under the omniscient: the structural, the repeatable, the lawful, existence and pattern. Reality itself, the domain of interaction in which we live, is immanent, and lies between.

Hard randomness is what the transcendent edge looks like from inside the immanent: possibility entering the real without any structure before it that fixed it. Determinism is what the omniscient edge would look like if a signal could come from it: structure complete, with no possibility left.

Chapter 7 placed the triple of the real in the same three: causality is omniscient, change is immanent, and choice is transcendent. The two boundaries of the domain are the two places where the immanent runs out, into the omniscient on one side and into the transcendent on the other, and the omniscient and the transcendent are conjugate to each other.

Once the two boundaries are placed under the omniscient and the transcendent, the habit of thought from earlier in this chapter shows what it is: an attempt to account for all change by causation, which is to say by the omniscient alone. At one time in history all change was put down to local gods: the river rose because its spirit willed it, the crop failed because a god was angry. The modern world has attempted the opposite reduction, and holds, mostly without saying so, that there is no change that is not causation. Quantum mechanics has forced a rethinking, and the account of comparison offers the way to reconcile the two accounts of change, by causes and by choices. Change is the larger class. Some change is accounted for by knowledge of causes; other change is accounted for by choices made by agents; and choice is as much a way of accounting for change as causation is. We know that we cannot predict choices; if we could, they would be causes and not choices.

Picture a friend about to say the next word of a sentence. From outside, no one can predict the word: nothing said so far fixes it. From inside, the friend means it. Both descriptions are of the same word. The unpredictability in a choice is hard randomness seen from inside. Seeing hard randomness from inside, as choice, is the modal rotation of Chapter 7, one event described from two positions, now met in physics. From the position of law, the omniscient position, the outcome that nothing fixed is hard randomness, a gap in the specification, without pattern and without meaning. From the position of the one engaged, the transcendent position, the same outcome is choice, meaningful to the process from which the outcome came. Neither description can replace the other.

An outcome that is left undetermined is left undetermined entirely, and so randomness of this kind, where it appears, is hard in the strictest sense. What a hard random outcome turns out to be could not be pre-

dicted, and could not have been predicted, on any basis at all: theoretical, conceptual or practical. Suppose one knew absolutely the position and the motion of every particle in the universe, and, like God, also knew absolutely every law of nature, every argument and every reason. All of that together would still not be enough to fix the value of any hard random outcome. The omniscient cannot contain the transcendent. The transcendent is the creative source out of which the immanent, the here and now, comes, and it must have and be and become things that are inherently and fully unknowable, in completeness and in totality, whatever else holds. In plain words: what the new comes from can never be laid out in advance as structure, however much structure is known. No amount of structure, of argument or of reasoning, however perfect, can pass that limit or trespass on it. The limit is law: involute (turned back on itself), inherent and absolute; and those three are themselves a triple, in the sense this book has used throughout: three that are distinct, inseparable and not interchangeable.

The Residual the Theorem Guarantees

At bottom, hard randomness rests on one result. Symmetry and continuity cannot both be perfected in any comparison, so no lawful description of anything specifies every finding exactly: something is always left over that the law does not fix. Chapter 6 named that remainder *the residual*: what no lawful description fixes. The theorem guarantees it in every description. Hard randomness is the residual in its strong form: the residual taken as unpredictable for every comparer whatsoever, even one with unlimited access to the state of things, and not only for us. It is the irreducible residual that the Incommensuration Theorem guarantees. Bell's theorem, the uncertainty principle and Planck's constant, read together, are the physical face of that residual; atomic decay and the amplification of tiny differences are the residual entering the domain.

The same residual is the starting point of a series of results that this book walks through, each of them the residual read in a different place. Here they are in brief; each has its own chapter, and none is needed yet. Chapters 5 and 11: because the residual is always there, determinism is impossible, and superdeterminism with it. Chapter 13: read in the order of events, the residual gives time a structural arrow, present in every single comparison. Chapter 12: read in the reach of a signal, it shows that sending a signal faster than light names nothing coherent. Chapter 7: read in the structure of the real, it shows that the real is the triple of causality, change and choice. The rest of this chapter: read in the stance of the comparer, it shows that the subjective is situated in time, in space and in possibility, and that subjectivity, directed time and hard randomness are bound together as one triple. Chapter 7 and the section before this one: read in the meaning of what is found, the residual shows that choice and randomness are one event described from two positions, and that meaningfulness cannot be removed from that event. Chapter 21: read in the claim that the world is a computation, it shows that a simulated world containing real subjects is a restricted region of the one world, and not a separate world. Chapters 20 and 22: read in the relation of mathematics to physics, it shows that physics, logic and computation form one triple, and that mathematics is effective in physics because both are aspects of one comparison, and incomplete in it for the same reason.

What Only a Stance Can Tell Apart

Hard randomness is distinguished from soft randomness by being a fact about the world and not about our knowledge. One further point shows why the idea is more than a gap into which physics has been forced to admit noise. Take a purely *third-person* standpoint: a description from outside, of structure alone, with no one's standpoint in it. From there hard randomness and soft randomness cannot be told apart. Both produce the

same statistics. Both are equally unpredictable. No pattern tells them apart. The distinction is real, and it is registered from a stance, from a perspective, in a present.

Look at one single outcome. A purely third-person description cannot tell an outcome that nothing before it fixed from an outcome that was fixed and is unknown: both are unpredicted, and the structural description has no category for the difference. The difference is the difference between a finding that is new and a finding that was only hidden, and it is a difference in the finding, which is made by a comparer, in the act. So hard randomness is the residual as found, and a finding needs a finder: the residual is hard randomness, and not ignorance, only where a stance finds it.

That hard randomness needs a stance to find it does not take hard randomness out of the world. A stance is a comparer, and a comparer, as Chapter 1 said, is a position inside the world; its finding is an event in the world, made in an act of comparison, which is an interaction. So the difference between a finding that is new and a finding that was only hidden is a fact about the world, and it is a fact of the kind that only a stance inside the world can register. The step from hard randomness to the stance that finds it ties hard randomness to two other things, and the three of them belong together.

Three That Come Together

The three are subjectivity, directed time and hard randomness, and I call them *the bound triple*: three aspects of one act of comparison proceeding in a direction, each requiring the next. *Subjectivity* is the subjective pole of comparison, the stance from which every comparison is made. *Directed time* is the structural arrow of time, the one-way order in which every comparison runs from the act to its record; Chapter 13 develops it. Hard randomness is the residual in its strong form. Each implies the next, and the three close a circle.

Directed time requires hard randomness. Without hard randomness every state is fixed by every state before it, and the future is contained in the present as fully as the past is. Then the order of states is an indexing, a label of earlier and later on one fixed whole, like the page numbers of the finished novel from Chapter 5, which can be read in either direction and mark no happening. A direction needs the next moment to hold something the present does not, and that something is the residual.

Hard randomness requires subjectivity. The reason that hard randomness requires subjectivity is the step of the last section. For a single outcome, a purely third-person description cannot tell an outcome that no earlier state fixed from an outcome that was fixed and is unknown. The difference between new and only hidden is a difference in a finding, which a comparer makes from a stance.

Subjectivity requires directed time. A stance compares a present state with an earlier one. To perceive is to set what is found against what was found; to remember is to set the present against a record; to anticipate is to set the present against what could come. Each of these runs one way, from and to. With no direction of time there would be no difference between remembering and anticipating, and no side from which to look.

So each of the three requires the other two, and none is present without the others. By the three axioms of Chapter 4, they are distinct, inseparable and not interchangeable. They are like three faces of one solid: there is no having a face without the solid. Directed time, the process, is the immanent member of the three, and the most fundamental.

The physics of measurement confirms that the three are bound together. An *interpretation* of quantum theory is an account of what its mathematics says the world is like. The interpretation called Copenhagen keeps the observer, the random single outcome, and the irreversibility of measurement (a measurement once made cannot be unmade, as a stopwatch click is never unclicked): it keeps all three together. The interpretation that Hugh Everett proposed in 1957, in which every outcome is kept in a branching state, drops all three together. In Everett's interpretation there

is no fact of which branch a comparer is in, there is no outcome that is not fixed, and the basic equation is symmetric in time, running equally well forward and backward. Every interpretation keeps the three together or drops them together. The interpretations differ on which side they take; none separates the three. So the bound triple is established twice: derived from the Incommensuration Theorem, and found in physics, where no account of quantum phenomena has ever kept one of the three without the others. The physical face of one member, hard randomness, is in this chapter; the physical face of a second, directed time, is in Chapter 13.

What We Now Know

We now know that hard randomness is real, and that the theorem guarantees it. Hard randomness is an outcome that no description of what came before fixes, however complete, for any comparer whatsoever. Its plainest instance is the single click of a quantum measurement; atomic decay shows it directly; the three-body problem, the weather and everything that amplifies what it cannot resolve carry it up into ordinary scale. It enters through the microscopic boundary of the domain, a boundary that is categorical, defined by the absence of any signal, and about what lies beyond it nothing can be said. Pure randomness, the idealization at the far end of the line, stays excluded, as determinism does; hard randomness is what the real world holds between them.

Hard randomness matters for the book's question of what cannot be. The residual is what makes a direction of time possible, and a stance from which anything is found. Within the domain there is a world that is lawful and causal and open at its microscopic edge to possibility, and so open to a perspective and to a direction of time, neither of which can be had without that openness to possibility. A world with no residual would be the clockwork that Chapter 5 excluded: every state fixed by the states before it, and no direction and no stance in it.

Reading one event as hard randomness from outside and as choice from inside also settles what a famous theorem says of a particle. John Conway and Simon Kochen proved the *free will theorem*, which the next chapter states in full: if experimenters' choices are not fixed by everything before them, neither are the particles' responses. A particle's response that nothing fixed is, described from outside, hard randomness; described from the position of the process from which it came, the same event is choice. The meaning of a choice belongs to that process alone: it cannot be extracted, communicated in its specificity, or reproduced by any other process.

Summary

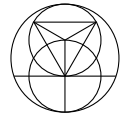
The central claims of this chapter are readings of known results: that Bell's theorem, the uncertainty principle and Planck's constant, read together, are the physical face of an incompressible residue that the Incommensuration Theorem requires, and that atomic decay and the amplification of tiny differences are instances of that residue entering the domain. That the microscopic limit is a categorical boundary, and that nothing can be assumed about what lies beyond it, are structural exclusions: they say what physics will not find, and why no theory can reach past the boundary.

Random has two meanings. Soft randomness is unpredictability from a limit of knowledge, which better knowledge could shrink; hard randomness is an outcome that nothing before it fixes, however completely described, for any comparer, even one with unlimited access to the state. Hard randomness is a fact about the world. It is the incompressible residue, what is left after every regularity is taken out, and it is the residual that the Incommensuration Theorem guarantees in every lawful description. Pure randomness and determinism are the two ends of one range, both idealizations, both excluded by the theorem; causation and choice describe the world between, with change, about half predictable, holding both, and hard randomness is the residual met there.

Causation is temporal, perceived, and holds to a degree; determinism is atemporal, mathematical, and holds without exception; so causation is not determinism, and a position that identifies the two is valid and not sound. Bell's theorem shows correlations that causation cannot cover; Heisenberg's principle shows a limit of access that no tool crosses; Planck's constant measures how much. The tests of 2015 removed every causal escape but one, superdeterminism, and that one asks so much of the world that Occam's razor takes it, before the structural argument against it is brought. Hidden variables cannot help, because a placeholder for a cause cannot restore causation. The single decaying atom, amplified in a Geiger counter, is hard randomness seen directly; three balls colliding at once, and the weather, carry it up to ordinary scale.

The microscopic limit is a categorical boundary of the domain of comparison, the twin of the horizon at the large-scale edge; what arrives through it can be nothing but random from inside, and about what lies beyond it nothing can be assumed. The number of Planck's constant is fixed by the definition of the kilogram; the limit is its reading as a boundary, and it sets no least length and no least duration. From outside, hard randomness; from inside, choice; one event, two descriptions, both correct, neither complete. From a purely third-person standpoint hard and soft randomness cannot be told apart; the difference is real, and it is registered from a stance, in a present. Because the difference between hard and soft randomness is registered from a stance, in a present, hard randomness belongs to the bound triple. In the bound triple, directed time needs hard randomness, hard randomness needs a stance, and a stance needs directed time. The three are distinct, inseparable and not interchangeable, and every interpretation of quantum theory keeps the three together or drops them together.

The free will theorem ties the particle's response to the experimenter's choice: is the person who sets the dial free, and what does that freedom carry with it for the particle?



Chapter 9

The Free Will Theorem

Two Laboratories and Two Dials

If I am free to choose what to measure, what does that say about the particle I measure?

In the setting where the question has an exact answer, two laboratories stand far apart. A source between them sends out pairs of particles, one particle of each pair to each laboratory. In each laboratory an experimenter turns a dial to choose the direction along which to measure, and the instrument records an answer. The question is whether the answers were already settled before the dials were turned.

Chapter 8 showed that a single atom decays when it does, and that nothing before the decay fixes the moment, however much is known. In this chapter I take the same question to the place where the question can be put most sharply: to the dials, and to the person who turns the dials.

The Enchanted Decks

To see what is strange about the answers, picture two ordinary decks of playing cards, enchanted together while they lie side by side on a table. One person keeps one deck; a second person takes the other deck and flies to Paris. Once the two have parted, each shuffles a deck thoroughly, takes

the top ten cards one at a time, and writes down each card. Each then emails the list to the other, so that neither can change a list after seeing the other's. The two lists are identical.

An ordinary agreement works differently. Seal a red card in one envelope and a blue card in another, shuffle the envelopes, and post one to each of two cities. Whoever opens red knows that the other envelope holds blue, and nothing traveled between the cities, because the agreement was laid out when the envelopes were sealed. The enchanted decks are not of the envelope kind: each person shuffles after the two have parted, so no order laid out in advance can make the lists match.

The enchanted decks have a second strange feature. Suppose one person, hoping to send a message, arranges a deck in a chosen order instead of shuffling it. Then the two lists no longer match. The decks agree only when both have been shuffled, so neither person can use them to send a message.

Pairs of particles can be prepared together so that they behave much like these decks. Such pairs are called *entangled*: two particles whose outcomes, measured far apart, agree in a pattern that no separate instructions could give them. Each pair is like one card drawn from each deck: one particle comes to one laboratory, and its twin goes to the other. John Bell showed in 1964 that the agreement between entangled particles is not of the envelope kind: no account in which the answers are laid out at the source reproduces it. The particles carry no answer fixed in advance. An attempt to fix the answer, by measuring a particle before it reaches the instrument, destroys the entanglement on which the agreement depends, as arranging a deck destroys the agreement of the cards.

The picture stops at the dials, because cards have no settings. In the experiments, a random number generator in each laboratory picks the setting at the last moment. So nothing at the source, and no signal traveling at the speed of light (the fastest that any influence can pass from one place to another), can know in advance which direction each instrument will measure. When the two dials are set to the same direction, the two

answers always match. When one dial is set partway between the direction of the other dial and the direction at a right angle to it, whether the answers match varies by chance from pair to pair. That chance agreement at the in-between settings is what the experiments on Bell's result measure, and Chapter 10 walks through them.

The Theorem

Does physics itself show that the world is not fixed in advance? It does, through a theorem proved by John Conway and Simon Kochen: if an experimenter is free to choose which measurement to make, the particle measured is free in the same sense. I will walk through the theorem with its starting points, show why a failure of determinism anywhere is its failure everywhere, and set out the limits that every finite observer meets. The walk arrives at one conclusion: *indeterminism*, the plain fact that what happens is not completely fixed in advance, is real at the scale of particles, and therefore at every scale.

Indeterminism is the denial of determinism. Chapter 5 defined *determinism* as the claim that everything is fully and exactly fixed: every value of every content for every context, with no observer and no time needed.

The *free will theorem*, proved by John Conway and Simon Kochen in 2006 and strengthened by them in 2009, answers the question of the two laboratories, whether the answers were settled before the dials were turned. Its answer is a conditional result of great reach: if the choices of experimenters are not fixed by anything that came before them (if, that is, the experimenter is free), then the responses of particles are not fixed by anything that came before them either. The particle, in this exact sense, is also free.

"Free" here has one meaning: not a function of anything that existed before the event. A *function*, in the mathematician's word, is a rule that gives exactly one output for each input; to be a function of the past is to be settled by the past, as an output is settled by its input. A free response is one that no input from the past settles.

The "if" belongs to the theorem. The theorem is a link between two freedoms, the freedom of the one who chooses the setting and the freedom of the thing that answers, and the link is what Conway and Kochen proved. The experimenter's freedom is a separate matter, and I secure it on two separate grounds, each enough by itself.

The theorem is a mathematical result. Conway and Kochen's papers of 2006 and 2009 give the proof, and the next two sections walk through its starting points and the shape of the step from them to the conclusion.

Four Starting Points

Every theorem starts from assumptions, the statements it takes as given and from which the proof proceeds. Conway and Kochen gave theirs short names, SPIN, TWIN and MIN; the Appendix gives the history of the names. In plain words, their content is four assumptions.

The first assumption is that quantum mechanics applies (SPIN and TWIN). Quantum mechanics is the theory of matter and light; its findings come in countable steps. Some particles have *spin*, a kind of rotation that belongs to the particle itself, and a *spin-one* particle is one whose spin can be measured along any direction. Measure the square of its spin (the spin multiplied by itself) along three directions at right angles to one another, like the three edges that meet at the corner of a box. Quantum mechanics says that the three answers are always one, zero and one, in some order (SPIN). Put another way: asked along any three directions at right angles, the particle answers "zero" along exactly one of them. Two such particles can be prepared together as entangled twins, so that, measured along the

same direction, they always give matching answers (TWIN). This first assumption is a constraint on the scale of particles: particles behave as quantum mechanics predicts.

The second assumption is that relativity applies: the speed of light is an upper limit on how fast any influence passes from one place to another. At root, the speed of light is the outer edge of the *domain of comparison*, the range in which comparisons can be made. The speed of light marks the edge of what any comparison can connect, and the speed of light is not the speed of any one thing. Light, which travels in empty space at exactly that speed, is only the perfected example of that edge. Because the speed of light marks the edge of what any comparison can connect, the limit that the speed of light sets holds for every influence, however the influence is carried, and not for light alone. Two events are *spacelike separated* when they are so far apart in space, and so close in time, that no signal traveling at the speed of light could pass between them. When the two measurements are spacelike separated, a choice of setting in one laboratory cannot affect the outcome in the other. The second assumption, applied to measurements that are spacelike separated, is the *locality* half of MIN: an influence reaches only as far as the speed of light allows.

The third assumption is that there is no backward causation: later events do not fix earlier settings. The setting an experimenter chooses is not arranged in advance by what the particle will later do.

The fourth is free choice: the experimenter's choice of setting is not a function of anything the particle could have had available to it beforehand. Free choice is the freedom half of MIN. Free choice is also the assumption from which the conditional form of the free will theorem comes: if the experimenter is free, the particle must be also. MIN as a whole joins the second and the fourth: when the two measurements are spacelike separated, each experimenter chooses a setting freely, and neither particle's response depends on the setting chosen in the other laboratory.

Each of the four rests on established physics and practice, and together they are among the most securely established starting points in the history of physics. Every science of causes presupposes that there is no backward causation. Special relativity, Einstein's theory of space and time in which the speed of light is the upper limit, has been confirmed to extraordinary precision. The matching answers of entangled pairs have been confirmed in every Bell-type experiment carried out to date; a *Bell-type experiment* is one that tests the agreements that John Bell showed no local theory with answers fixed in advance can reproduce. Free choice of settings is what every experiment presupposes when it treats its settings as independent of what it measures. The free will theorem shows that these four, taken together, lead to a conclusion that many people have resisted on philosophical grounds.

What the Theorem Concludes

How do the four starting points lead to the conclusion? The proof itself is mathematics, and the shape of its step can be told in words. Suppose each particle's answers were fixed by what came before it. Then the particle would carry, for every direction it could be asked along, a ready answer, zero or one. The first starting point says that along any three directions at right angles, exactly one of the three answers is zero. So the ready answers would have to form a list in which every set of three directions at right angles holds exactly one zero. In 1967 Simon Kochen and Ernst Specker proved that no such list can exist: there is a set of directions, finite in number, for which the rule of exactly one zero cannot be kept on every set of three at once.

One way out remains to be closed: a particle's answer along one direction might depend on which other two directions it is asked along together with it. The twin closes that way out. In the far laboratory the twin is asked along a single direction, chosen freely there; no influence can pass between the laboratories in time, and later events do not reach back to ar-

range earlier ones; and the twin, asked along the same direction, always gives the matching answer. The twin's answer cannot know which other two directions were chosen in the first laboratory, so the first particle's answer along that direction is one and the same answer, whichever two other directions it is asked along with. Then the list that Kochen and Specker ruled out would have to exist, and it cannot. So the answers are not fixed by what came before. The Appendix gives the technical form of the step.

From these assumptions the theorem derives that the responses of particles to measurements are not functions of anything earlier. The responses are not functions of the particle's history, nor of the state of the universe at earlier times. They are not functions of any *hidden variable* (a quantity the particle would carry, unseen, that would fix its answer in advance, as Chapter 8 described), nor of any combination of hidden variables. When the experimenter is free to choose which measurement to perform, the particle's response is not fixed by anything that existed before the measurement.

The indeterminism is *ontic*, a feature of the world, and not *epistemic*, a feature of our knowledge or our description of the world. The answer is missing from the world before the measurement, and not only from our notebooks. Given the experimenter's freedom, no deterministic account, no complete record of earlier states, however detailed, can produce the outcomes that are observed.

The experimenter is free, and that freedom is secured twice.

The experimenter's freedom is secured, first, by the Incommensuration Theorem, the theorem of Chapter 3: symmetry and continuity cannot both be applied, at once and absolutely, to one comparison. The Incommensuration Theorem is a structural fact about the nature of knowing. Chapter 5 showed that the Incommensuration Theorem excludes determinism. Determinism needs symmetry perfected (one law, the same everywhere) and continuity perfected (every state specified, with no gap) together, at once; that is the pair the theorem forbids, so total determinism is impossible, and no complete specification, written in advance, fixes the ex-

perimenter's choice along with everything else. That the experimenter's own choice, in particular, is not fixed follows with one more step, which I take in a later section of this chapter: a failure of determinism anywhere is a failure of determinism at every scale, the experimenter's scale included.

The experimenter's freedom is secured, second, by the practice of physics itself. Every experiment presupposes that its settings are independent of what it measures. That presupposition makes a reading evidence about the thing measured, rather than one more product of whatever produced the setting. A measurement is a measurement, and not a replay, only when the choice of setting is not fixed in advance.

The second ground uses nothing from the first. So the free will theorem confirms from physics that determinism is impossible, and it confirms this independently of the Incommensuration Theorem. With the experimenter's freedom secured, the condition of the theorem is met, and its conclusion follows: the particle's response is not fixed by anything that came before it.

The Chooser and the Chosen

Earlier chapters met three scales of description: the *microscopic*, the scale of particles; the *mesoscopic*, the middle scale, our own everyday size; and the *macroscopic*, the scale of the cosmos. To see what the theorem means across these scales, two terms are needed. A *person-mode*, first person or third person, names a relation between an observer and a scale; it is more than a convention of grammar.

To operate in the first person is to act at the scale of one's own size: for a human observer, the mesoscopic. First-person engagement is the mode of perceiving, acting and choosing, the mode in which the observer is inside the domain it acts in and cannot step away from its own position.

To operate in the *third person* is to observe a domain from a scale other than the domain's own: a human observer examining a particle from the everyday scale, or examining the cosmos from the everyday scale. Third-

person observation is the mode of describing, measuring and building theoretical models, the mode in which the observer stands at a distance from what is described and describes it from outside.

The two laboratories show both modes at once. The experimenter's hand on the dial, turning it at arm's length in a room of her own size, is the first person. The record of the particle's answer, read off a screen about something far smaller than any hand, is the third person.

The free will theorem establishes its result where these two modes meet. The experimenter operates in the first person, at the everyday scale: turning the dial is a first-person act at our own size. The particle's response is observed in the third person: an observer of everyday size examines something at the scale of particles. The free will theorem shows that the first-person freedom of the experimenter and the third-person openness of the outcome are tied together by structure. Each one brings the other with it.

Two Roads, and Only Two

The theorem creates a fork with two roads. On the first road the experimenter is free: the choice of measurement is not fixed by anything earlier. On the second road the experimenter is not free.

On the first road the theorem applies directly: the particle's response is not fixed by anything earlier, and indeterminism is real at the scale of particles.

On the second road the experimenter's choice, an everyday event such as the movement of a dial or the click of a button, is fixed by the earlier state of the universe, including the state of the very particles being measured. That road is superdeterminism, which Chapter 5 met: the settings of the detectors and the states of the particles share one deterministic history, and the choices of measurement are tied to the hidden variables of the system measured. It breaks *statistical independence*, the assumption that the settings chosen are unrelated to the states measured, and it implies that

every event, including every measurement and every design of an experiment, is fixed by the starting conditions of the universe. Chapter 11 sets out superdeterminism as its holders state it, and shows why it is excluded. First, the two roads are the only two, with no halfway house between them.

One Exception Breaks a Guarantee

A guarantee that covers everything is broken by a single exception. A shop that promises every item is perfect has broken its promise with the first cracked cup. Determinism is that kind of guarantee: by its definition, it is the complete specification of every content for every context, at every scale. There is no partial determinism, any more than there is a partial guarantee of everything.

The *catastrophic theorem of determinism* follows from that definition. If any amount of indeterminism is evident or available anywhere, to any degree, then no scale is completely deterministic. The reason is that a scale cannot be completely specified while the scales that make it up, or that it helps make up, are not. The step runs both ways. Upward: particles make up the everyday world, so if the particles are not completely specified, the everyday world they make up is not either; the next paragraphs show this step at work in physics. Downward: a free experimenter at the everyday scale brings with it a free particle at the scale of particles, which is what the free will theorem says; so a failure at the larger scale reaches the smaller scale that helps make it up. So one failure anywhere is a failure at every scale. The argument is one case of a general form, the catastrophic form: either every member of a kind has the property, or no member of the kind has it. Here: either every scale is completely deterministic, or none is. There is no partial determinism, and so there is no partial superdeterminism.

Physics shows this step at work: a failure at one scale becomes a failure at the larger scales that it helps make up. Among the processes that join the scale of particles to the scale of the experimenters are processes that magnify small differences into large ones. A difference in the air too small for any instrument to record grows, over days, into the difference between a storm and a calm. Edward Lorenz found the same magnifying of small differences in 1963, in a simple model of convection, the rolling of air heated from below: "slightly differing initial states can evolve into considerably different states". In the light of that finding, Lorenz examined whether the weather can be predicted very far ahead. A double pendulum, one pendulum hanging from the end of another, shows the same magnifying of small differences in the hand: release it twice from positions too close to tell apart, and within seconds the two swings bear no resemblance to each other. I call such magnifying *microstate amplification*: "the process by which small differences at a fine-grained level of a system are propagated and magnified into differences at the level that matters for the system's behavior".

Now put microstate amplification beside the second road, superdeterminism. A superdeterminist needs the settings the experimenters choose to stay tied to the states of the particles they measure, through the past that the settings and the particles share. An outcome at the scale of particles that nothing earlier fixes, carried up toward the scale of the experimenters by a process that magnifies small differences, is magnified like any other small difference, and in time it can reach the settings themselves, which the shared past then does not fix. The definition already makes the failure immediate at every scale, since a scale made of parts that are not completely specified is not completely specified either. Amplification shows the same failure in the world, at the scale where the experimenters choose their settings. So a determinism that held only approximately at the scale of particles is no determinism at the scale of the experimenters either.

Two independent routes supply the failure that the catastrophic theorem spreads to every scale. The first route is physical: the free will theorem, given the experimenter's freedom, shows that the responses of particles are not fixed by anything earlier, wherever measurement occurs. The second route is from the nature of comparison: the Incommensuration Theorem shows that determinism requires symmetry and continuity perfected together, which no comparison allows. Each route alone is enough. Two independent routes to one result are stronger than either alone, as a cable is stronger than its strongest strand.

Picture the whole argument as a drawing in three parts. On the left runs the physical route, the free will theorem, ending in a failure of determinism somewhere. On the right runs the route from comparison, the Incommensuration Theorem, ending in a failure of determinism somewhere. In the center stands the definition of determinism, which takes a failure anywhere and makes it a failure at every scale.

The definition works in the other direction as well, on a determinism that holds and not one that fails. A determinism that held absolutely, with no indeterminism at any scale to any degree, would be superdeterminism, and every event at every scale would be fixed by the starting conditions of the universe. There is no partial position. As Chapter 5 put it, determinism is impossible by the Incommensuration Theorem, and it cannot even be moderately held, because any crack in the specification is the end of the specification.

What Every Account of Quantum Theory Must Give Up

An *interpretation* of quantum mechanics is an account of what is happening in the world when the theory's predictions come true. The free will theorem constrains every interpretation: any interpretation in which outcomes are fixed in advance must give up one of the theorem's assumptions.

The best-known interpretations show the three ways of giving one up. Bohmian mechanics, the account David Bohm gave in 1952, in which particles have definite positions guided by a wave, gives up the locality half of MIN: the outcome in one laboratory depends on the setting in the other. The many-worlds interpretation, which descends from the account Hugh Everett gave in 1957, in which every outcome occurs in some branch of the world, gives up the single outcome the theorem speaks of. The single outcome is not a fifth starting point: the theorem takes it for granted in what its starting points speak of, each measurement giving one answer, the particle's response. An account in which every answer occurs gives up that one answer, and with it the thing the four starting points are about. Superdeterminism gives up free choice.

There is a deeper pattern among the interpretations, and Chapter 8 set it out: every interpretation keeps three things together or drops them together, the observer, the single outcome that comes by chance, and the fact that a measurement cannot be undone. Chapter 8 also showed why the three come together: they are the bound triple as quantum theory meets it, a stance (the observer), hard randomness (the single outcome by chance) and a direction of time (the measurement that cannot be undone), each of which requires the next.

Three Limits on Knowing

The free will theorem, the three scales and the two person-modes, taken together, show three limits on what a finite observer can know about a system at a scale other than its own. These limits come from how an observer stands toward the scale it observes, not from the instruments it holds, and no better instrument removes them.

The first limit belongs to the very large scale, observed in the third person: the absolute elsewhere. At every event there is a *light cone*: the set of every event that can be connected to it by cause, a forward cone of what it can affect and a backward cone of what can affect it. The *absolute elsewhere*

is every event outside both cones. An observer of everyday size, examining the cosmos, cannot reach events in its absolute elsewhere, nor receive anything back from beyond a *horizon*, an edge beyond which nothing can reach the observer. The first limit, the absolute elsewhere, is the limit in space.

The second limit belongs to the scale of particles, observed in the third person: the uncertainty relation that Werner Heisenberg found in 1927. Recall from Chapter 3 that two things are *conjugate* when each can be perfected only at the cost of the other, and neither comes first. An observer of everyday size, examining something at the scale of particles, cannot specify both members of a conjugate pair of quantities, such as position and momentum (where a thing is, and how much motion it carries), more sharply than the relation allows. The second limit, the uncertainty relation, is the limit on information.

The third limit belongs to the everyday scale, observed in the first person: no backward causation. An observer acting in the first person cannot reach information from the future of an event. The direction of time is what makes causality what it is, and the same direction sets causality's limit. No backward causation is the limit in time. The limit in time is the arrow present in every single comparison, the one-way order in which a record follows the act that made the record, as I set out in full in Chapter 13.

The limit in space and the limit on information also stand in the way of a finite observer's knowing its own future completely. To predict its own future completely, a finite observer would need two things: the present state of things far away whose influence will in time arrive here, and the present state of things very small whose differences will in time grow large. The state of things far away lies in the absolute elsewhere, outside the observer's light cone. The state of things very small lies below the limit that the uncertainty relation sets. So a finite observer cannot know its own future completely, and all three limits stand in the way: the direction of time directly, and the other two through what a complete prediction would need.

That such limits exist, and that no future theory will remove them, follows from the structure of comparison. Their form and their constants belong to physics: the speed of light, and Planck's constant, the number that sets the size of the smallest quantum step. The form of the uncertainty relation can be traced one step further. A standard tool of mathematics, the *Fourier transform*, turns a description of a thing by position into a description by wavelength, and the sharper the one description is, the broader the other. The same trade can be heard in sound: a handclap, sharp in time, is a broad spread of pitches, while a flute note held for a long time has one clear pitch. The uncertainty relation has that form. That form is the relation the first of the three axioms of *An Immanent Metaphysics* states between structure and possibility (the omniscient and the transcendent of Chapter 4): the two are conjugate. The Incommensuration Theorem derives that axiom.

The three limits are permanent, and the names the metaphysics gives them say so. That permanence matters for this chapter's question: the openness the free will theorem finds is no gap in present knowledge that a later finding will fill. *An Immanent Metaphysics* names three irreducible mysteries that answer to these three limits, mysteries that no further finding dissolves, since each sits at a limit that no future theory removes. They are the mystery of the horizon, at the very large scale; the mystery of the infinitesimal, the endlessly small, at the scale of particles; and, the most fundamental of the three, the mystery of the subconscious, at our own scale. The first two have faces in physics: the light cone and the uncertainty relation. The third has no such face. It "is the limit at the act itself: a comparison cannot take its own act as its content", and it "carries no constant of physics". The word subconscious has a special sense here, narrower than its everyday sense. It names the limit at the self's own act of comparing, which in my book The Basis of Physics is the horizon at the scale of the self: "the surface of unconsciousness, past which the self does not see". Within any domain, the three mark the edges of what *An Immanent Metaphysics* calls the Great Mystery, "the convergence of the unknown and the unknowable", the place where what is not yet known meets what

can never be known. No finite observer stands in the position from which a total description could be given, and no total description can include the finite observer's own view.

Every measurement meets these three limits. Chapter 1 showed what a measurement is: a comparison made from inside the world, which brings a finding into being and lays it down as a record. The limits are the edges of that act, wherever it is made.

Questions a Reader Asks

Does the free will theorem prove that we have free will?

The free will theorem is conditional in its own content: if the experimenter is free, the particle is free. I secure the experimenter's freedom on other grounds: the Incommensuration Theorem, and the presupposition of every experiment. What the theorem adds is that freedom reaches past the experimenter. Wherever freedom exists, it extends down to the particles measured; freedom is not confined to experimenters.

Is it not strange to call a particle "free"?

In the theorem the word has one meaning: not a function of anything earlier. Whether that openness, seen from within, is choice is a further question, and Chapter 7 answered it: hard randomness and choice are one event described from two positions, randomness from outside and choice from inside. Watching people leads to the same conclusion: what is random from outside is choice from inside. In Chapter 7 I told the picture of the program that shrinks computer files. In the picture, however good the program is, some part of a person's behavior will not shrink, and that leftover part cannot be told from randomness. That leftover part is the part the person counts as their choice.

Is the free will theorem only Bell's theorem again?

The two are close relatives. Bell's theorem shows that no local theory with answers fixed in advance reproduces the agreements of entangled pairs. The free will theorem draws the conclusion about the answers themselves: given free settings, the answers are not fixed in advance at all.

Could a determinist accept every observation and still be right?

Only as a superdeterminist, and then no observation would be evidence of anything: a reading would be one more product of whatever fixed the setting. The Incommensuration Theorem excludes superdeterminism in any case, with all the more reason than it excludes determinism, since superdeterminism asks for everything determinism asks for and more. Chapter 11 walks through both points in full, and shows that superdeterminism absorbs every possible piece of evidence in the same way as idealism (the view that all reality lies in the mind of the perceiver) and as the claim that the world was made with only the look of age.

No observation could ever tell superdeterminism, taken as a claim about all events, apart from its rival, since in its world every observation was itself fixed from the start. A reader may then ask: can any observation tell this book's own starting points apart from their rivals?

The test of observation applies to claims about observations, and superdeterminism is a claim of that kind: it is offered as a claim about the physical world. The axioms of a metaphysics are of another kind: they are "no law of any one domain". The Incommensuration Theorem is proved from the meaning of comparison. Claims of those kinds are judged by their own methods: by proof, and by the structure of the act of comparison. To bring the test of observation against them is to apply a test to claims of a kind it does not reach.

Why It Matters

The experimenter's freedom, which every experiment presupposes, carries the particle's freedom with it. So the old picture of a world that runs like clockwork is excluded twice over, from two directions independent of

each other: once from the structure of comparison, before any experiment, and once from physics, by a theorem whose starting points are among the best established in the science. And because one failure of determinism anywhere is a failure at every scale, the exclusion does not stay with particles. No scale, from particles to people to galaxies, is completely fixed in advance. What cannot be, met here in the laboratory, is a world whose every outcome was written before that outcome happened.

Summary

The claims of this chapter are of three kinds. The free will theorem is a result of mathematical physics, proved by Conway and Kochen, and I take the free will theorem as they proved it, conditional in its own content; read beside the structure of comparison, it is a confirmation from physics, independent of the Incommensuration Theorem, that determinism is impossible. That no scale is completely deterministic is the structural exclusion of Chapter 5, and the catastrophic theorem, that one failure anywhere is a failure at every scale, follows from determinism's own definition. That the three limits on knowing exist, and that no future theory will remove them, is a claim of the structure of comparison itself; their form and their constants belong to physics.

The free will theorem of John Conway and Simon Kochen (2006, strengthened in 2009) establishes that if experimenters are free, particles are free: the responses of particles are not functions of anything earlier, not of their history, not of the earlier state of the universe, not of any hidden variable. The theorem is conditional in its own content, and its four starting points (quantum mechanics applies; relativity applies; no backward causation; free choice) are among the most securely established in physics. The indeterminism it finds is a feature of the world, not of our knowledge.

The experimenter's freedom is secured twice: by the Incommensuration Theorem, which makes total determinism impossible, and, independently, by the presupposition of every experiment that its settings are independent of what it measures. The second ground uses nothing from the first.

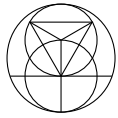
The catastrophic theorem of determinism follows from determinism's definition: one failure anywhere is a failure at every scale, microstate amplification shows that failure at work in the world, and two independent routes supply the failure. Superdeterminism, the only alternative, would leave no observation able to count as evidence of anything, and the Incommensuration Theorem excludes it with all the more reason than it excludes determinism; it absorbs every possible piece of evidence as idealism does, and as the claim that the world was made with only the look of age, which Chapter 11 shows in full. The theorem works where the experimenter's first-person act at our own scale meets the third-person observation of a particle, and it ties the freedom of the one to the openness of the other. Every interpretation of quantum mechanics in which outcomes are fixed must give up one of the theorem's assumptions.

The limits of knowing (the absolute elsewhere, the uncertainty relation, and no backward causation) are built into where a finite observer stands toward the scales it can observe, and no better instrument removes them. They answer to the three mysteries of the horizon, the infinitesimal and the subconscious. Indeterminism is real, at every scale.

Then what of the famous experiments on entangled pairs that John Bell's work made possible: what did they show, and why did their result shock everyone?

PART THREE

BELL AND THE ESCAPE THAT IS NOT THERE



Chapter 10

What Bell Found

Two Detectors, Far Apart

What did Bell's experiments actually show, and why did the result shock everyone? In each of Bell's experiments, a source in the middle of a laboratory makes two particles together, in one event, and sends them off in opposite directions. At each end of the laboratory, far apart, a detector waits.

Each particle has a direction that can be measured. For a particle of matter, such as an electron, that direction is called its *spin*. A useful picture is a spinning top, whose axis points some way in space, like an arrow set at an angle. The picture is an aid and no more: a particle has no inside that turns. For following the experiment, the arrow is enough.

A detector reads one thing about the arrow, and only one. The experimenter turns the detector to one direction of its own, and the detector then asks the arriving particle one question: does your arrow point along my direction, yes or no? Each detector is asked a single question, set by how it is turned, and each gives back a single answer, yes or no.

Now the pattern. When the two detectors happen to ask the same question, the two answers always agree, or always disagree, depending on how the pair was prepared at the source. When the two detectors ask questions at right angles to each other, whether the two answers agree is random, like the toss of a coin. The record at each detector, taken alone, is a random list of yes and no, like a run of coin tosses. Only when the two

records are brought together afterward, by ordinary means (a cable, a telephone call, a sheet of paper carried across the room) and at ordinary speed, does the pattern between them appear.

A pattern of agreement between two records, found when the two records are set side by side, is called a *correlation*. Two particles whose records, taken far apart, agree in a pattern that no separate instructions carried by each particle could give them are called *entangled*. Chapter 9 told a story of this kind: two enchanted decks of cards, each shuffled after the two owners parted, whose logs still agree card for card; and beside them two sealed envelopes, one red card and one blue, whose agreement was laid out when they were sealed. Entangled particles agree in the way the decks agree.

In 1964 John Bell proved that the agreement quantum theory predicts for such pairs is stronger than any account can supply in which each particle carries its answers with it from the source, and nothing done to one particle instantly affects the other. Experiments since then have found that nature gives the stronger agreement. The shock came from what had to go: one of two perfections that physicists wanted in their theories, the same law everywhere, or every detail fixed. In this chapter I follow Bell's result step by step, and I show that Bell's result is one case of the limit that Chapter 3 found in every comparison.

Bell's Question

Bell's theorem concerns the correlations between measurements made on pairs of particles that were prepared together and then separated. It asks a precise question. Could the correlations that quantum theory predicts be accounted for by a theory with two features: each particle carries, from the moment the pair is made, definite properties that fix what any measurement on it will find; and the measurement made on one particle has no in-

stant effect on the other? (Quantum theory, which physicists also call quantum mechanics, is the theory of matter and light, whose findings come in countable steps.)

A theory of that kind is called a *local hidden-variables theory*, and it makes two assumptions. The first assumption is *realism*: each particle has definite properties before the particle is measured, fixed by quantities that we may not know but that exist all the same. Those quantities, present in the particle though hidden from us, are the *hidden variables*. The second assumption is *locality*: what is done to one particle does not instantly alter what will be found at the other particle.

Bell showed, by a mathematical proof from the two assumptions alone, that any theory making both assumptions must obey a certain inequality in the tallies of the correlations: a ceiling on how strongly the two records can agree. Quantum theory predicts that the agreement goes over the ceiling.

The size of the excess can be said in numbers. In the form of the test that John Clauser, Michael Horne, Abner Shimony and Richard Holt gave in 1969, the detectors are set at in-between angles, four pairs of settings are used, and the four correlations are combined into one score. Any account that makes both of Bell's assumptions scores two at most; that ceiling of two is set by the structure of the two assumptions themselves. Quantum theory predicts that the score can reach twice the square root of two, a little more than two point eight. The higher value of twice the square root of two is set by quantum theory. The scores the experiments measure are above two.

Experiment, beginning in earnest with the work of Alain Aspect and his colleagues in 1982, and continuing through the tests of 2015 that closed the main gaps, confirms that the ceiling is broken. Physicists call such a gap in a test a *loophole*: a way some other explanation might slip through. In 2015 three groups of experimenters closed the loopholes of detection and of locality in the same year. In one of those tests, by Bas Hensen and colleagues at Delft, the two entangled electrons were one

point three kilometers apart, a little over three quarters of a mile, a distance the experimenters chose and measured. They set the electrons that far apart so that no signal at the speed of light could pass from one measurement to the other while the two measurements were being made; that is how a test closes the loophole of locality. At least one of Bell's two assumptions must therefore be given up. The standard summary is compact: reality cannot be both local and realistic. In this book, *Bell's theorem* means exactly this: no account in which every outcome is fixed by hidden local causes can match all the predictions of quantum theory, and the measurements confirm those predictions.

The theorem of physics is Bell's, and I take Bell's proof, and Bell's route to the inequality, as they stand. Bell's argument is rigorous within its own *formalism* (the mathematics in which a theory is written), and that formalism is quantum mechanics. What I offer is a reading of the result: an account of why the two assumptions cannot both hold, which does not depend on the formalism, and which applies to comparison as such. The claim that Bell's theorem is an instance of something more general is mine, and the rest of this chapter is the argument for that claim. The argument needs nothing beyond what the first chapters of this book have shown (that every measurement is a comparison, and the theorem about comparison found in Chapter 3) and what this chapter states.

Lawfulness Is a Kind of Symmetry

Consider first what realism assumes, and what science in general assumes when it takes the world to be lawful. Chapter 2 moved an experiment across the room and ran it on Tuesday instead of Monday, and the result came out the same. Science assumes that this is how things are: given similar conditions, the results of a procedure performed in one place and time will be the same as the results of the same procedure performed in other places and times. That assumption is what makes an experiment repeatable, and being repeatable is what makes an observation scientific.

Look at the form of the assumption. There is a sameness of content (the results are the same) and a difference of context (the places and times differ). Chapter 2 named that form. *Symmetry* is sameness of content across difference of context: change the setting, and what is found stays the same. To take an experiment as repeatable is to assume symmetry in advance.

The assumption goes deeper than method. Science regards the basic workings of physical process as the same everywhere in the universe, and it treats the laws of physics as unchanged when one shifts in time or in position. The mathematics of the laws of nature rests, at every level, on symmetry. The *conservation* of energy and of momentum (each an amount that stays the same through every change) is a law of symmetry, as the theorem the mathematician Emmy Noether proved in 1918 shows. The theories of relativity are built on relationships that stay the same when position and motion change. The largest claim of symmetry within science is that the same physical laws, as content, apply everywhere in the universe, as context. For science to regard the universe as lawful is to assert a basic symmetry. *Lawfulness*, the same physical law in every place and time, is symmetry applied to the physical world.

Now return to Bell's first assumption. To say that a particle has definite properties before it is measured is to say that what the measurement finds, its outcome, is fixed by the state of the particle alone, by a definite pattern, a law, whatever the context in which the measurement is made: whatever the setting of the apparatus, whoever performs the measurement, and wherever and whenever it is performed. An outcome fixed by the state of the particle alone, whatever the context, is sameness of content across difference of context. Realism, in Bell's sense, is lawfulness at the scale of the single particle. Realism is symmetry.

Locality Is a Kind of Continuity

Consider next what locality assumes. As physics uses the word, locality asserts that no physical influence, no interaction and no signal can travel faster than light. It asserts that no physical process crosses space instantly: nothing jumps between two far-apart points in no time. If an interaction spans any distance, it also spans some time greater than zero. A very small change in where an interaction reaches brings at most a very small change in when it reaches there. There are no jumps.

Look at the form of this assumption as we looked at the form of realism. Time and space are measures of context: they are the setting, the where and the when. Interactions, and the stuff that interacts, are content within that setting. So the claim that no signal can instantly cross any distance is the claim that for every small change in the content (a process, an interaction, a signal), there is a matching small change in the context, in time and in space. Chapter 2 named that form too. *Continuity* is sameness of content within sameness of context: hold the setting fixed and nothing jumps. Locality, in asserting that there are no abrupt, instant changes, is a special case of continuity. Locality is continuity applied to the physical world.

A further consequence matters for what follows. To assert locality absolutely, without exception at any scale, is to require complete continuity of interaction at every scale, down to and including the scale of zero distance and zero duration. It is to require that every physical process be fully specified, step by unbroken step, by smooth local rules. Absolute locality is therefore a large demand. It is the demand that the world be, at every scale, a process specified without a gap, and that is the same demand that Chapter 5 found can never be met: Chapter 5 showed that determinism would have to be made good all the way down, to the scale of exactly zero, and that it cannot be.

Physicists now state Bell's assumptions as three: local causality, realism, and measurement independence. *Local causality* says that every influence on what is found at a detector comes from nearby, and none from

what is done far away. *Measurement independence* says that the experimenters' choice of settings is free of whatever fixed the particles. Read with care, the three sort into the two demands already found: lawfulness and locality. Realism is lawfulness, which is symmetry. Local causality and measurement independence both belong with locality, which is continuity: each asks that what is found at one place not depend on what is set at another place, far away. The reading holds for the newer statement in three parts as it holds for Bell's first statement in two.

One more thing about the word locality. I use "locality" in two senses, and the difference will matter before the chapter ends. The broad sense is Bell's: every agreement between what is found in two places is carried by causes that travel from place to neighboring place, so that nothing done at one detector has any part in what is found at the other. The narrow sense is the one physics keeps, which this book calls locality "as physics keeps it, in no signaling": no signal crosses a distance in no time. A *signal*, here, is anything that passes from one place to another and lets a record at the second place depend on an act at the first.

The Same Shape and the Same Root

Chapter 3 found the one limit that every later chapter uses. In plain words, the Incommensuration Theorem says that symmetry and continuity cannot both be applied, at once and absolutely, to one comparison, and that neither can asymmetry and discontinuity. It is a structural fact about the nature of knowing. (*Asymmetry* is difference of content across difference of context; *discontinuity* is difference of content within sameness of context: a jump, a click, a step.) Of the four compounds of comparison, then, two pairings are forbidden when both members are applied absolutely and together: symmetry with continuity, and asymmetry with discontinuity. Two pairings survive, and Chapter 3 called them the two sur-

vivors: continuity with asymmetry (smooth and going somewhere, like a wave), and symmetry with discontinuity (the same everywhere and countable, like a particle).

The word "absolutely" matters, and it is never dropped. To apply a compound absolutely is to have it hold exactly, at the foundation of a description, with no exception at any scale. Within any finite setting, symmetry and continuity are both present, each to some degree. What they cannot do is both be perfect at once.

Place the two readings of the last two sections beside the theorem. Lawfulness is symmetry. Locality is continuity. A local hidden-variables theory asserts realism and locality together, and asserts them without limit; so it asserts symmetry and continuity together, and asserts them absolutely. Symmetry and continuity, asserted together and absolutely, are the first of the two forbidden pairings. The match is direct. The Incommensuration Theorem holds that no comparison can support that pairing. Bell's theorem holds that no theory that matches the quantum correlations can support that pairing. The two results have the same shape because, in the reading I give here, they have the same root.

One step in this argument should be made explicit, because it is the step that rests on the metaphysics. The Incommensuration Theorem forbids the pairing of symmetry with continuity when both compounds are claimed of the same thing, in the immanent modality. The immanent modality is the modality of interaction, process and reality, of what is happening here and now. (Chapter 4 named the three modalities: the immanent, process, what is happening here and now; the omniscient, structure, what holds the same for everyone, laid out as a pattern; and the transcendent, possibility, what makes things possible and new.) A theory of reality is a theory of what is. It has to start from the act of interaction, the immanent, and not from a pattern seen whole from outside, the omniscient, because in this book reality belongs to the act. To exist is to stand as structure, the same for every comparer, and that is omniscient; to be ob-

jective is a matter of degree, the agreement of many findings across many comparers, and that is transcendent; to be real is to be in an actual act of comparison, here and now, and that is immanent.

The first axiom of An Immanent Metaphysics says why a theory of reality must start from the immanent:

> Axiom I: The immanent is more fundamental than the omniscient and/or the transcendent. The omniscient and the transcendent are in conjugate relationship to one another.

In plain words: the act of interaction, here and now, comes first, before the pattern laid out whole and before the possibility that stands beyond; and the pattern and the possibility are *conjugates*, two things each of which can be perfected only at the cost of the other, neither one prior. Symmetry is omniscient: it is the same pattern across every setting. Continuity is transcendent: it is the unbroken, the fully specified, the same content held without a gap through a single setting, as Chapter 2 found. Interaction is immanent: it is the act happening here and now. So lawfulness and locality are conjugates, and the two of them, perfected together, are exactly what cannot be found in the immanent, which is where a theory of reality has to live. The idea that reality is absolutely lawful is inconsistent with the idea that reality is absolutely local.

In effect, any idea of what reality is must allow one of two things. It can allow several sets of physical laws that disagree with one another, each holding in different places and times; that is a failure of lawfulness, and so of symmetry. Or it can allow interactions and changes that are, at bottom, not local, and that are, at bottom, not fixed in advance; that is a failure of locality, and so of continuity. No theory of reality can assert both total generality and total determinism. In that sentence, total generality is the same law everywhere, which is symmetry perfected; total determinism is every interaction and change fixed in advance, step by unbroken

step, at every scale, which is continuity perfected. (Determinism in the full sense of Chapter 5 asks for both perfections at once, and so it is the forbidden pairing itself.)

That no theory of reality can assert both total generality and total determinism is what Bell's theorem establishes within physics. Reality can be completely lawful and somewhat non-local, or completely local and somewhat non-lawful. Reality cannot be both absolutely lawful and absolutely local. The theorem about comparison gives the form of Bell's result from the structure of comparison alone, without any reference to what the theories of physics contain. Bell found the result in quantum correlations. The Incommensuration Theorem says that the result had to be found somewhere.

Bell's Words Beside the Theorem's

The match can be laid out piece by piece, each piece of Bell's result beside the piece of the Incommensuration Theorem to which it corresponds. Realism, the definiteness of a particle's properties whatever the setting of the measurement, corresponds to symmetry. Locality, the absence of any instant influence, corresponds to continuity. The breaking of Bell's inequality corresponds to the theorem's finding that symmetry and continuity cannot both be applied absolutely together. The two options Bell leaves open, lawful but non-local, and local but non-lawful, correspond to the two survivors of the theorem, symmetry with discontinuity, and continuity with asymmetry. The general result does not depend on quantum theory. Bell's specific result does. The structural reason is the same in both.

In Bell's result	In the theorem about comparison
Realism	Symmetry
Locality	Continuity
The inequality broken	Symmetry and continuity cannot both be absolute together
Lawful but non-local	Symmetry with discontinuity
Local but non-lawful	Continuity with asymmetry

The Two Ways Bell Leaves Open

The Incommensuration Theorem does more than forbid two pairings. It names the survivors, and the survivors are the two options that Bell's result leaves standing.

The first survivor is symmetry with discontinuity: completely lawful, somewhat non-local. The laws of physics are universal, the same everywhere, and being universal is symmetry; but locality fails at some scale, and a failure of locality is a discontinuity, a break. On this option the quantum correlations are read as an instant influence: the measurement on one particle affects the correlations observed at the other. The laws are kept. Locality is not. The first survivor is the option that most physicists have taken in describing what the broken inequality means, and the first survivor is what Albert Einstein's phrase "spooky action at a distance" names.

The second survivor is continuity with asymmetry: completely local, somewhat non-lawful. All interactions are local, and that is continuity; but the laws that govern outcomes depend on the setting, and depending on the setting is asymmetry. The outcome depends on the particular measurement performed, the particular apparatus, the particular moment. There is no statement of the particle's state that stands apart from the setting in which it is measured. In practice this is the option that most inter-

pretations of quantum theory adopt, whether or not they say so: the outcome of a measurement cannot be fixed apart from the setting of the measurement.

The two majorities answer two different questions. The first majority's answer is what most physicists say the broken inequality means: an instant influence between the two particles of a pair. The second majority's answer is what most interpretations of quantum theory do in practice when working out an outcome. Most interpretations fix the outcome only together with the setting of the measurement, whatever each interpretation says the broken inequality means.

Bell's dichotomy, local or realistic but not both, is therefore recovered as a special case. The general structure has four compounds and two survivors. In the domain of quantum correlation, the four compounds present themselves as lawfulness, its failure, locality, and its failure; and the two survivors present themselves as Bell's two options. Nothing in the mathematics of quantum theory is needed to see that there must be exactly two options, or to see which two. What quantum theory and experiment supply, and what only they can supply, is the finding that nature actually breaks the inequality, and is not only permitted to break it.

Three Wishes, and Which One to Give Up

The same result can be stated a second way, and the second way makes visible a choice that the first way hides.

Picture a bank that keeps copies of its accounts on computers in three cities, so that a customer anywhere can be served. The bank would like three things at once. It would like every copy to give the latest balance whenever anyone asks, which engineers call consistency. It would like every request to get an answer, which they call availability. And it would like the system to keep working when the network between the cities is cut, which they call tolerance of partitions. In 2000 the computer scientist Eric Brewer stated that a store of data copied across machines cannot

guarantee all three at once, and in 2002 Seth Gilbert and Nancy Lynch proved it. Engineers know the result as the CAP theorem, from the first letters of the three wishes. Pick at most two.

Bell's theorem can be read as a three-way choice of the same form. A choice among three things, of which at most two can be kept, is called a *trilemma*. The first of the three properties is the *universality of causation*: every change in the world is accounted for by causes, and only by causes. The second property is the *consistency of causation*: the laws of cause, wherever they apply, are the same type of law everywhere. The third property is locality: every causal process is local, bounded by the speed of light. (A *cause*, as Chapter 6 defined it, is a consistency that a finite comparer finds between two things in the world, one following the other, again and again.) Bell's result, read this way, says: keep at most two.

The community of physicists, faced with the result, kept the first two and gave up the third. Their reasons make sense. If the same causation does not apply everywhere, then the idea of knowing the causation that holds in general, throughout the universe, fails, and with it the idea of what it means to be scientific: we have found a lawfulness, we can describe it by causes, so it can be understood and can be treated as an explanation. All of that ties back to causation as a concept. To keep causation as a concept, universal and the same everywhere, one gives up locality, and one lives with an influence that has no speed and no mechanism.

The result of Bell's theorem was very hard to accept, because the result touched a belief underneath physics: the belief in a perfect idea of causation, perfectly applied. People want their theories to be perfect, either perfectly general, with the same law everywhere, or perfectly specified, with every detail fixed. Bell required that one of the two be surrendered. Which one to surrender was a choice, and it still is.

The framework of this book describes the same choice more exactly, and here the two senses of locality set out earlier in this chapter both come into play. The third wish, that every causal process be local, is locality in the narrow sense, no signaling: every cause, and so every signal, stays

within the speed of light. Locality in Bell's broad sense asks for more: that every agreement between distant records be carried by causes traveling from place to neighboring place. So the broad sense needs the first wish as well, that every change have a cause. The framework keeps the consistency of causation wherever causation applies, and it keeps locality as physics keeps it, in no signaling; what it gives up is the universality of causation, and with it locality in Bell's broad sense. Giving up universality leaves a gap in what law covers, and a gap is a discontinuity. Keeping the consistency of law is keeping symmetry. So the resolution is symmetry with discontinuity: the survivor physics also took, described without the word "non-local". Physics keeps the first two of the three wishes, and this framework keeps the last two. Each gives up locality in Bell's broad sense: physics by admitting an influence that has no speed and no mechanism, and this framework by admitting change that has no cause. Because physics and this framework each give up locality in Bell's broad sense, physics and this framework stand on the same survivor, the first survivor: completely lawful, and somewhat non-local in Bell's broad sense. The second survivor, completely local in Bell's broad sense and somewhat non-lawful, is the other option, in which an outcome is fixed only together with the setting of its measurement. The reason this framework chooses as it does comes from looking closely at what the correlations are.

What the Correlations Are

Go back to the two detectors of the opening, each asking its one yes or no question, each keeping a record that, taken alone, is a random list, with the correlation appearing only when the two records are brought together afterward, by ordinary means and at ordinary speed. And this is the point: the correlation cannot be used to send anything. No choice made at one end changes the *statistics*, the running tallies of yes and no, recorded at the other. The experimenter controls only the question asked, and never the answer. Any attempt to fix the answer in advance, by measuring the

particle before it reaches the detector, destroys the entanglement on which the correlation depends, just as, in the story of the enchanted decks, a deck stacked in a chosen order instead of shuffled loses its agreement with the other deck. Nothing crosses the distance between the two detectors. There is no signal. The correlation is observed after the fact, as something that has happened, and it is not an influence traveling from one place to the other.

What, then, has been broken? What has been broken is the expectation that a correlation of this strength must have a causal account, either in the shared past of the two particles, or in an influence passing between them. Bell's theorem shows that neither account is available. There are more correlations in nature than any coherent causal account can explain. The correlations are, in a strict sense, non-causal: no cause produced them. Locality, in the sense that physics also holds (no signal, no influence, no interaction crosses a distance in no time), is untouched: no signal does cross. An agreement between separated records that goes beyond every causal account, that cannot be used to signal, and that leaves the tallies at each end independent of the setting at the other, is a *non-causal correlation*: real, and not an influence.

The mainstream option reads a correlation that has no cause as a cause that is not local, and it does so because of an assumption made earlier: that every change is causal, that there is no such thing as change that is not causation. Given that assumption, the correlation must be an influence; and since the influence cannot be local, locality falls. Only the assumption that every change must be causal turns a correlation into an influence. Remove the assumption, and locality does not need to fall. What falls instead is the universality of causation: there is change in the world that causation does not account for.

The metaphysics has a name for change that is not causation. The name is *choice*, in the general sense that *An Immanent Metaphysics* gives the word, a sense far wider than a person's choosing; and *change* is the larger class that holds both causation and choice. Causality is omniscient, because it is the repeatable pattern, the same for every comparer; choice is

transcendent, because it is the arrival of what was not already contained in what came before; and change is immanent, because it is the actual happening here and now, in which pattern and the new meet. Some of what changes is accounted for by causal knowledge, and some is not, and what is left over is real. Chapter 6 named that remainder the residual, what no lawful description fixes, and called the outcomes in it hard randomness: outcomes that no earlier specification fixes, however complete, for any comparer whatsoever.

That some change is real and not caused is why this framework chooses as it does. Locality, taken as physics takes it, is the statement that nothing is compared across the edge of comparison: no signal crosses a distance in no time. Locality as physics takes it is a statement about the edge of the *domain of comparison*, the whole range in which comparisons can be made, and the correlations stay inside that edge. The universality of causation is the demand that every change be pattern, and pattern is omniscient; the first axiom places the act, in which change happens, before the pattern. When the two demands conflict, the framework keeps the edge of comparison and gives up the claim that all change is causal. In the words physics uses, the correlations are non-local, and they are confined inside the rule of no signaling. In the words of this book, they are not caused, and they carry nothing.

The two senses of locality now fall into place, as the three wishes showed. Locality in Bell's broad sense is a demand on causes, and I give it up together with the claim that every change has a cause. Locality in the narrow sense, no signaling, is kept by physics and kept by this book. Physics and this framework stand on the same survivor, symmetry with discontinuity, and they describe the gap in that survivor in two ways, the second the more exact: physics calls the gap a non-local influence that can send nothing, and I call it change that no cause produced.

Hidden Instructions, and the Last Escape

Bell's inequality is exactly the ceiling on how much correlation a shared past can supply, and nature goes over that ceiling. Two attempts to keep both lawfulness and locality deserve a direct reply, and hidden variables come first.

The idea of hidden variables was that something in the source, some definite state we do not know, might account for the correlations: that the two particles carry, from their common origin, the answers they will later give. The idea of hidden variables is the shared-past account. An event at the source gives rise to two later events, one at each detector, and the two later events agree because both descend from the one earlier event, as two children share a parent. Bell's ceiling is exactly the most agreement that such a shared parent can give, and the measured agreement is more.

There is a deeper point beneath the arithmetic. Hidden variables were proposed to be the basis of the causes that produce the outcomes. Once the observed correlations have gone past what any causal account allows, it no longer matters what the variables are. A variable cannot restore causation after the correlations have exceeded what any causal account allows. The program of local hidden variables was a passing attempt to account for the result by causes, and it failed; that failure is what the summary "no local hidden variables" records.

Two accounts remain open to a determinist, someone who holds, in the sense of Chapter 5, that everything is fully and exactly fixed. The first of the two accounts is the *pilot-wave theory* that David Bohm published in 1952. It gives each particle a definite position at every moment, hidden from us and guided by a wave, and it gives up locality: the outcome in one laboratory depends on the setting in the other. The second of the two accounts is *superdeterminism*: determinism extended to the experimenters' choice of settings, so that the setting and the particle share one causal history. It tries to keep both lawfulness and locality by asserting that the experimenter's choice of question was fixed, together with the particle's answer, in a past that determined both. On this account the correlation is

neither an influence nor a hidden state of the two particles alone; it is a shared fixing of everything, the choices of the experimenters included, at every scale. Superdeterminism thus fixes the experimenters' settings together with the particles' states, and so it gives up the independence of the settings on which every experiment relies.

The Incommensuration Theorem excludes both accounts as accounts of a whole world. A whole world completely fixed by one law perfects symmetry and continuity together, and the theorem forbids that pairing. Superdeterminism adds a second fault: it folds choice into causation, and choice and causation are distinct, and neither can stand in for the other. The Bell experiments, by themselves, narrow both accounts and leave each one possible by logic alone; the exclusion of both comes from the structure of comparison. A deterministic formalism remains available as a description, as Bohm's is. The findings that such a formalism predicts have the statistics of quantum theory, and no finding ever shows that the formalism is deterministic. An account excluded as an account of a whole world stays available as a description, because a description is a map of the world, and not the world it maps; Chapter 4 put the point this way: the determinism of a formalism is the determinism of a map, and not of a territory. The Incommensuration Theorem excludes a whole world completely fixed by one law.

Read by its structure, superdeterminism is the attempt to keep both symmetry and continuity at once, and absolutely. It keeps lawfulness by making law universal, reaching down even to the choice of what to measure. It keeps locality by making every process continuous, with no gap between the past and any outcome. Superdeterminism is determinism extended to the observers, and so the argument of Chapter 5 excludes it before any experiment is consulted: the pairing superdeterminism requires is the pairing the theorem forbids. Bell's theorem and the Incommensuration Theorem meet at the same point. The superdeterministic escape is not available. Superdeterminism is not a further option on the list of three; it

is the demand to have all three wishes of the trilemma at once, and that demand is exactly what was ruled out. Chapter 11 takes up superdeterminism in full.

Where the Method of Science Reaches

Reading lawfulness as symmetry and locality as continuity has a consequence for science itself, beyond Bell.

The method of science can make observations about causes. It can establish that similar conditions give similar results, that regularities hold, that patterns repeat. What the method of science can establish is causal knowledge, and gaining causal knowledge is what the method is for. The method of science cannot, even in principle, observe continuity or locality directly. An experiment observes that in this case there was no jump; no experiment observes that there are no jumps at all. Nor can the method make a direct observation of symmetry, even though it must assume symmetry in everything it does, because the assumption that a result will repeat is what allows the result to count. No single experiment, and no finite collection of experiments, could establish the absolute and universal truth either of continuity, that all interaction and change is local, or of symmetry, that reality is lawful everywhere and always.

No finite collection of experiments could establish the absolute truth of the beliefs that science takes on without observing them, either. The background beliefs that science has taken on are six: material monism, that everything is matter; causal determinism, that everything is fixed by cause; locality; strong objectivity, that the facts stand wholly apart from any observer; universality, that the laws hold everywhere; and epiphenomenalism, that the mind is a by-product that does nothing. Each of the six is challenged by findings within science itself, and the challenges cannot be settled from inside the framework of science. Bell's finding is the challenge this chapter has followed, and it presses on two of the six, causal determinism and locality.

This challenge marks where the reach of the method ends, and it is no complaint against the method. Those parts of the physical world that are real and yet by their nature cannot be repeated, cannot be observed more than once, would be non-lawful, and not symmetric in time and space, and so they lie outside what the method of science can study. Bell's theorem is one place where the method met that edge and reported what it found there: a correlation it could observe and could not explain. The framework says that the method must meet such an edge, because the method rests on symmetry, and symmetry perfected is not to be had in the world where the method works. Science can establish causal knowledge. The fully deterministic and the fully symmetric are beyond its reach by its structure, and so is the fully continuous.

What Bell Found, Read Through the Theorem

We now know what Bell found. Bell's theorem is an instance of the Incommensuration Theorem. Lawfulness is symmetry. Locality is continuity. The theorem holds that symmetry and continuity cannot both be perfected in the same comparison. Bell shows that lawfulness and locality cannot both be kept in the face of quantum correlations. The general limit contains the specific one. The two options Bell leaves open are the two pairings the theorem allows.

The impossibility is a feature of comparison itself; it is not peculiar to quantum theory, and it is not a product of the mathematics Bell used. Any domain of comparison that demands both laws that never vary anywhere and interaction that is perfectly continuous will meet the same limit. Bell met it in the correlations of separated particles. The Incommensuration Theorem says the limit is there to be met everywhere.

For the question of this book, what cannot be and why, Bell's result, read through the Incommensuration Theorem, matters at once. The strongest agreement that nature shows between two distant places is real, and it carries no signal. Nothing chosen at one detector reaches the other.

The correlations are confined inside no signaling, and reality is where signals travel. To be real, in this book, is to be in an actual act of comparison; two records made far apart are compared only when a signal brings them together, at ordinary speed, as the records of the two detectors were brought together. So the real is the region that signals reach, and the correlations, which appear only after such a signal has joined the two records, lie inside it.

Summary

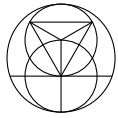
Bell's theorem (1964), confirmed by experiment from 1982 onward, is an instance of the Incommensuration Theorem: realism is a symmetry, locality is a continuity, and the two cannot both be absolute together. The structure of comparison gives Bell's two remaining options and which two they are; quantum theory and experiment give the finding that nature actually breaks the ceiling. No theory of reality can assert both total generality (the same law everywhere) and total determinism (every interaction and change fixed in advance, at every scale). Faced with that, one may keep the universality of causation and give up locality in Bell's broad sense, keeping only no signaling, which is what most of physics has done; or one may keep locality as physics keeps it, in no signaling, keep the consistency of law wherever law applies, and give up the claim that all change is causal, which is what this framework does, because the correlations are real but they are not signals, and because reality is where signals travel. The two descriptions name one survivor, symmetry with discontinuity, completely lawful: physics calls the correlations non-local and confines them inside no signaling, and I say the same correlations are not caused and carry nothing.

Each claim of this chapter is stated together with the kind of claim it is. That Bell's theorem is an instance of the Incommensuration Theorem is a reading of known physics: it gives the structural reason the result holds, and it shows that the reason reaches beyond the mathematics of quantum

theory. Finding Bell's two options as the theorem's two survivors is a reading of the same kind. The account of the correlations, which keeps no signaling and gives up the universality of causation, is grounded in the metaphysics, in the first axiom, which places the act before the pattern, and it describes the survivor physics took, symmetry with discontinuity. That no whole world is completely fixed by one law, so that Bohm's account and superdeterminism are both excluded as accounts of a whole world, is a structural exclusion: it says what physics will not find, and it says why.

Whichever of the two descriptions is chosen, the method of science, built on the assumption of symmetry, cannot establish that assumption absolutely, nor the continuity it also relies on. Science establishes causal knowledge. The conflict of lawfulness and locality is a structural necessity in the nature of comparison, and no accident of the quantum world; its scope is every domain in which comparison is made.

Superdeterminism still has its defenders, who hold that everything, the experimenters' choices included, was fixed together from the start, so that nothing in the world need ever be uncaused. What exactly do they claim, and why is there no way out along that road?



Chapter 11

The Escape That Is Not There

A Quiz Arranged in Advance

Could everything, including the experimenters' choices, have been fixed together from the start? Chapter 10 ended on that question. Bell's experiments showed that two particles measured far apart agree in a pattern that no separate instructions, carried by each particle from their common source, could give them. One way out has been offered: say that the experimenters' questions were never free, and that what each experimenter chose to measure was fixed long ago, together with what each particle would answer.

Picture a quiz on television. A contestant answers question after question correctly, on every subject, faster than anyone could look anything up. The producers want to be sure that no one is feeding her the questions, so they stop writing questions in advance. Each question is now picked at the moment it is asked: first by the roll of a die, then by a machine that draws each number from a single quantum event, then by the color of the light arriving that second from a distant star, then by whichever button thousands of viewers happen to press. She still answers correctly.

One suspicious viewer still says she is being fed. The die, the machine, the star and the viewers, he says, were all arranged from before the quiz began, so that each of them would pick exactly a question she was ready for. Nothing anywhere else in the world shows the arrangement. It shows only in the quiz.

In a Bell experiment, the questions are the settings of the two detectors (the angle at which each detector asks its yes or no question), and the answers are the outcomes the particles give. *Superdeterminism* is the proposal that the suspicious viewer is right about the world: the setting an experimenter chooses and the state of the particle the experimenter measures share one causal history, so that the choice of question is fixed along with the answer. Put shortly, superdeterminism is determinism extended to the experimenters' choice of settings. Its purpose is specific. Bell's theorem shows that no account that is both lawful and local (lawful: one rule holds the same everywhere; local: the outcome in one laboratory does not depend on the setting chosen in the other) can give the measured matches, and superdeterminism lets such an account give them after all, by fixing the questions together with the answers.

Superdeterminism is excluded by the structure of comparison, before any experiment is run. It asks for everything determinism asks for, symmetry and continuity perfected together, which is the very pairing the Incommensuration Theorem forbids; and then it asks for more: it asks that choice be nothing but cause, and that request is a mistake of kind.

The Case as Its Holders Make It

Superdeterminism is the one refuge left to a determinist after the catastrophic theorem of Chapter 9, which showed that determinism has no partial form: if determinism fails anywhere, at any scale, it fails at every scale. A position that is the last refuge of a whole view of the world deserves to be heard as those who hold it make it.

Gerard 't Hooft, in his book *The Cellular Automaton Interpretation of Quantum Mechanics* (2016), holds that at its deepest level the universe is a deterministic system, like a vast grid of cells, each cell updated at each step by a fixed rule, and that the probabilities of quantum theory describe our ignorance of the grid's state. On his view, the settings experimenters choose and the states of the particles they measure are correlated through their common past, so that the assumption of free choice fails without anyone noticing. Two things are *correlated* when they are linked, so that knowing one tells you something about the other.

Sabine Hossenfelder and Tim Palmer, in "Rethinking Superdeterminism" (2020), have argued the position with care, and Hossenfelder set it out again that year in "Superdeterminism: A Guide for the Perplexed". Their argument turns on one assumption of Bell's theorem. *Statistical independence* is the assumption that the detector settings are not correlated with whatever hidden state the particles carry. John Bell's theorem of 1964 makes this assumption, and Bell himself discussed it in his essay "Free Variables and Local Causality". Hossenfelder and Palmer argue that statistical independence is a physical assumption like any other, that it can fail without any conspiracy, and that a theory which drops it could restore a deterministic and local account of what quantum theory describes. They propose experiments to look for such a theory.

I give the case of Hossenfelder and Palmer at full strength. Bell's theorem and the free will theorem (Chapter 9: if the experimenters' choices are not fixed by everything before them, neither are the particles' responses) both assume that the settings are independent of what is measured. Drop that assumption, and a deterministic, local world is consistent with every observation made so far.

Hossenfelder and Palmer have made the clearest recent statement of superdeterminism, and in this chapter I answer their statement. In public discussion the position is often set inside a proposal of three stages, and I take the proposal in that form, because each stage bears on the others. The first stage says that the paradoxes of cause usually blamed on influences faster than light can be avoided if time has a preferred direction, one that

prevents anyone from building a loop of causes, in which an effect comes round to become its own cause. The second stage says that quantum theory is not fundamental: it is a statistical average over a deeper theory (what you see when you add up many cases and look only at the overall rates), the deeper theory is deterministic, and the settings of the measuring devices are among the things it fixes. That second stage is superdeterminism, and it is the stage that the works just named develop. The third stage says that where the deeper theory departs from the predictions of quantum theory, those departures might let information be sent faster than light, and that this might even explain why we hear nothing from other civilizations: we are listening on the wrong channel.

The proposal gets more right than its critics usually allow, and I affirm three of its points. Time has a direction. Quantum theory is not the final description. The usual account of why the speed of light is a limit runs in a circle. Each of these three, followed through, tells against superdeterminism, and later sections of this chapter show how.

Each stage has its own answer. The first stage is right in what it says of direction, and I agree with it; a single preferred way of slicing the world into moments, one "now" cut through all of space and shared by every observer, would ask for more than a direction, and I do not grant it. Directed time, as a later section shows, cannot be had together with superdeterminism. The second stage is excluded by structure, and in most of this chapter I show why. The third stage concerns signals, and Chapter 12 answers the third stage: the speed of light is, at root, the speed of no one thing. The speed of light is the outer edge of the range within which anything can be compared with anything, so a request to signal past that outer edge is badly formed (as distinct from merely hard to meet). Superdeterminism and the signal question are two halves of one proposal, answered in this chapter and the next. The answer to the signal question holds whether or not superdeterminism could be had. Superdeterminism cannot be had, and a superdeterministic world, were one available (none is), would still open no channel.

The impossibility of superdeterminism rests on the structure of comparison, and quantum theory plays no part in it. That impossibility would hold in any physics in which anything can be measured at all. The impossibility of superdeterminism is structural in exactly this sense: it holds in any physics in which anything can be measured at all. Because the impossibility holds in any physics, the experiments that superdeterminism's defenders recommend, searches for departures from the quantum predictions, cannot bear on that impossibility either way.

Everything Determinism Asks For, and More

Set the two positions side by side. Ordinary *determinism*, as Chapter 5 described it, is the claim that everything is fully and exactly fixed. Stated in full: the state of every physical system is fixed by the conditions before it together with law; a complete specification of the state at one time, carried forward by the same lawful rule in every setting, yields exactly one state at every other time. Laplace's demon, who knows every position and every speed at one instant, reads every other instant off that knowledge.

Chapter 5 read this as a demand for two things at once. That the same lawful rule hold in every setting is sameness of content across difference of context: the same rule, whatever the place or the moment. Sameness of content across difference of context is *symmetry*, and determinism asks for symmetry perfected. That the state be specified without a gap at every point is sameness of content within sameness of context: hold the setting fixed and nothing jumps. Sameness of content within sameness of context is *continuity*, and determinism asks for continuity perfected too.

The Incommensuration Theorem, a structural fact about the nature of knowing, says that of the four ways of joining sameness and difference to content and context, symmetry and continuity are one of the two pairs that cannot both be applied absolutely at once. Within any limited region, the two can both be present to some limited degree. What they cannot do

is both be perfected together. Determinism asks for exactly that. So determinism is not false as a matter of fact; it is out of reach as a matter of structure.

Superdeterminism keeps all of that and adds a further clause. Determinism fixes the states of the systems under study; superdeterminism fixes also the choices of the experimenters about what to study: which detector setting to use, which question to ask. The apparent freedom to choose the measurement is declared an illusion. The choice of setting and the state of the particle are declared to share a common causal history, reaching back, if need be, to the initial conditions of the universe, the state of everything at the very beginning.

The clause has one job. Bell's theorem, together with the loophole-closing experiments that Chapter 8 described, leaves no room for an account of the matches that is both lawful and local, provided that the settings are statistically independent of whatever hidden state the particles carry. (The loophole-closing experiments, recall, closed every named channel by which a cause could have linked the two sides: the detectors were set far enough apart, the settings were chosen fast enough and randomly enough, and the particles were detected efficiently enough.) Superdeterminism denies the independence. Fix the settings along with the states, and the matches could have been arranged in advance, with lawfulness and locality both kept.

The argument from determinism is now short, and I give it whole before adding anything to it. Determinism is out of reach. Superdeterminism demands everything determinism demands, and then demands more. So superdeterminism is out of reach *a fortiori*, a Latin phrase meaning "from the stronger reason", or, in plain words, all the more. A position that requires a forbidden pair does not become available by requiring the same pair over a larger region. Whatever cannot be had for the particles cannot be had for the particles, the apparatus and the experimenter together.

Superdeterminism is therefore excluded before any experiment is run. The Incommensuration Theorem excludes it directly, and for the same reason it excludes determinism: a reason about comparison, and not a reason about quantum theory.

The argument carried over from Chapter 5 settles the question by itself. A reader who accepts Chapter 5 needs nothing further. The added clause has a character of its own, though, and there is something specific to say about it. The added clause is a claim of a different kind from determinism's claim, and not only a larger claim of the same kind. The added clause also fails in a different way.

Choice Made Into Cause

Ask how many kilograms a promise weighs. The question has good grammar, and it has no answer of the kind it asks for. A promise can be made, kept or broken; weight is no part of what a promise is. Someone who insists that a promise is really a kind of weight has not said something false about weights. He has failed to say anything about promises at all. A question or a claim that applies a concept where the concept does not belong is a *category error*: an error of kind, as distinct from an error of fact.

Chapter 7 showed that the real is understood through a triple of ideas: choice, change and causality. The three are *distinct, inseparable and non-interchangeable*: three different things, never found apart, none able to stand in for another. The three are distinct, inseparable and non-interchangeable in the same way as the three modalities of Chapter 4 are, and by the same axiom.

Causality is the lawful, repeatable pattern that relates events to one another. Causality is what an observer sees when the same kind of cause is followed by the same kind of effect, again and again, in many settings. Causality has the nature of the omniscient (the modality of structure: what holds the same for everyone, laid out as a pattern). Choice is the new: the arrival of what was not already contained in what came before,

the single act that is only itself. Choice has the nature of the transcendent (the modality of possibility: what makes things possible and new). Change is the actual happening, the passing from one state to another: the interaction in which pattern and the new meet, and in which either of them can be noticed at all. It has the nature of the immanent (the modality of process: what is happening here and now), and it is the wider kind within which the other two are found.

I take that alignment from the metaphysics as fixed: causality goes with the omniscient, change with the immanent, choice with the transcendent. The third axiom, which Chapter 4 quoted, holds that the kinds and the instances of the immanent, the omniscient and the transcendent are distinct, inseparable and non-interchangeable. The triple of choice, change and causality inherits that standing directly, because each of its members is an instance of one modality in the domain of the real.

Superdeterminism, in its added clause, is the attempt to collapse choice into causality. It says that what appears as the experimenter's choice of setting is in fact nothing but a mechanical consequence of the conditions before it: that the new is only the old unrolling, that the transcendent is the omniscient seen from too close. I identify this as a category error, and I want to say exactly what kind.

The error is the attempt to reduce one modality to another across a boundary that the axioms declare impassable. It is not the error of saying something false in a domain where truth and falsity apply. Choice cannot be reduced to causality, any more than causality can work without change, or change can happen with nothing new in it. The three remain three. A proposal whose content is that one of them is really another has not succeeded in saying anything about the real, just as the claim that a promise is really a weight says nothing about promises. Such a proposal is not one that could turn out to be true.

The same conclusion can be reached from the theorem itself, without leaning on the axioms as stated. Symmetry and continuity cannot both be perfected. So what happens cannot be both fully lawful and fully spe-

cified: being fully lawful is symmetry, which is causality, and being fully specified is continuity, everything fixed with nothing left open. There is therefore always a *residual*: as Chapter 6 first put it, what no lawful description fixes; a range in which the outcome is not fixed by the pattern. Chapter 6 named that residual choice. (Chapter 7 showed that one and the same unfixed event is randomness when seen from outside and choice when seen from inside.)

The Incommensuration Theorem requires the residual to be there, in any world in which anything can be compared at all. The residual is not our ignorance of a pattern that is there all along. Superdeterminism is the claim that the residual is empty. The Incommensuration Theorem holds that the residual cannot be empty.

So superdeterminism fails twice. It fails as determinism fails, by asking too much, and it fails in its own way, by asking the wrong thing. Ordinary determinism already collapses choice into causality for the systems it describes: it asks that pattern be total, and a total pattern leaves no residual for choice. Superdeterminism extends the collapse to the comparer, the one who makes the comparison: it asks that the stance from which a measurement is made, the choice of what to compare, be pattern too. Determinism denies that anything new arises in what is measured. Superdeterminism denies that anything new arises in the act of measuring too, and so in the very comparison that makes a measurement a measurement.

Whoever, or Whatever, Turns the Dial

The argument of this chapter stands without any appeal to free will. Superdeterminism is often presented as a denial of the experimenter's free will, and defended on the ground that physics has no business assuming free will in the first place. Nothing here requires that a human being who chooses a detector setting does so by an uncaused act of the will.

The quiz at the start of this chapter was drawn from what physicists have done. In the strongest Bell experiments, the settings have been taken from quantum random number generators, machines that draw each number from a single quantum event; from the colors of light arriving from distant stars, in the Cosmic Bell Test of Johannes Handsteiner and colleagues in 2017, and from quasars, the brilliant centers of galaxies so far away that their light left them before the earth existed, in the test of Dominik Rauch and colleagues in 2018; and from the button presses of many thousands of people, in the BIG Bell Test of 2018. The point of every one of these choices was the same: that no plausible common cause could reach both the setting and the particle.

The claim of this book is about the residual, and not about the experimenter. In any world in which comparison happens, something is not fixed by the pattern, and whatever picks the setting draws on that something in the same way the particle does. Superdeterminism has to say that nothing is left unfixed by the pattern. That nothing is left unfixed by the pattern is the claim that the Incommensuration Theorem excludes, and the theorem excludes that claim whether the setting was chosen by a person, by a photon (a single particle of light), or by a star.

Why the Residual Is Not Ignorance

The defender of superdeterminism has a natural reply, and it should be met directly. The residual, the reply says, is only what we do not yet know. The initial conditions of the universe are unknown to us in detail, the hidden variables are hidden, and a complete enough description would fill the gap. On this reply, the residual is a gap in our knowledge, and the book has merely given ignorance a dignified name.

The reply mistakes the kind of gap in question. A gap in knowledge is a gap between a description that exists and our having it, like a sealed letter whose words are already written and which we have not yet read. The theorem says something of another kind. It says that a complete descrip-

tion, with every state specified and every rule holding without exception, describes no comparison at all, and a world is what comparisons are of. It does not say that the complete description is hard to obtain.

Follow the complete description through. Suppose every state were specified without a gap, and every rule held without exception. Then nothing would remain to be measured, because a measurement is the interaction in which content is compared across contexts, and both kinds of sameness, sameness within a setting and sameness across settings, would already be total. A description in which the residual is closed is a description in which nothing happens. A description in which nothing happens cannot serve as the superdeterminist's account of what happens in a laboratory.

The same reasoning shows why *hidden variables*, whatever they are taken to be, cannot help. (Chapter 10 met them: some definite state in the source, unknown to us, which might carry from the particles' common origin the answers they later give.) Chapter 10 showed that a variable cannot restore causation once the matches have gone past what causation can supply. The program of local hidden variables tried to account for the matches by causes, and the measured matches went beyond what any account by causes allows. Superdeterminism is a response to that failure, and it is the wrong kind of response. It adds no variable. It removes the independence of the settings from what is measured, the independence that made the experiment an experiment. A variable that is correlated with everything is not a hidden cause. It is the whole of the past, renamed.

The suspicious viewer of the quiz has not found a hidden earpiece. He has taken the whole history of the studio, the die, the machine, the star and every viewer at home, and given all of it a new name: "the earpiece".

The Conspiracy and the Uncaused

Ordinary economy of explanation settles the choice between superdeterminism and hard randomness before the structural argument is even brought to bear. That economy has a name, *Occam's razor*, after William of Ockham, an English philosopher of the fourteenth century: when two accounts fit what is seen, prefer the one that needs fewer assumptions.

Chapter 8 told how the loophole-closing tests left two options standing. One is superdeterminism: everything fixed at literally every scale, including the settings of the detectors, the random number generators that chose them, and the experimenters who built the generators, with the arrangement of the laboratory and the state of the source fixed together by a common past. The other is *hard randomness*, the subject of Chapter 8: an outcome that no prior specification fixes, however complete, for any comparer whatsoever.

Set the two side by side, as hypotheses. The first says: the whole of the past, at every scale, has been arranged so that the choices of experimenters happen to conspire with the states of particles to produce exactly the matches that a non-causal account predicts (an account in which no cause produces the matches), while leaving no trace of the conspiracy anywhere else. The second says: some of what happens is not caused. Of the two hypotheses, the second is simpler by a very large margin. We accept hard randomness. We live with it.

The razor cuts sharply here for a plain reason. A position that has to correlate every setting with every state, including settings taken from starlight that left its star before the earth existed, has made itself safe from any experiment whatever, and ordinary economy takes it off the table before the theorem does. The theorem then says why the razor was right.

Chapter 10 met the two deterministic accounts offered to keep every outcome fixed after Bell: David Bohm's pilot-wave theory, which gives up locality, and superdeterminism, which gives up the independence of the settings. Chapter 10 also gave two means of settling whether either of those two accounts can stand. The Bell experiments narrow both accounts

and leave each one possible by logic alone. The Incommensuration Theorem excludes both as accounts of a whole world; a deterministic *formalism*, a body of mathematics used to describe something, stays available as a description, as Bohm's is, because the determinism of a formalism is the determinism of a map, and not of a territory. To those two means this section has added a third, Occam's razor, which bears on a different pair of options, superdeterminism and hard randomness, and settles the choice between those two before the structural argument is brought to bear.

A World That Swallows Its Own Evidence

Every experiment presupposes that its settings are independent of what it measures. That independence makes a reading evidence about the thing measured, and not one more product of whatever produced the setting. The quiz proves something about the contestant only because the questions were not picked to fit her. Superdeterminism, as a claim about all events, removes that condition everywhere. Within a superdeterministic universe, the very idea that one thing is evidence for another, which requires that the evidence be something other than an outcome fixed in advance, has no place. No observation could count as evidence for superdeterminism, and no observation could count as evidence for anything else either.

Superdeterminism is offered as a claim about physical correlations, a claim about the world as it is for everyone, and claims of that kind are judged by observation. As a claim about every event, it is out of observation's reach: no observation could tell a superdeterministic universe from one in which not everything is fixed, because in the superdeterministic universe the observations themselves are fixed by the initial conditions, and any evidence that seemed to count against the position is itself a product of those conditions. Evidence never refutes the position, because the position has already accounted for every possible piece of evidence as a product of the very determinism it asserts.

In the way it absorbs all evidence, superdeterminism is equivalent to *idealism*, the view that all reality lies within the mind of the one who perceives it. Idealism places all reality within the mind of the perceiver. Superdeterminism places all events within the unfolding of the initial conditions, and makes every appearance of freedom a product of necessity. Neither can be refuted, and for the same reason: both have absorbed every possible counterexample into themselves in advance. The two are alike in how they absorb evidence and opposite in where they place the final authority. Superdeterminism places it in the omniscient, the fixed structure of the initial conditions, which is the realist limit, the extreme at which all that is real is fixed structure, the same for everyone. Idealism places it in the transcendent: in the mind of the one who perceives, the single self that persists as itself, which Chapter 2 placed in the transcendent.

By the same token, both are equivalent in structure to an older thesis. In 1857 the naturalist Philip Gosse published *Omphalos*, a Greek word meaning navel, arguing that the world was created with the appearance of age, of history and of natural process, none of which had happened: the first man made with a navel, though he was never born. Superdeterminism, idealism and the appearance of age share one shape: each has explained every possible clue in advance, as part of itself.

Superdeterminism is also a *totalization*: a claim that could hold only from a view over all events at once, the view from everywhere at once that Chapter 5 described, with no possible comparison behind it. It places all creation, the coming into being of what is new, in the initial conditions of the universe. Science, for its own convenience, sets creation aside, as I put it in my book *Comparison; The Necessary Basis*, "long ago and far away"; superdeterminism completes that move. The metaphysics of this book does not make it.

The rational position is that *indeterminism* is real: not everything is fixed in advance.

The Arrow That Needs an Open Future

The first stage of the proposal is right, and the agreement turns out to be fatal to the second. The first stage observes that time has a preferred direction, and invokes that direction to prevent loops of causes. I agree. Time has a direction, and that direction comes before the geometry of space; it is not obtained from that geometry. The first axiom, in its own wording, holds that "the immanent is more fundamental than the omniscient and/or the transcendent": process comes before structure and before possibility. Change, which is time as process, has the nature of the immanent, while pattern, which is space as structure, has the nature of the omniscient. Read through the axiom, the direction of time comes before the geometry of space. Chapter 13 develops this in full, and shows that there is an arrow of time prior to the one given by heat spreading out and disorder growing.

I hold the direction of time as one member of the *bound triple* of Chapter 8, with hard randomness and subjectivity (the stance from which a comparison is made): three things, distinct, inseparable and non-interchangeable, each requiring the next. The direction of time is never a lone fact that can be added to any picture. The consequence for the proposal is direct. The proposal cannot have the directed time it needs in its first stage together with the determinism it demands in its second. The direction of a process is bound up, in its very make-up, with the arrival of the new: a process that goes somewhere is one in which what comes next was not already wholly there. A universe in which everything is fixed by the conditions before it has no direction in time of its own.

Such a universe is a static block. Picture it as a solid slab of glass with every moment of history set inside it, layer by layer. Nothing in the slab moves. One end of it is called "earlier" and the other "later", but those words are labels that a reader of the block attaches to it; they are not a feature of any process, because in the block nothing is in process.

So the first and second stages of the proposal conflict, and they conflict for exactly the reason the Incommensuration Theorem gives. They are not two separate commitments that happen to be held together. Directed time is *asymmetry*: difference of content across difference of context, the later not the same as the earlier. Asymmetry is the survivor that pairs with continuity (Chapter 3: continuity with asymmetry, smooth and going somewhere). Superdeterminism asks for symmetry and continuity together, and so leaves no asymmetry over to be the arrow. The proposal needs the arrow to save causality from loops, and then needs the block to save locality from Bell, and it cannot have both, for the same reason that no one can.

The Circle, and Where to Break It

The strongest part of the case made for the proposal is its diagnosis of the usual position, and the diagnosis is right. The usual position runs in a circle. The speed of light must be a limit on signals; therefore no deeper theory can permit signals past it; therefore quantum theory, whose own no-signaling result respects the limit (its correlations, however strong, carry no message), is fundamental; therefore the speed of light is a limit. Each step is taken as support for the next, and the circle closes without anything outside it bearing the weight. The usual position is circular, as the diagnosis says.

The superdeterministic escape breaks the circle at the claim that quantum theory is fundamental: it denies that quantum theory is the deepest description, posits a deterministic theory beneath it, and accepts whatever follows about signals.

I break the circle at a different place, and I think at the right one. The speed of light is a limit because of the structure of comparison: it is the large-scale edge, the macroscopic boundary, of the range within which comparison, and therefore interaction, can happen at all. It is not a limit because quantum theory is fundamental. Chapter 12 builds that argu-

ment. Once the limit rests on the structure of comparison, it no longer needs quantum theory to hold it up, and the circle has nothing left to close.

On the other point, whether quantum theory is fundamental, the book agrees with the proposal that it is not, though not in the proposal's sense. Picture two craftsmen who have traded workshops: each works with his own tools on the other's material. The two great theories of physics stand that way, and the next two paragraphs name the tools and the materials. Quantum theory is a partial description. It is symmetry with discontinuity (the same everywhere and countable), the tool proper to the omniscient, applied to the microscopic, the domain of the very small, which is the transcendent domain: the domain of possibility and probability, of what may happen rather than what has. Its tool, the countable pattern that is the same everywhere, is a tool of the omniscient; its material, what may happen, belongs to the transcendent.

General relativity, the theory of gravity as the shape of space and time, is the other partial description. It is continuity with asymmetry, the tool proper to the transcendent, applied to the macroscopic, the domain of the very large, which is the omniscient domain: the domain of pattern, of space, of what holds across. Its tool, the smooth curve that goes somewhere, is a tool of the transcendent; its material, the pattern of space, belongs to the omniscient. Each theory, then, brings the tool of one modality to the material of the other. Quantum theory and general relativity are the two survivors of the theorem, each finding one of the two pairings the theorem allows. They are *conjugate* descriptions within a single domain of comparison (Chapter 3: each can be perfected only at the cost of the other, and neither is prior): each captures what the other must leave out, and neither can be folded into the other. Chapter 17 shows why this crossing, each tool applied to the other's domain, is required.

So quantum theory is not the final word, and the reason is that quantum theory is one of two survivors, not that a deterministic theory lies beneath quantum theory. There is no theory beneath quantum theory whose findings perfect both survivors at once, and such a theory is what

the Incommensuration Theorem precludes. The superdeterministic escape from the circle replaces one error with another. The usual position errs by resting a structural boundary on one particular body of mathematics. The superdeterministic escape errs by proposing, as the deeper ground, exactly the pairing that no theory can find: symmetry and continuity perfected together. I rest the boundary on comparison, and leave the two theories where they are, each doing what the other cannot.

Whatever the Experiments Find

Hossenfelder proposes that superdeterminism be tested by looking for departures from the quantum predictions where a deeper theory might show through: in small systems, in repeated measurements, at short times. I have no objection to the experiments. Physics should always look where it has not looked.

The argument of this chapter stands whichever way their outcome falls. Suppose no departure is found. Then quantum theory stands where it stood, and the book's reading of it stands as well. Suppose a departure is found. Then there is a physics beyond the present one, and I say of it in advance exactly what I say of the present one: that it is a description of comparisons; that it therefore has the six intrinsics (sameness and difference, content and context, the subjective and the objective); that symmetry and continuity are compounds of them; and that the two cannot both be perfected in it.

A departure from quantum theory would show a new theory at work. It would not be a deterministic one, and it would not be a superdeterministic one, because no comparison can be an instance of either.

The argument against superdeterminism, then, depends only on the theorem. The argument against signals faster than light, which Chapter 12 gives, depends only on reading the speed of light as the large-scale boundary of the range of comparison. Neither depends on quantum the-

ory. The conclusions of the two arguments are conditions on any body of mathematics that describes the world, and they follow from no body of mathematics.

The argument rests on two things only. One is the *Root Tautology* of Chapter 1: interaction, comparison and relation are one act, named three ways. The other is the Incommensuration Theorem: symmetry and continuity cannot both be perfected at once. A reader who wishes to keep superdeterminism must deny one of the two, and neither denial can be made without comparing. To dispute the Root Tautology or the Incommensuration Theorem, one must compare, and comparison has the structure that the Root Tautology and the Incommensuration Theorem state. The denial has the shape of the argument Chapter 3 recalled from René Descartes, "I think, therefore I am", which philosophers call the *Cogito*, Latin for "I think": to doubt that you are thinking, you must think, and in the same way the doubt enacts what it doubts.

The parts of comparison that the two rest on are these, and anyone who denies them uses each one in the denial. Sameness and difference go together inseparably, in each respect in which a comparison gives a verdict; content and context are distinct and inseparable; and none of this can be denied without being used. What is not available is to accept both the Root Tautology and the theorem and keep the position, because superdeterminism is nothing other than the demand for the pairing the theorem forbids, extended to the observer.

What We Now Know

We now know that there is no superdeterministic escape. Superdeterminism is out of reach by structure. It asks for symmetry and continuity perfected together, which is what determinism asks for, and the Incommensuration Theorem holds that this is the pairing no comparison can support. It then extends the demand to the observer, and in doing so asks for something further: that choice be causality, that the transcendent be the

omniscient, that the residual the theorem requires be empty. The further demand of superdeterminism is a claim across a boundary that the axioms hold impassable. The further demand fails as a category error before the further demand could fail as a hypothesis about the world.

The claims of this chapter are of two kinds, and each is stated at its kind. The central claim, that superdeterminism is excluded by structure, is a structural exclusion: it says what physics will not find, in advance of any experiment. The identification of superdeterminism as a category error across the triple of choice, change and causality is a claim of the same kind. The reading of quantum theory and general relativity as the two partial descriptions, and the reading of the circle in the usual position, are readings of known physics. Against superdeterminism the chapter holds exactly two things: regularity has a limit, and the residual always has content.

The two things matter for the book's own question, what cannot be, and why. Bell's matches stand as measured, and no hidden arrangement of the past stands behind them. The independence of the settings, which every experiment presupposes, stands with them. Whatever turns the dial, a person, a machine or a star, draws on the same residual that the particle draws on. And the proposal's three true points stand, each now telling against superdeterminism. Directed time requires that not everything is fixed. Quantum theory is partial because it is one of two survivors, not because there is a deterministic floor beneath it. And the circle is broken by resting the boundary at the speed of light on comparison, so that the boundary no longer needs a body of mathematics to hold it up, and it remains a real boundary.

The universe is a process, lawful and open at once, and it is open in exactly the measure that the theorem requires for anything to be compared at all. It is no block in which the label "earlier" is a convention. The arrow of time is real. Hard randomness is real. Subjectivity is real. The three are inseparable, and superdeterminism is the attempt to have the arrow of time while denying hard randomness; having the arrow without hard randomness is the one thing the structure of the real does not permit.

Summary

Superdeterminism is determinism extended to the experimenters' choice of settings, so that setting and particle share one causal history, and the choice of question is fixed along with the answer. It is excluded twice over, before any experiment. It asks for everything determinism asks for, perfect symmetry with perfect continuity, a pairing the Incommensuration Theorem forbids, and whatever cannot be had for the particles cannot be had for particles, apparatus and experimenter together. It then asks that choice be nothing but cause, that the transcendent be the omniscient, which is a category error across the triple of choice, change and causality, three things that are distinct, inseparable and non-interchangeable. The residual the theorem guarantees is no gap in our knowledge; a description with the residual closed is a description in which nothing happens. A variable correlated with everything is the whole of the past, renamed.

Of the three means of settling the question, Chapter 10 gave two: the Bell experiments narrow Bohm's account and the superdeterministic one and leave each possible by logic alone, and the theorem excludes both as accounts of a whole world, while a deterministic formalism stays available as a description, a map and not a territory. In this chapter I added the third means: Occam's razor settles the choice between superdeterminism and hard randomness before the structural argument is brought to bear. The Incommensuration Theorem says why the razor was right.

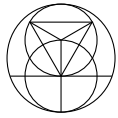
The proposal's true points stand, and each tells against it: time has a direction, which needs an open future; quantum theory is partial, as one of two survivors; the usual account of the light speed limit is circular, and the circle is broken by resting the limit on comparison. The argument depends only on the Root Tautology and the theorem, neither of which can be denied without being used, and it holds whatever any experiment finds, in any physics in which anything can be measured. The world is a process, lawful and open at once: the arrow of time, hard randomness and subjectivity are real together. That superdeterminism is no escape is a

structural exclusion, the third kind of claim the Preface names; the experiments the chapter reports are results of physics, given with their authors and years.

If no hidden plan links distant events, what does link anything to anything, and how far can a link reach?

PART FOUR

SIGNALS AND TIME



Chapter 12

The Shape of What Can Reach What

A Rule on a Stage

Why can we not send a message faster than light? Is it only that no one has yet found the trick?

The subject of this chapter is signals, and travel is a topic for another day. A *signal* is anything that passes from one event to another and lets a record at the second depend on an act at the first: a shout, a letter, a radio pulse, a lamp switched on and off. An *event* is a single happening at one place and one moment: a switch thrown in a laboratory, a click in a kitchen, a flash on the face of the Sun.

In ordinary talk, and in most of what is written for the public, the speed of light is presented as a speed limit: no signal and no piece of information, it is said, can go faster. The figure usually given is a little under three hundred thousand kilometers each second. Exactly, it is two hundred ninety-nine million, seven hundred ninety-two thousand, four hundred fifty-eight meters each second, and it is exact by agreement: the meter is now defined by the distance light goes in a set fraction of a second, so the number expresses the size of our unit of length. I quote the figure only to describe the speed, and nothing in this chapter depends on it. Einstein's relativity adds that, for a body with mass, durations lengthen and lengths shorten as that speed is approached.

The usual framing invites a picture. Space is a stage, objects are the actors on the stage, and the speed of light is a rule that limits how fast the actors may move about. The framing is not wrong, exactly. But it misleads, and it misleads in a way that raises a question that does not deserve the effort spent on it. Once the picture of the stage is in place, one naturally asks whether the rule might be broken. Can we find a loophole, a shortcut, a channel that the rule does not cover? The answer is no.

The reason is that the picture is wrong; how strict the rule might be has nothing to do with it. In place of a stage there is only the pattern of connection that the word "space" names, and the speed of light is a feature of that pattern, not a restriction laid upon it. There is no rule, and there is no stage. In the rest of this chapter I say that there is no rule and no stage, and I say so precisely enough that the question of signaling past the limit can be seen to answer itself.

Every Cause Is a Comparison

Start where the book started. Chapter 1 showed that interaction, comparison and relation are one act, named three ways; I call that identity the Root Tautology. So whatever is true of comparison as such is true of every physical event as such, because every physical event is an interaction, and a measurement is that same interaction described as one in which the comparison is recorded. Whatever limit applies to comparison as such applies to every physical event as such.

A cause is one particular kind of interaction. It is the interaction in which the pattern of relations between events is steady and one-way: causes come before their effects, like causes are followed by like effects, and the relation does not run the other way. Chapter 4 named the three modalities: the immanent, which is process, what is happening here and now; the omniscient, which is structure, what holds the same for everyone, laid out as a pattern; and the transcendent, which is possibility. In the sorting this book keeps throughout, causality belongs to the omniscient.

Causality is the repeatable, the structural, the shareable: what any comparer sees when the same kind of thing is followed by the same kind of thing in setting after setting. Causality is a pattern read from inside time, which is why it runs one way, and it is a pattern, which is why it is lawful. Chapter 6 showed that causality in this sense is what science establishes, and that it is a different thing from determinism, the claim that everything is fixed in full, every value for every setting, with no observer and no time. Here causality is the thing whose reach is in question.

Interaction is comparison, so causation is a form of comparison. Causation is a form of comparison, so the structure of comparison shapes the structure of causation. In particular, wherever comparison has an edge, causation has the same edge: the edge of what causation is, and not a restriction placed on causation from outside.

What Space Is

Take two events, and call them A and B. In the usual picture, A and B sit at positions in a space that is already there, and to ask whether A can affect B is to ask whether a signal can get from the one position to the other within the speed limit. Space comes first. Events are set into it. Causes are laid on top, and then trimmed by the rule.

I turn that order around. Two events are related exactly when one can affect the other: when a comparison can be made between them, when an interaction connects them. The whole set of such connections, the complete pattern of which events can affect which other events, is what the word "space" names. Space is the pattern of which events are causally connected with which, and it is no container that exists before the events it holds. A mathematician would call that pattern the *connectivity topology* of causal relations; a *topology* is the pattern of what hangs together with what, taken apart from any measure of distance.

In the terms of Chapter 4, space is the omniscient context: the pattern that holds across, the structure in which what is here can be compared with what is there. Time, by contrast, is the immanent context: the process in which the comparison is made, the change through which anything is here rather than there. The first of the three axioms holds that the immanent is more basic than the omniscient, than the transcendent, and than the two together, and so I hold time prior to space; the next chapter, on the direction of time, takes that up. What matters here is that space is a structure of relations, that the relations are causal, and that causal relations are comparisons.

This reading comes from inside physics: it is the content of general relativity, Einstein's theory of gravity as the shape of space and time, stated without the scaffolding of the container picture. In general relativity, spacetime (space and time taken together as one whole, as Hermann Minkowski described them in a paper published in 1909) is a changing structure, shaped by where matter and energy are, and it has no properties apart from those relations. The geometry is the physics. The curvature is the gravitation. Whatever in the theory could be called space is made of the relations that make it up, and of nothing else. And there are well-known results within the theory showing that the causal structure alone, the bare list of which events can reach which, already fixes the topology. David Malament proved one such result in 1977, and Stephen Hawking and George Ellis gave the standard account of this causal structure in their book of 1973; they are named here so that a reader who wants the proofs can find them. I take that seriously and trace it to its root. Space is made of causal relations, and causal relations are comparisons; so space is made of comparison, and the edges of space are the edges of what can be compared.

The Two Edges of Comparison

Comparison has limits built into its structure. They belong to what comparison is, and they are not limits of technology. There are two. I call the range within which comparisons can be made the *domain of comparison*. Its inner edge, the *microscopic boundary*, is marked by Planck's constant. Its outer edge, the *macroscopic boundary*, is marked by the speed of light.

Planck's constant is what physicists call the quantum of action: in plain words, the smallest package of process, of happening, that nature deals in. Below the scale it sets, for a given energy, no comparison can be made finer. The scale depends on the energy brought to the comparison; it is a floor under how fine a single comparison can be, and not a smallest piece of space or of time. At that scale what is found becomes grainy: an atom holds only certain sharp, separate energies, electric charge comes in whole units, and the numbers that label a quantum state are whole numbers. Grainy findings of this kind are what comparison looks like when comparison is taken to its lower bound. Chapter 18 develops this lower edge as the physical face of quantum indeterminacy, the fact that outcomes at that scale are not fixed in advance. In the terms of the theorem of Chapter 3, the lower edge is what one finds where symmetry is asked for absolutely: discontinuity, the jump, the click, the step. Chapter 2 called this the shadow of symmetry: symmetry and continuity cannot both be perfected, so a content that is symmetric, followed to its limit, comes in separate pieces that can be counted.

The speed of light marks the other edge. Beyond the reach it sets, for a given elapsed time, no comparison can be made at all. In one second, a comparison made here reaches out a little under three hundred thousand kilometers, and no farther; in one year, it reaches one light-year (the distance light covers in a year), and no farther. The cosmological horizon, the farthest surface in the sky from which anything has reached us, is the large-scale partner of the small-scale limit. It is the surface beyond which the structure of comparison does not extend from here. Comparison stops there because nothing from beyond the horizon has arrived, or can; dis-

tance is not the reason. Chapter 18 sets this horizon beside the other horizons of knowing, and Chapter 19 develops the two constants as the pair that bounds every domain.

The two edges define the domain. They are the edges of the arena (space and time together, the setting of every measurement) within which comparison, and therefore interaction, can happen, and neither edge is a fence put up around a space that is already there. Physics is the structure that exists between the two edges. Physics does not happen in a larger space that extends beyond them and is merely cut off at them. Planck's constant is the smallest package of process; the speed of light is the greatest rate of causal connection, the farthest reach of process. Together they bracket the domain, from the finest grain to the farthest reach. Both are properties of comparison itself, and neither is a property of the things being compared.

The Cone Comes First

Picture a single flash: a firework bursting at one place and one moment in a dark sky. Its light spreads out as a sphere that grows by about three hundred thousand kilometers every second. Now picture all the moments of that sphere at once: the small sphere of the first instant, the larger one a second later, the larger still a minute later. Physicists draw this with time running up the page and only one or two directions of space across it, and in such a drawing the growing sphere looks like a cone opening upward from the flash. Run the same picture the other way, and every place and moment from which light could have reached the flash makes a second cone, opening downward.

In relativity the *light cone* of an event is the set of all events that can be causally connected to it. The *forward cone* holds every event the given event can affect. The *backward cone* holds every event that can affect it. Events outside both cones are *spacelike separated*: they are so far apart in space, and so close in time, that no signal traveling at light speed could

pass between them. Spacelike separated events cannot affect each other, and no signal passes between them. Light from the Sun takes a little over eight minutes to reach the Earth, a measured figure that follows from the measured distance of the Sun, so a flash on the Sun and a click in your kitchen four minutes later are spacelike separated. The light cone could as well be called the cone of precedence, since it sorts, for each event, what can come before that event and what can come after it. The cone at each event is the shape of causal relatedness itself, and an act reaches its record only within its cone: whatever records an act lies in the forward cone of that act.

The usual account presents the cone as a result of the speed limit. On that account, two events are spacelike separated because a signal between them would have to exceed the limit, and it cannot. The cone is then a secondary structure: take the limit, apply it to every possible path, and the cone is what remains.

I reverse the order. The cone is primary, and it is not obtained from the limit. The cone is the structure of causal relatedness, which is to say the structure of which comparisons can be made from here.

The idea of causality in general relativity is a very perfected idea of causality. At each event there is a cone: everything inside it can be reached from that event, nothing outside it can, and the cone points one way. The cone at each event, unbroken and pointing one way, is what causation is. In the terms of Chapter 2 it is continuity with asymmetry, smooth and going somewhere. Causality can be described in two ways, as a pattern and as a reach, and I place each way differently. As a pattern, the same kind of thing followed by the same kind of thing in setting after setting, causality belongs to the omniscient, as the second section of this chapter said. As a reach, the connection from one event to every event it can affect, unbroken and one-way, causality has the shape of continuity with asymmetry, and the reach is what the cone of general relativity draws to perfection. Chapter 3 showed that symmetry and continuity cannot both be applied, at once and absolutely, to one comparison, and that neither can asymmetry and discontinuity; that is the Incommensuration Theorem, a struc-

tural fact about the nature of knowing. Continuity with asymmetry is one of the two pairings the theorem allows to be carried all the way. General relativity carries it all the way. To try to think of causality in any way that is not in some sense drawn from the general relativity idea of causality, the cone at each event, unbroken and one-way, is probably to have an inconsistent idea of causality. Because an idea of causality not drawn in some sense from the cone of general relativity is probably inconsistent, in this chapter I take the cone as my picture of what a cause can reach.

Two events are spacelike separated because there is no causal connection between them, and not because a signal between them would have to be too fast. They are not in the domain of each other's comparisons. No interaction joins the two events. No information passes between them, because the connection does not exist to carry it, and not because a rule forbids it.

An Exchange Rate Between Time and Space

Stated as relativity states it, the speed of light is an upper limit on how fast any influence can spread. At root, the speed of light is the macroscopic boundary of the domain of comparison itself, the outer edge of what any comparison can connect, and not the velocity of any one thing in particular. It is the greatest speed at which any signal can pass, and it is the conversion factor between the time side and the space side of the pattern of causal connection, the one ratio by which that connection is shared between time and space: the constant that mediates between the omniscient context, which is space, and the immanent context, which is time.

Astronomers already use the speed of light as such a conversion. A light-year is a distance named by a time: the distance light covers in one year. A light-second is a little under three hundred thousand kilometers. The speed of light turns an amount of time into an amount of space, the way a rate of exchange turns one currency into another. It tells you how the edge is shaped: how the large-scale limit of comparison is shared

between the one direction of time and the several directions of space. It describes the geometry of the domain. That light happens to travel at that speed is a fact about light: light is the kind of thing whose connections lie on the edge rather than inside it. Light, which travels in empty space at exactly that speed, is the perfected example of the edge; light is not the reason the edge is where it is.

Throw a ball forward at twenty kilometers an hour in a train moving at a hundred. For a rider, the ball moves at twenty; for someone standing beside the track, it moves at about a hundred and twenty. The speeds of ordinary things change with the motion of whoever compares them. The speed of light behaves otherwise: every comparer, however it moves, finds the same speed of light. Albert Einstein made that sameness a starting point of his special theory of relativity in 1905, the theory of how lengths and durations compare for comparers moving steadily relative to one another. Physicists call the sameness *invariance*: an invariant quantity is one that every comparer finds the same. So it is sometimes asked, as if it were a puzzle, why the speed of light should be a constant, the same for every observer, when the speed of everything else is not.

On this reading the answer is plain, and nothing is left to puzzle over. The speeds of things are ratios of content within the domain, and they differ from context to context, which is what one expects of content: the ball's speed is one finding for the rider and another for the bystander. The speed of light is the shape of the domain's edge, and no ratio within the domain; and the edge is the same from everywhere because it is not a thing in the domain at all. The edge is a property of comparison, and comparison is what every observer is doing. The invariance is no additional fact that a theory had to make room for. The invariance follows from the speed of light being an edge: an edge of comparison is the same for every comparer. This explanation is a reading of a known result: it gives the reason the result holds, and the theory that established the result stands as it is.

Beyond the Farthest Point

Asking whether one can signal faster than light is therefore a *category error*, a mistake about what kind of thing is being spoken of, and I want to say exactly what kind. The question treats the speed of light as a limit on motion through space. The speed of light is the shape of the domain within which motion is defined, and so it is no limit on motion.

A *constraint*, a limit laid on something, could in principle be otherwise: a rule that might be relaxed, a wall that might be climbed, a limit that better instruments might exceed. The edge of a domain is none of these. The edge is not a rule about what is in the domain; the edge is what makes the domain the domain it is.

An analogy helps. Picture someone walking on a globe, who asks for a point on its surface farther from where she stands than the *antipode*, the point directly opposite her on the far side. The answer is no. The reason is that on the surface of a sphere the antipode is the farthest point there is; the reason is not a rule forbidding anyone to go beyond the antipode. At the antipode, "beyond" is not defined. The surface does not extend past the antipode. The question is not hard. It is malformed, and the person who asks it has not yet understood what kind of thing a sphere is. A request to go north of the North Pole is the same: it defeats itself in its statement, because on the Earth north ends at the Pole.

The same globe gives a second picture. A traveler on open ground sees only so far, and the reason is the curve of the earth, which makes the traveler's horizon; no fence stands at the end of a flat meadow. My book *Incommensuration; The Theorem and the Three Axioms* puts the point in a phrase: the limit the theorem draws is "the curvature of the field itself, rather than a wall at the edge of an otherwise open field". In that phrase the field is a stretch of ground, like the meadow, and its curvature is the bend of that ground, the shape of the domain itself; the curvature of space and time that is gravitation is a different thing. The edge of the domain of comparison is like the curve of the ground, and not like a fence.

Signaling faster than light is the causal counterpart of going beyond the antipode. The structure of the domain rules it out, and no rule forbids it. There is no "there" to signal to. Events outside the cone of an event are outside the causal domain of that event, and they are not far-away places that better technology might reach. Events outside the cone are not places that are hard to get to. They are not places at all, in the sense that matters here, because a place is a position in the pattern of connection, and they have no position in the pattern of connection that runs from that event.

At this point the argument of this chapter meets the argument of Chapter 11. Chapter 11 took up superdeterminism, determinism extended to the experimenters' choice of settings, so that setting and particle share one causal history, as Sabine Hossenfelder and Tim Palmer defended it in 2020. The proposal to fix the experimenter's choice in that way was a claim of a different kind from determinism, and not only a stronger claim of the same kind: it crossed from one of the three modalities to another, from choice, which belongs to possibility, into causality, which belongs to pattern, and so it was a category error rather than a hypothesis. The proposal to signal past the edge has the same shape. Signaling past the edge is a different claim from signaling fast, and not a stronger one. Signaling past the edge is the claim to compare beyond comparison, and a comparison beyond the domain of comparison is not a comparison. A causal connection outside the causal structure is not a connection.

So signaling faster than light fails structurally, which is to say that it fails before there is anything for an experiment to test. It does not fail by experiment.

A Signal That Would Arrive Before It Was Sent

Chapter 10 showed that the linked outcomes of entangled particles carry no signal: the correlation cannot be used to send anything. The geometry of space and time shows the same result. Picture a long train car moving steadily along a straight track, with a lamp in the exact middle of the car.

The lamp flashes once. For a rider in the car, the light reaches the front wall and the back wall at the same moment, since both walls are the same distance from the lamp and the light goes at the same speed both ways. Now take a comparer standing beside the track as the car goes by. For her, too, the light goes at the same speed both ways; that is the invariance of the last section. But while the light is on its way, the back wall moves toward it and the front wall moves away from it. So, for her, the light reaches the back wall first. A third comparer, on a faster train overtaking the car, sees the car fall back past him, and for him the light reaches the front wall first.

The two events, the light reaching the back wall and the light reaching the front wall, are spacelike separated, and they have no fixed order: the same moment for the rider, back first for the bystander, front first for the overtaking passenger. Comparers in steady motion relative to one another disagree about which distant events happen at the same moment, and so two events outside each other's cones stand in no fixed order. For one comparer the first is earlier, and for another the second is.

The order of acts runs one way. In every comparison the range of what is possible comes first, then the act, then the record, and the record follows the act it records. Chapter 4 set out that order as the second axiom, and the next chapter shows that it holds in every single comparison. A signal joining two events outside each other's cones would, for some comparer, arrive before it was sent, since for some comparer the event of receiving comes first; and its record would come before the act it records. A signal sent at the back wall and received at the front wall would be received, for the overtaking passenger, before it was sent. The order of acts cannot be reversed, so such a signal is no signal.

The Edge Rests on Comparison Alone

The usual argument in physics against signaling faster than light uses the *no-communication theorem* of quantum theory. Giancarlo Ghirardi, Alberto Rimini and Tullio Weber gave a general form of it in 1980. Recall the entangled pairs of Chapter 10: two particles whose outcomes, measured far apart, agree in a pattern no separate instructions could give them. The theorem says that one cannot use an entangled pair to send information, because the statistics of a measurement on one member, its long run of yes and no answers, do not depend on which measurement is made on the other. The detector at each end gives a mix of answers that looks the same whatever question is put at the far end; the agreement shows only when the two lists of answers are set side by side. The *correlations*, the agreements between the two lists, are there, and they exceed what any account built from local causes allows, as John Bell showed in 1964 and Chapter 10 described; but they cannot be steered from a distance, and so they carry nothing.

That argument is correct. But it is not fundamental, and it is worth seeing why. The no-communication theorem is a result within one specific *formalism*, one particular mathematical framework for describing what is found. One can imagine proposing a different formalism: a change to quantum theory, a hidden-variables theory, in which hidden causes fix each outcome, or a regime in which the standard predictions fail; and one can ask whether the no-signaling result survives the change. Chapter 11 met exactly that strategy, where the hope was that a deeper deterministic theory might permit departures from quantum theory, and that the departures might open a channel. Whatever one thinks of the deeper theory, the hope rests on an assumption: that the edge is held in place by the formalism, so that changing the formalism might move it.

The argument of this chapter stands free of that assumption. It depends only on three things: that interaction is comparison, that comparison has structural limits, and that the speed of light marks the outer one. It does not call on the no-signaling result, on the formalism of entangle-

ment, or on the structure of the wavefunction, the smooth mathematical wave by which quantum theory describes a system before a finding. The edge was never a result of quantum theory, so the edge of the domain would remain if quantum theory were modified or replaced. No theory, quantum or otherwise, can extend comparison beyond the edge of what can be compared, for the same reason that no theory can make a comparison that is not a comparison. The physics we have is consistent with the edge; but the edge is not resting on that physics. The argument proceeds from the Root Tautology, that interaction, comparison and relation are one act, and the argument depends on no particular formalism.

An Arrow Opens No Door

A related proposal holds that signaling faster than light could be made consistent if the universe had a preferred direction of time, one that prevents anyone from building causal loops. A *causal loop* is a chain of causes that comes back around to its own start, so that an effect helps bring about its own cause. I agree that the universe has a preferred direction of time. Time is prior to space, and the next chapter shows that the direction of time is structural before it is a matter of heat and disorder. But agreeing that time has a direction does not make a signal faster than light possible.

A preferred direction of time confirms only that the causal structure is directed: that the cone has an arrow. The arrow tells you which half of the cone is the future and which the past. The arrow does not enlarge the cone; it does not stretch either half to take in what is spacelike separated. Directed time and a bounded reach of causal connection are features of the same structure, not alternatives between which one must choose, and directed time cannot be used to escape the bound on the reach of causal connection.

What Stays Exactly as It Was

A reader may ask what the argument leaves standing: the galaxies said to move away faster than light, and the entangled pairs of Chapter 10. The argument excludes one thing: signaling that passes beyond the edge of the causal domain. Everything within the cone is untouched.

Untouched, too, are the things sometimes called faster than light in the geometry of the universe at its largest scale: the expansion of the universe, and the moving away of distant galaxies at effective rates that, computed in a natural way, exceed the speed of light. The natural way is this: take the distance to a very distant galaxy as it stands today, and ask how fast that distance grows; for the most distant galaxies it grows by more than a light-year each year. These rates are features of the pattern of connection as a whole, of how the domain is shaped at its farthest reach. The moving away of distant galaxies is a growth of the arena between them, and not a signal across it. The expansion and the moving away make no signal and no causal connection; indeed they are the reason the cosmological horizon is where it is. Tamara Davis and Charles Lineweaver set out these rates, and the horizons they make, in 2004.

Untouched, last, are the correlations of entangled particles. Those correlations are real. Chapter 10 read them as the Incommensuration Theorem showing through: lawfulness and locality cannot both be had together absolutely. But a correlation is not a connection. Nothing passes, and the argument here concerns what can pass. The expansion, the moving away of galaxies and the entangled correlations have this in common: none of them compares anything across the edge.

The Silent Sky

The universe is very large and very old, and yet we have heard from no one else in it. That puzzle is called the Fermi paradox, after the physicist Enrico Fermi, who asked where everyone is. It has been proposed, as an

application of the superdeterministic program, that other civilizations may communicate faster than light, and that such communication would explain the silence of the sky: we hear nothing because we are listening on the wrong channel. The last chapter and this one answer two halves of that one proposal: the last chapter, the claim that the experimenters' choices were fixed from the start; this chapter, the claim that a channel faster than light could be open.

Every signal stays within the macroscopic boundary of comparison; no signal can cross it. So no civilization communicates faster than light, because no civilization can. The limit is not technological, and so it is not a limit that a more advanced technology might have surpassed. It is structural. It applies to any being, in any universe, that interacts with anything else, because interaction is comparison, and comparison has the edges that the theorem and the domain describe. A civilization that has understood physics better than we have has understood, among other things, what kind of thing the edge is; it has not found a way past the edge, because there is no "past" to be found.

The silence therefore needs some other explanation, and there is no shortage of candidates. I do not offer one. The argument of this chapter rules out the faster-than-light channel on which, the proposal says, we fail to listen. I would only add that the proposal has the shape this book has now met twice: the attempt to get outside a structure by means of that structure, to fix the chooser by causality, to connect the unconnected by connection. Each time, the attempt is defeated by its own formulation, and not by evidence.

What Kind of Impossible

The impossibility argued here is structural: the request to signal past the edge is the request to compare beyond comparison, and there is no content to the request once its terms are understood. It is a different kind from a logical impossibility: a formalism in which some quantity exceeds

the speed of light is a perfectly consistent object, and several have been written down (the hypothetical particles called tachyons, which would always move faster than light, are the best known). It is a different kind, too, from an empirical impossibility: no experiment is cited and none is needed, though every experiment is consistent with it.

The conclusion rests on two things: that space is the pattern of which events are causally connected with which, and, beneath that, that interaction, comparison and relation are one act. It rests on no particular physics; the physics already says what I read it as saying. The central claim of this chapter, that no signal can cross the macroscopic boundary of comparison, is a structural exclusion: it says what physics will not find, and it says so before looking. The claim that signaling faster than light is malformed, rather than merely impossible, is a claim of the same kind.

The other readings of this chapter are readings of known results: the cone as primary; the speed of light as the conversion factor between the time side and the space side of connection; and its invariance as the invariance of an edge. Each reading gives the structural reason the known result holds. The number for the speed of light expresses a unit, fixed by the definition of the meter; what I read is the edge.

What We Now Know

We now know why no one sends a message faster than light. To signal faster than light is to make a comparison beyond the edge of what can be compared. The concept of a signal faster than light is structurally incoherent, and not empirically difficult. The trouble is that the request is malformed, as the request to go beyond the antipode is malformed, and not that we lack the technology. One cannot interact with an event that is outside the domain of interaction. One cannot compare beyond the reach of comparison. So the request to send such a message names nothing, and no cleverer engineer will find the trick, because there is no trick to find.

The exclusion in this chapter is the first of the book's exclusions in space: what cannot be, here, is a signal past the outer edge of comparison. The chapters that follow meet the same structure in time.

Summary

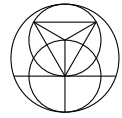
The speed of light is the macroscopic boundary of the domain of comparison, and not, at root, the speed of any one thing. Every cause is a comparison, because interaction, comparison and relation are one act, so the edges of comparison are the edges of causation. Comparison has two such edges: an inner one, marked by Planck's constant, below which, for a given energy, no comparison can be made finer; and an outer one, marked by the speed of light, beyond which, for a given elapsed time, no comparison can be made at all. Space is the structure of causal connection, the totality of relations that join events to one another, and since those relations are comparisons, the edges of space are the edges of what can be compared. Space is no container. The light cone of each event, the set of every event that can be causally connected to it, is the shape of the causal domain itself, and the speed of light is how that shape is shared between the one direction of time and the directions of space. The speed of light is the same for every comparer because it is an edge, and not a thing within the domain. The light cone is no speed limit applied to signals moving through a container.

A signal faster than light asks for a comparison beyond the edge of comparison, and for some comparer it would arrive before it was sent. No signal can cross the macroscopic boundary of comparison, for any being, in any universe, that interacts with anything else. The impossibility is structural: it is not a matter of logic, since consistent formalisms in which some quantity exceeds the speed of light have been written down; and it is not a matter of experiment, since none is needed, though every experiment agrees with it. The result depends on the structure of comparison, which is the structure of interaction, which is the structure of physics. It

does not depend on quantum theory, on special relativity as a set of postulates (its starting assumptions), or on any particular formalism; and a preferred direction of time does not enlarge the cone.

The edge is the edge of the domain, and it is not a wall. Beyond the edge there lies no forbidden territory. For the one who compares, "beyond" is not defined at the edge: the edge is a limit of comparison, and the limit the theorem draws is the curvature of the field itself, rather than a wall at the edge of an otherwise open field: the curve of the ground that makes a traveler's horizon, and not a fence at the end of a flat meadow. No comparison made from here reaches past the edge, no causal connection crosses it, and no finding comes through it. The scope of the result is signals across the edge. Three things stand exactly as they were, because none of them compares anything across the edge: everything within the cone; the expansion of the universe and the moving away of distant galaxies; and the correlations of entangled particles. The universe is not a cage with the light cone for its bars. The light cone is the shape of what it means for anything to happen at all.

The cone points one way: its future half and its past half are not alike. Why, then, does time itself run only one way, when the laws of physics seem to run as well backward as forward?



Chapter 13

Why Time Has a Direction

Eggs Break and Do Not Unbreak

The laws of physics run as well backward as forward; why does the world run only one way?

Everyone has seen an egg roll off the counter and break, and no broken egg ever gathers itself up and leaps back onto the counter whole. An ice cube dropped into warm water melts, and the water never gathers itself back into the cube. Heat flows from a hot cup into the cool air of the room, and never from the cool room into the cup. Memory holds the past and not the future: you remember yesterday, and you do not remember tomorrow. Causes come before their effects. That time has a direction is among the most immediately felt features of anything we would call real, and among the most universal.

Now picture a film of a rubber ball dropped onto a floor. Each bounce is a little lower than the one before, until the ball lies still. Run the film backward, and you see a ball lying still on the floor that begins to hop, and hops a little higher each time, until it rises into the hand that dropped it. Anyone watching knows at once that the film is running backward. Yet each arc of the ball in the backward film is an arc a ball can fly, and each push the floor gives it is a push the laws of motion allow, provided that the molecules of the floor, at each bounce, all move together in just the right way. Nothing in the backward film breaks a law.

The puzzle is this: anyone can tell that the film is running backward, yet nothing in the backward film breaks a law. Nearly all of the basic equations of physics are exactly the same run backward as run forward; the one exception, in the weak interaction, departs from that sameness only slightly. Newton's laws of motion, Maxwell's equations of electricity and magnetism, Schrödinger's equation for the waves of quantum theory, and Einstein's equations of gravity each allow motions that run equally well in either direction. If a process is possible under one of these laws, the same process run backward is possible under the same law. Nothing in the form of the equations tells the past from the future. Physicists call the one-way character of the world the *arrow of time*: the direction from past to future along which every breaking egg, every melting ice cube and every memory points. In this chapter I find where the arrow of time comes from, and I find the source of the arrow in every single act of comparison.

The Usual Answer, and Two Worries

The usual answer comes from the science of heat, thermodynamics, and its central quantity, entropy. The *entropy* of a thing is a count of how many ways its small parts could be arranged while the thing looks just the same from outside. An egg spread across the floor, or heat spread through a whole room, can be arranged in vastly more ways than an egg whole in its shell or heat gathered in one cup, and so each of the spread-out ones has the higher entropy. The *second law of thermodynamics* says that the entropy of a closed system of many parts does not decrease. On the usual account, the arrow of time is the direction in which entropy increases. I call that the *entropic arrow*. The second law supplies the one-way character that the equations lack.

Two features of that account trouble careful thinkers. The first feature that troubles careful thinkers is that the second law of thermodynamics is a law of chances. It describes what very large collections of parts will do

with overwhelming likelihood; it is not a strict rule that governs each event. The second feature goes deeper. The law of chances, with no added assumption about how things began, runs both ways. From any given moment, the same counting of arrangements that predicts rising entropy toward the future predicts rising entropy toward the past as well. So the account needs one more assumption: that the universe began in a state of extraordinarily low entropy. That assumption is called the *Past Hypothesis*. The philosopher David Albert argued in 2000 that it should be admitted as a basic law of physics in its own right, because nothing else in physics supplies it.

So the entropic arrow rests on a fact about how the universe began, a fact that could have been otherwise. Resting the arrow of time on such a fact has troubled philosophers and physicists alike, and rightly: why should something so basic to experience depend on an accident of the cosmos? The philosopher Huw Price, in 1996, pressed the same trouble from the opposite direction, by asking whether the direction of time belongs to the world at all. He argued that our sense of a direction to time may be a quirk of our own position inside time, and that the fair view is the view from nowhen: a vantage point outside time altogether, from which no direction of time is favored, like the fixed point outside the world from which Archimedes said he could move the Earth.

I accept the force of both worries, and both are answered by something more basic than thermodynamics: the structure of comparison. The arrow of time rests on the structure of comparison itself, and not on entropy.

Two Ways to Run Backward

Recall the words Chapter 2 built from the parts of every comparison. Every comparison carries six parts it cannot do without: sameness and difference, content and context, the subjective and the objective. Content is what a comparison finds; context is the setting that makes one finding distinct from another, such as two places or two moments. Applying same-

ness and difference to content and context gives four compounds. *Continuity* is sameness of content within sameness of context: hold the setting fixed, and nothing jumps. *Discontinuity* is difference of content within sameness of context: a jump, a click. *Symmetry* is sameness of content across difference of context: change the setting, and the content stays the same. *Asymmetry* is difference of content across difference of context: change the setting, and the content changes with it.

Chapter 3 found the one limit every later chapter uses. The Incommensuration Theorem holds that two pairs of these four cannot be applied absolutely and together: symmetry with continuity, and asymmetry with discontinuity. Two pairs survive: continuity with asymmetry, and symmetry with discontinuity. The words "absolutely and together" are part of the theorem, and they are never dropped. Within any finite domain, symmetry and continuity are both present, each to some finite degree. Symmetry and continuity cannot both be perfected at once. The Incommensuration Theorem is a structural fact about the nature of knowing.

Time-reversal symmetry is a symmetry in exactly the sense the four compounds gave: sameness of content across difference of context. The content is the laws of physics. The change of setting is the reversal of time: every time reading is replaced by the same reading with its sign flipped, so that ten seconds after becomes ten seconds before. To say that the laws are symmetric under time reversal is to say that their content stays the same across that change of setting. The symmetry of the laws under time reversal is the definition of symmetry itself, applied to one particular change of setting, and not something that only resembles the definition.

What would a perfect time-reversal symmetry require? It would have to hold for every process, at every scale, in every setting, without exception. Here two claims separate, and the usual treatment of time reversal treats them as one claim.

The first claim is that the laws are symmetric: if a process is lawful, its reverse is lawful too. I call that *lawful reversal*. The second claim is that the reversal is faithful: the reversed process passes back through exactly the

states the original passed through, and arrives at exactly the condition the original started from. I call that *faithful reversal*. The *state* of a system is the full description of it at one moment: where each part is, and how each part is moving. The first claim is about the form of the laws. The second claim is about how complete the description of the state is, the state being what the laws act upon.

Think of the film of the ball again. A lawful reversal is any backward motion the laws allow. A faithful reversal is the film itself run backward, frame by frame, every molecule of the floor and of the air retracing its exact path, until the ball is back in the hand that dropped it, exactly as it was.

Symmetry of the laws guarantees lawful reversal. It does not guarantee faithful reversal. For a reversal to reproduce the original, every aspect of the state must be captured in the description that is reversed. Physicists call each independent way a system can vary a *degree of freedom*: the position of one particle along one direction is one degree of freedom, and its speed along that direction is another. If even one degree of freedom is left unspecified, the reversed motion is free to carry that degree of freedom somewhere other than where it came from, and the reversed process will not retrace the original. Faithful reversal needs a description of the state with no gaps and no ambiguities, one that leaves no part of the system unspecified and free to go its own way.

A full specification of the state, with no gaps and no ambiguities, is exactly continuity, in the sense the theorem uses. It is sameness of content (the state, fully defined) within sameness of context, at every point and every scale, without a break. So a perfect time reversal, one that is faithful as well as lawful, needs two things at once: the symmetry of the laws under reversal, and the continuity of the specification of the state. What a perfect time reversal needs is symmetry with continuity, applied absolutely and together. It is the pair the theorem forbids.

Perfect time reversal is therefore impossible. Perfect symmetry of time, under which every process could be run backward faithfully, through exactly its own states to exactly its own starting condition, cannot be. The impossibility rests on the Incommensuration Theorem and on the bridge from lawful to faithful reversal alone; it uses no entropy, no starting condition, no Past Hypothesis, and no particular theory of physics.

I have now given the whole of the argument that perfect time reversal is impossible. It turns on one step, the bridge from lawful reversal to faithful reversal, and the bridge holds because a faithful reversal needs a fully specified state, and a fully specified state is continuity. Every other step is the Incommensuration Theorem forbidding symmetry and continuity to be perfected together.

Why the Equations Hide the Gap

The mathematics of physics assumes that every state is fully specified, and so physicists usually do not notice the step from lawful reversal to faithful reversal. Classical mechanics describes a system by one complete list of the position and the motion of every part; physicists call such a list a point in phase space. Quantum mechanics describes a system by one complete quantum state, which physicists call a state vector. In both, the description is complete by construction, and once the description is taken to be complete, lawful reversal and faithful reversal look like one and the same thing.

The Incommensuration Theorem shows that the complete description is an idealization: it assumes perfect continuity, and continuity cannot be perfected at the same time as the symmetry of the laws. The mathematics assumes perfect continuity together with perfect symmetry of the laws, and the theorem shows that the two cannot be achieved together. The mathematics remains an excellent tool. The theorem sets a limit on what the tool can be taken to say about the real world.

In 1876 Josef Loschmidt raised an objection that is an early sign of the same gap, the gap between lawful reversal and faithful reversal. Picture a box of gas whose molecules all began crowded in one corner and have since spread through the whole box, so that the entropy of the gas has risen. Now imagine stopping every molecule at once and sending it off again at exactly its own speed in exactly the opposite direction. If the laws of mechanics run backward as well as forward, every molecule retraces its path, and the gas gathers itself back into the corner; its entropy falls. So for every path along which entropy rises, there is a path, obtained by reversing every velocity, along which entropy falls. Loschmidt's objection is an argument about the equations, and not something anyone has observed.

Ludwig Boltzmann answered that the reversed condition is overwhelmingly improbable. His answer concedes, without saying so, that what keeps the gas from gathering itself back is the need to specify the state with every velocity exactly reversed, and not the symmetry of the laws. The theorem shows that the obstacle Boltzmann conceded is required by the structure of comparison. The specification of the state cannot be perfected, because perfecting it alongside the symmetry of the laws would perfect continuity and symmetry together, the pair the theorem forbids. The reversed path exists in the mathematics, as a solution of the laws; a gas in the faithfully reversed state never occurs in the world.

The philosopher John Earman argued in 2002 that physics writing habitually runs together two things: what time-reversal symmetry is, and what it implies for the direction of time. The distinction between lawful reversal and faithful reversal is the clarification Earman asked for, grounded in the structure of comparison, and his point is named here because it shows that physics writing itself runs together the claim of lawful reversal and the claim of faithful reversal.

An Arrow Beneath the Arrow of Heat

The entropic arrow is a result about the chances of large systems, and it needs four things. First, the entropic arrow needs thermodynamics. Second, the entropic arrow needs the difference between the full detail of every particle, which physicists call the *microstate*, and the overall look of the whole, such as the temperature and the pressure of the whole, which physicists call the *macrostate*. The rule that lumps together every detailed arrangement sharing one overall look is called *coarse-graining*. Third, the entropic arrow needs very many particles. Fourth, the entropic arrow needs the Past Hypothesis.

The arrow found from the gap between lawful and faithful reversal needs fewer things. I call it the *structural arrow*. It follows from the parts of comparison alone, and it needs none of the four. It needs no large number of particles, since it is present in a single comparison; it needs no coarse-graining and no concept of entropy, since it is present in the act of comparison itself; and it needs no beginning of low entropy, and no starting condition of the universe at all, since it belongs to the order in which acts of comparison follow one another and not to the statistics of states.

The structural arrow applies to every system in which comparison happens. By the Root Tautology of Chapter 1 (interaction, comparison and relation are one act, named three ways), that is every system in which interaction happens. An interaction is a comparison, and a comparison has the structure that the Incommensuration Theorem limits. So the structural arrow is already present before the first entropy is calculated, and the structural arrow is present in every single interaction, not only in large collections.

Because the structural arrow is already present before the first entropy is calculated, and present in every single interaction, I say that the structural arrow is more fundamental than the entropic arrow. Thermodynamics needs the structural arrow first; the structural arrow is not taken from

thermodynamics. Every process of heat is built from interactions, and every interaction already carries the one-way character the theorem requires.

The structural arrow lies beneath the entropic arrow, and the two are joined in every act of recording. A record is itself a change that cannot be undone, and making a record raises the entropy of the instrument that holds it. The stopwatch of Chapter 1, whose click is never unclicked, gains a little entropy in the very act of recording the click. So the two arrows are different things that point the same way, and each record joins them. The entropic arrow is how the structural arrow shows itself at our own scale, in systems of many parts, where reasoning by chances applies. The structural arrow is the ground: the reason that any process has a direction at all. The structural arrow supplies a one-way character that the structure of comparison requires and that nothing can remove, and the entropic arrow magnifies that one-way character at the scale of eggs and engines. Wherever the entropic arrow is found, it is found pointing the same way as the structural arrow.

On this reading the Past Hypothesis is a statement about the starting conditions of the universe: it fixes the particular course that the direction set by the structural arrow takes in the history of this particular universe. The Past Hypothesis is not the source of the arrow. The argument shows that the direction of time never depended on that starting condition; why the starting condition was what it was is a topic for another day.

The claim of the structural arrow is that reversal cannot be faithful: the past cannot be recovered from the future by running the laws backward, because the state from which the laws would run is never completely specified. The claim of the structural arrow is a different kind of claim from the entropic arrow's claim that entropy keeps growing. The direction of time is the direction in which the unspecified remainder of the state, the part no description fixes, is left to go its own way, rather than recovered as faithful reversal would have to recover it.

Chapter 8 met that remainder under other names. There, hard randomness was an outcome that no prior specification fixes, however complete, for any comparer whatsoever, and the bound triple was subjectivity, directed time and hard randomness, each requiring the next. The remainder that faithful reversal would have to pick up is the same irreducible remainder that shows itself as hard randomness, and as the first-person stance that cannot be reduced to anything else. The arrow of time is that same irreducible remainder, as it shows itself in time rather than in outcomes or in the first-person stance.

The Order of Every Comparison

A second road reaches the same arrow, from the order in which every comparison proceeds. Chapter 4 named the three modalities: the immanent (process, what is happening here and now), the omniscient (structure, what holds the same for everyone, laid out as a pattern), and the transcendent (possibility, what makes things possible and new). In this book's alignment, symmetry belongs to the omniscient: the structural, the repeatable, the shareable, the same content in every setting. Continuity belongs to the transcendent: the unbroken, the fully specified, the same content held without a gap through a single setting. Interaction belongs to the immanent: the act happening here and now, in which symmetry and continuity are each present to some degree and neither is perfected. In *An Immanent Metaphysics* (aphorism 1.85-4), causality is omniscient, change is immanent, and choice is transcendent. Time, the measure of change, belongs to the immanent.

The axioms of the metaphysics already carry a direction in their own form. The second axiom (aphorism 1.3-1) reads: "A class of the transcendent will precede an instance of the immanent. A class of the immanent will precede an instance of the omniscient. A class of the omniscient will precede an instance of the transcendent". A *class* here is a kind, and an *instance* is one case of a kind. In plain words: a range of possibilities comes

before any act; an act comes before the record of it, laid out as structure; and structure laid out comes before the next range of possibilities. The cycle runs from the omniscient to the transcendent to the immanent, and back to the omniscient. It is a cycle of "comes before", not a cycle of "is the same as", and its order runs one way and cannot be reversed.

The third axiom reads: "The (classes/instances) of the immanent, omniscient, and transcendent are distinct, inseparable, and non-interchangeable". Not being interchangeable is asymmetry at the root: swap any two of the three, and what is found changes, which is difference of content across difference of context. And "comes before", the order of the second axiom, gives that asymmetry its one direction. The structural arrow of time is this order of precedence, applied to physics: the direction of time is the direction in which comparison proceeds. The arrow and the order of the axioms are one and the same structure, not two things that resemble each other, and the sameness shows in the simplest fact about any record. In every comparison the range of the possible comes first, then the act, then the record, and the record follows the act it records. The click of the stopwatch comes after the press of the button, and the record of the click comes after the click.

Two roads therefore reach one arrow: faithful reversal is impossible, and the cycle of comparison has a direction.

The philosopher Tim Maudlin argued in 2007 that the passage of time is primitive: a bare given, with nothing beneath it, not reducible to entropy or to the history of the cosmos, and a basic feature of the world that physics should make room for rather than explain away. I agree with his conclusion and differ on its standing. Maudlin says the arrow cannot be grounded in anything more basic. I say the arrow is grounded, in the parts of comparison, and so it is structural rather than primitive. The arrow is what it means for anything to interact with anything else. The difference matters: a brute fact could have been otherwise, and a requirement of structure could not.

Time as Happening, and Not Only as a Line

Chapter 1 called the outcome of one act of comparison a finding, laid down as a record. More exactly, a finding is one outcome selected from a set of separate possible outcomes and laid down as a record. I call the order in which a comparer's records are made *event time*, the time of findings. The findings make that order; they do not happen inside a time given before them. Event time is not a line along which findings are placed. Event time is the order that the findings themselves make, one record after another. Its order is the order of the second axiom: a precedence, and not a clock reading, and it cannot be reversed.

Physics uses the one word "time" for three different things, and Chapter 14 walks through all three. *Geometric time* is time as a direction in the shape of space and time, the time that clocks measure along their paths; it is the dimension, and it exists. Event time is the order of findings; it is the process, and it is real. *Evolution time* is the smooth dial along which the equation of quantum theory carries a state forward. The three are distinct, inseparable, and not interchangeable: time is itself a triple. The direction belongs to one of the three: event time.

The *block universe* is the view that past, present and future are equally real, that time is a fixed dimension rather than a directed process, and that the felt direction of time belongs to us and not to the world. Picture the whole history of the world as one solid block, with every moment laid out in it at once, and our attention as a light that slides along a line already fully drawn. The view is the natural companion of determinism. If the past fixes the future uniquely, then the future exists in the same sense the past does, and the only thing that moves is our attention. The block universe takes geometric time for the whole of time, and drops event time.

Chapter 5 found that determinism, the claim that everything is fully and exactly fixed, needs perfect symmetry and perfect continuity together, and so cannot be had. Perfect time reversal cannot be had, for the same reason. Together the two results undercut the block universe from two sides. Determinism is excluded, so the future does not exist in the sense

the past does: there is a real residual, a part of the future that the past leaves unfixed, in the world itself and not only in what we know. Faithful reversal is excluded, so the two directions of time are not equivalent: the difference between forward and backward is built into the structure of comparison, and it is not a convention.

Those two results are why I place time among the immanent and not among the omniscient. The omniscient is the domain of the structural, the shareable, the repeatable, the same in every setting. A dimension is an omniscient notion: a line along which every point has the same standing as every other. Change is immanent. Change happens here, now, in the interaction, and a change run backward is not the same as the change run forward. Time is the measure of change, and the arrow is the difference found when a change is compared with the same change run backward.

Price's view from nowhen asks us to take the omniscient stance on time, which is immanent. The stance is available as an abstraction. It is not available as a place from which a comparison can actually be made, because comparison is made only in the present, and the present is not one point on a line; the present is the site of the interaction itself.

One more meaning of time belongs beside the three times of physics, geometric time, event time and evolution time, so that it is not taken for any of them: time as it is felt by the one who lives it. A separate and older account concerns felt time. The proportional theory of Paul Janet, from 1877, holds that the felt length of a stretch of time is in inverse proportion to the length of life already lived: a stretch is felt as the fraction it makes of the life lived so far. A year is a tenth of the life of a ten-year-old and a fiftieth of the life of a fifty-year-old, so on Janet's theory a year feels one fifth as long at fifty as it did at ten. Equal stretches of felt time then match stretches of clock time that grow in proportion to the age at which each stretch begins: the ten years from ten to twenty feel as long as the twenty years from twenty to forty, and those feel as long as the forty years from forty to eighty. A quantity that counts doublings, rather than years, is said to grow as the logarithm, so on Janet's theory felt time grows as the logarithm of age. The theory accounts for the common observation that time

seems to speed up with age: each year feels shorter than the last, because it is a smaller fraction of the time already lived. Janet's theory is a claim of psychology, and not a result of the framework of this book; felt time is how long a stretch of life seems to the one living it, and the direction of time this chapter finds belongs to event time, the order of findings.

What Can Be Run Backward, and What Cannot

The laws of mechanics, of electricity and magnetism, and of the smooth change of quantum states read the same with time reversed, apart from the small breaking of time-reversal symmetry observed in the *weak interaction* (the force behind some kinds of radioactive decay). I take up that small breaking in a later section of this chapter. That sameness is often taken to mean that any process could, in principle, be run backward. The inference runs together the two claims this chapter has separated, lawful reversal and faithful reversal. A second question follows: can physics reverse the order of findings? Going back to an earlier event, as the event it was, would need the order of findings to run backward, and the order of findings cannot run backward. So it matters what physics can reverse, and whether what it can reverse is the order of findings.

Quantum theory carries a state forward by a smooth rule, and the rule has an exact undo: for every stretch of smooth change there is a rule that runs that stretch back. The undo acts on the state, which is given in evolution time. It does not act on a finding, which is given in event time. A finding is one discrete outcome, and its record is not undone, while the smooth rule is continuous and has an undo; so a finding can never serve as the dial of the smooth rule, and event time can never stand in for evolution time. Everything I argue in the rest of this chapter turns on the distinction between a state, which the undo of the smooth rule reverses, and a finding, which no undo reverses.

A measurement is the place where quantum theory seems to run only one way, so the account of what happens at a measurement bears directly on what can be reversed. The standard formulation of quantum mechanics describes a measurement as a sudden change of state, from a *superposition* (a description that holds several possible outcomes at once, each with its chance) to the one outcome found, and calls that change the collapse of the wavefunction. In my account, the update of the state at a measurement is not a physical process, and it is not an extra law of motion. It is an update of what is known, in one fixed sense: what changes is the laid-out record, which is shared and the same for everyone who has the outcome; it is not anyone's private degree of belief. The word "collapse" is kept only as the name of the act in which one outcome is found.

A measurement is irreversible because of the order of the second axiom, and not because of any motion that the equations describe. Each step of that order, from the range of possibilities to the act to the record, goes from many to one. The superposition before the finding is not destroyed by a physical process; it is superseded by the new content that the comparison has produced. In my account of quantum theory the time of findings is also called measurement time: the time of separate, irreversible, chancy findings. Measurement time is event time under a second name, and not a fourth time. Evolution time is the time of smooth, reversible change, at the level of the mathematics. So one thing in quantum mechanics is reversible, in the sense that an undo exists in the mathematics: the smooth change of a state. A finding is not reversible in that sense.

Physics also has a laboratory reversal, and a picture shows how it works. Picture runners on a wide track, all starting together at a line when a whistle blows, each running at a slightly different speed. After a while they are spread out along the track. The whistle blows again, and every runner turns around and runs back at the same speed as before. The fastest runner, now farthest out, has the farthest to come back, and comes back fastest; after the same time again, all of them cross the starting line together.

The *spin echo* is the laboratory version of the runners, and it is the standard laboratory example of a process that is run backward. The nuclei of atoms in a liquid each carry a spin, like a tiny top, and in a magnetic field each spin turns about the direction of the field. Nuclei in slightly different fields turn at slightly different rates, so a set of spins lined up together drifts out of step, and the signal the spins give together dies away. A second short pulse of radio waves, applied some time after the first, reverses the direction in which the spins have drifted apart, and after a further equal time the spins come back into step, and the signal returns: the echo. Erwin Hahn found the echo in 1950, with pulses that each turned the spins through a right angle. In 1954 Herman Carr and Edward Purcell gave the form of the echo that uses a single refocusing pulse, which turns the spins through a half turn. In 1971 Won-Kyu Rhim, Alexander Pines and John Waugh extended the reversal to the interactions of the nuclei with one another. That history shows how far the reversal has been carried: first the drift of each nucleus on its own, then the interactions of the nuclei with one another; and even carried that far, the echo undoes the falling out of step only in part.

The echo comes back because the drift was held in each nucleus's own angle, one angle for each nucleus, and one pulse could reverse every one of those angles. I call what a motion leaves in one particle, such as the angle by which one nucleus has drifted, a *trace*, and I keep the word *record* for a finding that is not undone. Each reversal acts on the traces that the motion has made among the particles a pulse can reach. No reversal gathers back the traces that the light of a warm body, or the collisions of the molecules in a glass of water, carry off to countless particles. The echo undoes the falling out of step only in part.

Two things are therefore true of the spin echo. First, the spin echo is a lawful reversal of the drift of the spins, possible because each trace of the drift is held in one particle, and it is partial, as the gap between lawful and faithful reversal requires. Second, the spin echo undoes no finding. The finding that the signal had faded was made before the pulse, and its record stays in the laboratory; the returning signal is a new finding, made

after the pulse. In event time the experiment runs forward throughout: a finding that the signal has faded, then a pulse, then a finding of the echo. What runs backward is the evolution of the spins, in evolution time, and not the order of findings.

So we can say exactly what a reversal can reach. A reversal of evolution, whether by an undo in the mathematics or by a pulse in a laboratory, acts on a state, given in evolution time, or on the physical traces of a motion among particles. It does not remove a finding from the order of findings, and it does not place a comparer back before a finding the comparer has already made. Three steps show that a reversal neither removes a finding nor places a comparer back before a finding. First, the undo of the smooth rule acts on the state, and a finding is not a state. Second, a laboratory reversal is a process, and every process a comparer brings about is followed by further findings, which take their places after the earlier ones in event time: the fading, the pulse, the echo. Third, the order of findings is the order of the second axiom, which cannot be reversed.

The order of findings is also the order of causes. Chapter 6 defined a cause as a consistency that a finite comparer finds between two things in the world, one following the other, again and again. The order in which one thing follows the other is the order of findings. So there is no backward causation: a comparer at our own scale, seen from its own first-person stance, cannot reach information from the future of an event. The direction of time, which makes causality what it is, also limits it. That limit in time is the structural arrow, present in every comparison. A later choice of what to measure is a new stance, which follows any finding already made and cannot reach back to that finding. Sending a signal into the past would require running the cycle of the second axiom backward, and the axiom's order runs one way only.

An Anchor in the Free Will Theorem

The two roads so far stand on the structure of comparison. They meet a result from the mathematics of physics itself: the free will theorem of Chapter 9. John Conway and Simon Kochen published that theorem in 2006, and its strong form in 2009. It holds that if the choices of experimenters about which measurements to make are not fixed by prior information, then the responses of the particles measured are not fixed by prior information either. Their theorem has the form of an "if": it does not establish that experimenters are free. It establishes that freedom, wherever it is present in an experiment, cannot be confined to the experimenter's side of the experiment. Either the whole is fixed, or the particle is as unfixed as the experimenter. The middle ground, a free choice above and a fixed response below, is what the theorem excludes.

On that theorem I build an argument called the *Epistemic Sandwich*, a sandwich of knowing. It runs across three layers of physical description, each at a different scale. Our own scale is the filling; the very large and the very small are the two slices, and everything we know of either slice reaches us through the filling. The three layers are distinct, inseparable, and not interchangeable, and their scales have the natures of the immanent (our own scale), the omniscient (the very large), and the transcendent (the very small).

Layer one is the *mesoscopic* scale: the scale of observers and their instruments. In this book "mesoscopic" names the observer's own scale, and not, as in the physics of materials, the scale between single atoms and bulk matter. Layer one is described by Newton's mechanics, by the theory of electricity and magnetism, and by thermodynamics. Layer two is the *macroscopic* scale, the very large, described by general relativity. Layer three is the *microscopic* scale, the very small, described by quantum mechanics. The argument needs only the three scales and the theories that describe them. The Appendix gives the layers in their technical form: the terms in which each layer is described, and the key concepts of each.

At layer one, the scale of laboratories and instruments, knowledge depends on the entropic arrow. To measure, to record, to tell an ordered state from a disordered one, requires that there be records, and records require an arrow. Without the entropic arrow at our own scale, nothing could be told apart from anything at all. And all of our knowledge of the very large, described by general relativity, and of the very small, described by quantum mechanics, reaches us through this knowing at our own scale. Knowing at our own scale is the only channel connecting the mathematics of general relativity and of quantum mechanics to the real world.

Now apply the free will theorem. If experimenters at layer one make choices that are not fixed in advance, which is what it means for their measurements to be measurements rather than replays, then the theorem carries that unfixedness down to layer three. I read an unfixed response at layer three as a state that is not fully specified, which is to say a state whose reversal cannot be faithful, by the same bridge between lawful and faithful reversal that carries the first road. That the reversal of an unfixed response cannot be faithful is a one-way character in time at the quantum scale. The one-way character at the quantum scale is present in single events, before any collection of events, and owes nothing to entropy. And a one-way character of that kind at the smallest scale sets the direction for the entropic arrow at the layers above it. The reason is the one this chapter gave earlier. The entropic arrow is a law of chances, and a law of chances, left to itself, runs both ways; every process of heat is built from single interactions; so the one-way character of the single events beneath is what gives the entropic arrow above them its one direction.

So the conclusion has the form of a choice. Either an arrow beneath the entropic arrow is real at all three layers, or the entropic arrow, on which all knowledge depends, is itself absent, and we know nothing at all. The second half of the choice follows from what I said of layer one, where knowledge is bound to the entropic arrow. Take away the arrow beneath, and nothing gives the entropic arrow a direction. Without the entropic arrow there are no records, and without records nothing can be told apart from anything, so nothing is known.

This anchor has the form of an "if", and that form is where its force lies. The anchor shows that denying the arrow requires denying the very knowing in which the denial would have to be made; it does not assert the arrow by decree. The choice the anchor offers is between knowledge and its absence, and not between two theories of physics. The "if" is the free will theorem's own. The two roads to the arrow stand on three things: the Incommensuration Theorem, which is a structural fact about the nature of knowing; the bridge between lawful and faithful reversal; and the order of the second axiom. The Appendix gives, for the specialist, which step of the anchor is the result of Conway and Kochen and which step is my reading. The view that every experimenter's choice was fixed in advance is answered on other grounds: Chapter 11 showed that this view asks for determinism, extended to the experimenters' choices, and determinism cannot be had.

What Physics Already Shows

The structural arrow sits comfortably with the time symmetry of the laws, and the reason is the words "absolutely and together" in the theorem. Within any finite domain, symmetry and continuity are both present, each to a finite degree. The laws show time-reversal symmetry to an extraordinary degree, and the argument agrees. What the argument says is that the symmetry cannot be perfect. There is a residual one-way character in time: small, difficult to measure, and required by structure.

Known physics already shows such a residual, and I read the known physics as an instance of the argument. The weak interaction, the force behind some kinds of radioactive decay, breaks a symmetry physicists call C P. The C, charge conjugation, swaps every particle for its antiparticle, its twin with the opposite charge; the P, parity, reflects space as in a mirror. In 1964 Christenson, Cronin, Fitch and Turlay found the decays of particles called neutral kaons breaking the combined symmetry of the two. A general result, the *CPT theorem*, says that swapping particles for antiparticles,

mirroring space and reversing time, all three together, is an exact symmetry of every local quantum field theory, the kind of theory physicists use for particles and forces. So a break of C P entails a break of time-reversal symmetry. There is therefore already a time-reversal asymmetry, confirmed by measurement, in the accepted description of particle physics, and direct observation of it has since been reported as well: by the CPLEAR collaboration in 1998, again with neutral kaons, and by the BaBar collaboration in 2012, with particles called B mesons.

I read these findings as the kind of thing the argument expects: a place where the impossibility of perfecting the symmetry of time has become visible to instruments, and not a chance feature of the mathematics of one force. What the argument requires is that some asymmetry exist; the size of this one, the interaction that shows it and its mechanism are found by measurement. The Appendix states for the specialist exactly what the reading of these findings claims and what it leaves to measurement.

The CPT theorem itself fits the argument. The argument rules out perfect time-reversal symmetry taken by itself. It says nothing against a combined symmetry that mixes time reversal with the other two flips. Nature keeps the combined symmetry and breaks the symmetry of time alone, and that pattern is what the argument allows: the direction of time cannot be made symmetric by itself, though it can take part in a combination where the other flips make up for its one-way character. The exactness of the CPT theorem, too, is exactness inside a mathematics, and that mathematics assumes the fully specified state met earlier in this chapter. What is exact in the model is not thereby perfected in the real. The CPT theorem thus meets the same limit that time-reversal symmetry meets, the limit that symmetry and continuity cannot be perfected together, approached here from the side of an exact mathematical theorem.

One further consequence belongs here, and Chapter 18 develops it. Chapter 12 found that the domain of comparison, the range in which comparisons can be made, has two edges of reach: an outer edge marked by the speed of light and an inner edge marked by Planck's constant. A third limit of the domain, of another kind, lies at our own scale: the present, the

one moment in which any measurement is made. So the domain has three limits, the two edges and the present, and Chapter 18 develops the limits at the three scales as the three horizons of what can be known from here. The three limits of the domain are the same three limits on knowing that I found in Chapter 9, seen from the domain: the absolute elsewhere at the outer edge, the uncertainty relation at the inner edge, and the limit in time at the present. Each of the three has to do with time. The outer edge is fixed by how much time has elapsed: for a given elapsed time, the speed of light sets the farthest reach of any comparison. The inner edge is fixed by the bound Planck's constant sets on how finely a comparison can resolve, for the energy brought to it. The present is fixed by the fact that measurement happens only now. The three limits cannot be swapped for one another, and, as the third axiom showed earlier in this chapter, not being interchangeable is asymmetry at the root; so the three limits, too, show the structural arrow.

What We Now Know

We now know why time has a direction. Time has a direction because comparison has a structure. The structure of comparison, in which symmetry and continuity cannot both be perfected together, makes sure that no process, at any scale, can be reversed faithfully. The laws run backward lawfully; no process runs backward faithfully. The arrow does not need entropy. It needs only the fact that something compares with something else.

The claims of this chapter are of three kinds, and each kind rests on a different ground. That there is an arrow of time beneath the entropic arrow, and that perfect time reversal cannot be, is a structural exclusion. That this arrow is the order of precedence of the second axiom, applied to physics, is a claim of the metaphysics itself. That the observed breaking of time symmetry in the weak interaction fits the impossibility of perfect time-reversal symmetry is my reading of a known result. The argument

establishes the direction of time; the Appendix sets out, for the specialist, the further premises the entropic arrow needs and what the argument does not establish.

The result of this chapter matters for my question about time travel because the record follows the act it records, always, and any "going back" would have to undo that order. A time machine, in the sense the phrase intends, would return a traveler to an earlier event as the event it was. For that, the findings made since that event, by the traveler and by everything the traveler has compared, would have to be not yet made: each record laid down after that event would have to be undone, and undone faithfully, so that the earlier event stands again as it stood when it was found. Physics offers two kinds of reversal, and neither undoes, faithfully, the records laid down since an earlier event. The undo of the smooth change of a quantum state reverses a state in evolution time, and does not touch a finding. A laboratory reversal, such as the spin echo, reverses the drift of a few particles, in part, and its result is itself a new finding, later in event time. No reversal is perfectly faithful, and no reversal removes a finding from the order of findings. The record is not undone.

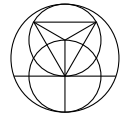
Summary

The arrow of time is structural at its root, and not entropic. The laws of physics are time symmetric to an extraordinary degree, and that symmetry guarantees lawful reversal: the reverse of a lawful process is lawful. Faithful reversal, which retraces exactly the states of the original back to exactly its starting condition, needs a state specified completely, with no gap, and that complete specification is continuity. Perfect time reversal would therefore need the symmetry of the laws and the continuity of the state, perfected together, which is the pair the Incommensuration Theorem forbids. Lawful reversal is not faithful reversal, and the gap between them cannot be closed.

The structural arrow is present in every single comparison: the range of the possible comes first, then the act, then the record, and every record follows the act it records. That order is the order of the second axiom, and the direction of time is the direction in which comparison proceeds. The structural arrow needs no large number of parts, no coarse-graining, and no beginning of low entropy; the structural arrow is a precondition of every process, including the processes from which entropy is calculated. The entropic arrow points the same way, and every act of recording joins the two. The arrow belongs to event time, the order of findings, which the block universe drops when it takes geometric time for the whole of time.

Physics can reverse two things: the smooth change of a quantum state, by an undo in the mathematics, and, in part, the traces of a motion held in particles a pulse can reach, as in the spin echo. Neither reversal removes a finding from the order of findings, and neither reversal places a comparer back before a finding that the comparer has made; there is no backward causation. The Epistemic Sandwich, an argument anchored in the free will theorem of Conway and Kochen, reaches the same arrow from what knowing requires: the choice it leaves is between a real arrow at every layer of description and no knowledge at all. The observed breaking of time symmetry in the weak interaction is consistent with the requirement that time-reversal symmetry cannot be perfect, and I read it as an instance of the requirement; its size is found by measurement. I establish the direction of time, and with it one plain fact: the record is not undone. That the arrow is present in every single comparison is a claim of the structure of comparison itself; that time-reversal symmetry is broken in the decays of kaons and of B mesons is a result of physics, measured; and that no reversal of time is faithful is a structural exclusion.

When physicists say that Einstein's equations allow a trip back to last Thursday, which kind of time do physicists mean: geometric time, event time or evolution time?



Chapter 14

Three Kinds of Time

Last Thursday

When people talk about going back to last Thursday, which "time" do they mean? The question has a definite answer, and the answer carries the whole of this chapter.

The everyday picture treats time as a kind of place. You are "at" Tuesday the way you are "at" London. If only you had the right vehicle (a DeLorean, a TARDIS, a wormhole), you could go to Wednesday, or back to last Thursday, the way you go from London to Paris. In that picture the days are cities along a road, and a calendar is a map of the road.

Time travel is usually discussed as a very hard engineering problem: a matter of strange kinds of matter, enormous energies, or physics not yet discovered. This part of the book shows that it is none of these things. The phrase "travel to an earlier time", taken in the sense the people who use it intend, names no idea at all: nothing that anyone could describe would fit it. Chapter 16 makes that exact, as a chain of steps.

A phrase can be like that. A prime number is a whole number that no whole number divides exactly except one and the number itself: two, three, five, seven, eleven, and so on. "The largest prime number" sounds as though it names a number. Every prime has a larger prime beyond it,

so the phrase names no number. The phrase has the form of a name, and nothing answers to it. "Travel to an earlier time" is a phrase of the same kind.

The picture of days as cities has a real source. It comes from *general relativity*, Einstein's theory of gravity, which he completed in 1915: the theory that gravity is the shape of space and time. Matter and energy bend the geometry of space and time around them, and a body left to itself follows the straightest path the geometry allows. In that theory one of the times of physics is laid out as a fourth direction, joined to the three directions of space in one smooth shape with four directions. I call space and time together the arena: the setting of every measurement, in which every length and every duration is measured.

Einstein's equations, taken as laws, work the same run backward. Take any solution of the equations, which is one complete shape of space and time with its matter, and flip its direction of time: the result is again a solution of the same equations. Chapter 13 showed the same thing in ordinary mechanics with a film of a ball rolling downhill: run backward, the film shows a motion that the laws of mechanics also allow.

That backward sameness belongs to the laws, the mathematics in which the theory is written. The findings that measurements return have no such sameness. A given solution of Einstein's equations, the shape of one particular world with its particular matter, generally looks nothing like itself run backward. And a law, read in either direction, is not itself a happening. Things happen. Eggs break. People age. Records are made, and they are not unmade.

The time travel picture takes two steps. The first step goes from the backward sameness of the equations to a backward sameness of the one solution that describes our world. The second step goes from that sameness to the possibility of being present again at an earlier occasion. Neither step holds. The first fails because the solution of one actual world

generally has no backward sameness at all. In this chapter and in Chapters 15 and 16, I show why the second step of the time travel picture fails as a matter of concept, and not as a matter of engineering.

Three Times

Physics uses the one word "time" for three different things, each a different kind of thing in the mathematics. Each of the three goes with one of the three modalities that Chapter 4 set out: the immanent (process, what is happening here and now), the omniscient (structure, what holds the same for everyone, laid out as a pattern) and the transcendent (possibility, what makes things possible and new). These matches hold in what I call the *focus of time*: the frame of argument in which time is the thing being considered, with all three axioms of Chapter 4 held together. Said plainly, the three axioms are that process is more basic than structure or possibility; that the cycle runs one way, from possibility to act to record; and that the three modalities are distinct, inseparable and not interchangeable.

The first time is *geometric time*: time as part of the shape of space and time. The shape carries two things together. One is the time a clock shows when it is carried along a path through space and time, the path of anything that moves slower than light; physicists call it the proper time along that path. The other is the order of which events can reach which, the order Chapter 12 drew as light cones (at each event, a forward cone of every event it can affect, and a backward cone of every event that can affect it).

What a comparison finds of geometric time is smooth and uneven. It is smooth, continuous in Chapter 2's word, because the measured times change gradually from one path to the paths beside it. It is uneven, asymmetric in Chapter 2's word, because the times differ from one region to another with the matter there. A clock close to a heavy mass runs slower than a clock far from it, and the difference changes gradually as the clock is moved. The backward sameness of Einstein's equations belongs to the mathematics of general relativity, and not to geometric time as it is found.

In the focus of time, geometric time goes with the omniscient modality: time laid out as structure, the whole pattern of space and time held in view as one.

An everyday question shows why the word "time" has to be sorted before anything else is said about it. To speak of how fast time flows, one has to say which time is meant. For someone perceiving, the felt passage of time goes with the number of changes, and with how much is new, in what is perceived. The rate of a clock is another matter. It is set by the clock's place in the arena: by its height in a field of gravity, and by its motion. A clock raised higher runs a little faster than one below it, and a clock in fast motion runs a little slower than one at rest. A busy clock and an idle clock, at the same height and moving the same way, keep the same rate: how many interactions a clock takes part in has no part in setting its rate.

The second time is *event time*, the time of findings: the order of findings for a comparer. I recall Chapter 1's words here: a comparer is anything that makes a comparison (an instrument, a detector, any interaction), and the word 'comparer' never means a mind. A finding is what one act of comparison yields: one outcome chosen from a set of separate possible outcomes and laid down as a record.

What is found of event time is jumpy and even. It is jumpy, discontinuous in Chapter 2's word, because a finding is one outcome with no case in between: the counter clicks, or it does not. It is even, symmetric in Chapter 2's word, because a record, once made, reads the same for every later comparison of it: the click recorded on Monday is the same click when the record is read on Tuesday. The order of event time is the order of the second axiom: possibility, then act, then record. In the words of my book The Unquantized Arena, which shows why the arena of space and time is not made of quanta, and what a measurement is, "That order is a precedence and not a clock reading, and it cannot be reversed". A precedence is a this-before-that. In the focus of time, event time goes with the immanent modality: it is the time of the act. Event time is also called measurement time.

A picture of an empty universe shows why findings make the order of event time. Imagine a completely empty universe, and put two things in it: a capacity to perceive (say, a head with an eye in it) and a book. Let nothing change. The scene is completely still. A carving of a head and a carving of a book would then serve as well as a real head and a real book, because nothing is changing. Take two moments in that universe, and call them moment zero and moment one. The two moments cannot be told apart, because nothing is different between them.

To tell moment zero from moment one, something has to change. Only three things are there: the head, the book, and a channel of interaction between them. The book is still, and what goes on inside the head cannot be seen, so the difference can be found only in the flow between the head and something else: in the perceiving itself, which is a comparison. Comparing is telling sameness from difference. So time and comparison are bound together as ideas. Time cannot be thought of without comparison already thought of, and comparison cannot be thought of without time as a part of it. Time comes with the comparison, not before it.

Comparison, interaction and perception (perceiving being a kind of interaction and a kind of comparison) are ideas of one and the same form: whatever belongs to one of them belongs to the other two. Time is an intrinsic of comparison, in the sense Chapter 1 gave that word: a part of what comparison is. Time is likewise an intrinsic of interaction, and an intrinsic of perception; those are one statement, made three ways. The time the picture shows is event time, and each act of perceiving in it is a finding. In the words of The Unquantized Arena, "Findings define the order of event time; they do not occur within a time given before them".

The third time is *evolution time*, and it belongs to *quantum theory*. Quantum theory is the theory of matter and light, and it says that what a measurement finds comes in separate, countable steps: light is absorbed in whole packets, an atom holds only certain energies, a detector clicks or it does not. Before a measurement, the theory describes a system by its *quantum state*: a smooth mathematical wave that spreads, overlaps itself, and sets the odds of each possible finding. The theory has a rule for how

the quantum state changes, and the rule has a single dial: turn the dial forward by some amount, and the rule carries the state forward by that amount. Physicists call the rule unitary, which here means that it keeps the odds of all the possible findings adding up to exactly one, and that every turn of it can be undone exactly. Evolution time is the setting of that dial.

Evolution time belongs to the mathematics of quantum theory, and it is smooth there. Every turn of the dial has an exact undoing, the same rule with the dial turned back by the same amount, so that within the mathematics a change of state can be run backward exactly. In the focus of time, evolution time goes with the transcendent modality: it is the time of possibility, of what could be found, and with what odds.

The three times are distinct, inseparable and not interchangeable, which are the three conditions of the third axiom. Each of the three times stays distinct from the others, for plain reasons. A position in the geometry, a date and a place in the shape of space and time, picks out no result of any measurement. A finding cannot serve as the dial of quantum theory's rule, because a finding comes in a jump and its record is never undone, while the dial turns smoothly and every turn of it has an exact undoing. The space of possible quantum states, every state the wave could be in, is not a direction of space and time. And no one time is all three at once.

That no one time is all three at once is the center of this part of the book, and it follows from the Incommensuration Theorem. Chapter 3 stated the theorem: symmetry and continuity cannot both be applied, at once and absolutely, to one comparison, and neither can asymmetry and discontinuity. The Incommensuration Theorem is a structural fact about the nature of knowing. Geometric time is found smooth and uneven. Event time is found jumpy and even. A time found both as geometric time and as event time would be asked for the characters of both at once. Set the smoothness of geometric time beside the evenness of event time, and the pair is smooth and even together, symmetry with continuity; set the jumpiness of event time beside the unevenness of geometric time, and the

pair is jumpy and uneven together, asymmetry with discontinuity. Symmetry with continuity and asymmetry with discontinuity are the two pairings that the Incommensuration Theorem forbids, and a time found as both geometric time and event time would carry at least one of the two pairings. So no unified time exists: no one time is geometric time, event time and evolution time together, with each perfected. Chapter 16 makes that sentence, no unified time, the hinge of its chain of steps.

What Exists and What Is Real

Chapter 4 drew a distinction between two ways of being, and the distinction does most of the work in what follows. In the words of The Unquantized Arena: "To say that something exists is to say that it stands as structure: timeless, the same in every instance. To say that something is real is to say that it is given in an actual act of comparison: temporal, occurring now, with the future still open. A law exists; a measurement is real".

Geometric time, laid out as structure, exists in this sense. The pattern of a solution of Einstein's equations is the same pattern however it is consulted, and every event in it has its place in the pattern. Event time is real in this sense. A finding is given in an actual act of comparison, and the record it leaves follows that act.

The two times stand side by side, each in its own way of being. Geometric time is not a mere description of some more genuine time beneath it, and event time is not an illusion laid over the geometry. Each is one of the three times, and neither does the work of the other.

The distinction answers a question that the popular picture of time travel never asks. Take last Thursday. When someone says that last Thursday "is still there" in space and time, the sentence is true of last Thursday as a region of the geometry: the stretch of space and time between the two events on the Earth's path that mark the start and the end of that day. That region exists, as structure. When someone says that one could "go back" to last Thursday and be present at what happened,

the sentence is about last Thursday as findings: the acts of comparison that were given then, and the records they left. Those acts were real when they were given, and what remains of them now is their records: a photograph, a receipt, a memory, a scuff on a step. The sentence that promises time travel uses last Thursday as a region of the geometry to secure the destination, and last Thursday as findings to describe the arrival.

Time Is Not a Place

What is a place? In the ordinary sense, a place is a location to which one can go, from which one can leave, and to which one can return. London is a place: you can go there, leave, and come back next year to the same city. Geometric time has locations: every event has its place in the order of what can reach what. A comparer meets those locations in one way only: one after another, along its own path. Anything that moves slower than light moves along such a path, and along every such path the time its own clock shows increases in one direction. The comparer's own findings, the readings of its clock among them, follow one another in the one-way order of event time. So a location in geometric time is not something a comparer reaches as a place is reached.

To "travel backward in time" would therefore have to mean one of three things. The phrase could mean reaching an earlier location in geometric time, or reversing the order of one's own findings. Or it could mean what the phrase is used to mean: arriving at an earlier occasion as the occasion it was, present among its findings, so that last Thursday is both a place in the structure and a happening one takes part in again.

Chapter 16 examines each reading in turn, and shows that each one fails for its own reason. The first reading, where the geometry allows such a trip at all, describes a trip that is not to an earlier time in the sense that the phrase intends. The second asks for the order of findings to run back-

ward, the one thing that order excludes. The third asks for one time found as geometric time and as event time together, which is the demand the theorem excludes, since no one time is found as both.

The third reading's demand is the sense in which time travel is a *category mistake*: a mistake of placing a thing in the wrong kind. It takes what is real, an occasion given in an act of comparison, and treats it as though it existed, standing as structure to be visited. And, in the reading people intend, it asks one time to be found as both geometric time and event time. The confusion is between two of the three times, each of which is in good order on its own terms; it is not a quarrel between a true theory and a false one.

Physics has examined several proposals for building a time machine: paths through space and time that loop back to their own starting point; wormholes, tunnels through space and time, whose two mouths are moved apart; and shapes of space and time in which a region with such loops is made where none existed before. Chapter 15 reports them, with the theorems that bear on them. Those theorems hold without the argument of this part of the book, and the argument holds without them, apart from one principle. The argument uses one principle from that literature, and only in examining the first reading: the principle of self-consistency, set out by John Friedman, Igor Novikov, Kip Thorne and their colleagues in 1990, that in a world with such loops the only histories allowed are those that agree with themselves everywhere, along every loop.

Two Levels of a Theory

What about the places where Einstein's equations themselves stop settling what happens? The mathematics of gravity has an edge of its own, called a Cauchy horizon (after the nineteenth century French mathematician Augustin-Louis Cauchy): an edge beyond which everything given at one instant no longer settles what comes next. Physicists have proved a great deal about such edges, and Chapter 15 reports what they have proved. In

this chapter I ask a different question: what a Cauchy horizon is, in the terms of comparison, and what bearing a Cauchy horizon has on the claim that time travel names nothing.

The first thing needed is the distinction between what exists and what is real, applied to a theory. A physical theory has two levels. Its mathematics, the equations and the space of all their solutions, stands as structure, and so it exists. Its findings, the outcomes of actual comparisons laid down as records, are given in acts, and so they are real. A *solution* of a theory's equations is one full description the equations allow; for Einstein's equations, a solution is one complete shape of space and time with its matter.

The two levels matter for the Incommensuration Theorem. The theorem forbids two pairings in what a comparison finds, when the finding is perfected at its foundation: symmetry with continuity, and asymmetry with discontinuity. The theorem permits smooth symmetries in a theory's mathematics. In the words of my book Incommensuration, my book on the theorem itself, "The theorem permits the continuous symmetries of which physics is full", and it permits them "because a physical theory has two levels, its mathematics and its findings, and the theorem concerns what a comparison finds".

I identify general relativity as the description of context. Chapter 1 called context the respect in which two things compared are two, such as two places or two moments; the arena is the context of physics, and general relativity is the description of the arena, whose findings are smooth and uneven. Quantum theory is the description of what is found in that setting, whose findings come in countable steps that repeat exactly, jumpy and even. The match of the theorem's two survivors to the two theories of physics is an identification that I make, and not a theorem.

The smooth symmetries of general relativity belong to its mathematics. Its equations keep the same form however the points of space and time are labeled, the way the coastlines on a globe stay where they are however the lines of latitude and longitude are drawn across it. Its equations also

keep the same form when the direction of time is reversed. Both of those samenesses belong to the mathematics. My book Two Theories, One World, on gravity, the quantum and the mistake behind the search for one theory, puts it in one sentence: “General relativity’s symmetry lives in its mathematics, and its continuity in its findings”.

Determinism in the Equations and in the World

Chapter 5 showed what determinism is, and why no world can have it. Determinism is the claim that one law carries every state into exactly one next state, without a break and without exception. That claim asks for total symmetry (one law, the same everywhere and always) and total continuity (every state joined to the next with no gap), applied at once and without remainder: exactly the first pairing the theorem forbids. So no world is purely and completely deterministic.

That conclusion is about findings, and it allows equations to be deterministic. An equation that carries each state into exactly one next state is deterministic in its mathematics, which exists as structure, and not in the findings, which are real. A deterministic equation is therefore no counterexample: its determinism is a property of the mathematics, and no world is deterministic in what it finds. Causes are established, and determinism never. A cause is a connection without a break along an order that runs one way, smooth and uneven, and that is one of the two pairings the theorem allows.

General relativity already shows the limits of determinism inside its own mathematics. In 1987 John Earman and John Norton set out what is called the hole argument: one and the same physical situation admits many different mathematical descriptions that agree outside a region and differ inside it. Symmetry carried to its limit, the freedom to relabel every point of space and time, undoes the uniqueness of the description. On the globe of the last section, redraw the lines of latitude and longitude inside one small patch only, and leave them as they were everywhere else. Inside

the patch the labels of the points now differ; outside it nothing differs. Einstein's equations accept both drawings, so what lies outside the patch cannot settle which drawing holds inside it. Chapter 5 met that result among the places where deterministic mathematics comes apart; here it shows Einstein's equations marking where their own determinism stops. That Einstein's equations mark where their own determinism stops is needed for what comes next: a Cauchy horizon is another place where the equations stop settling one answer.

The Cauchy Horizon

Picture a complete description of the universe at one instant: everything, everywhere, at that moment. Physicists call it the initial data; I will call it the starting data. Einstein's equations take the starting data and grow from them the largest region of space and time that the data settle: every event in that region is fixed by the starting data, with nothing left over. Physicists call that region the maximal globally hyperbolic development of the data; its plain name is the predicted world grown from the starting data. Yvonne Choquet-Bruhat and Robert Geroch proved in 1969 that this predicted world is unique: there is exactly one, up to a relabelling that leaves every length and every duration the same.

That uniqueness is a theorem about the mathematics of general relativity: about the space of solutions of its equations, which exists as structure. That uniqueness is the determinism of the mathematics, in the sense of the last section. And it is the determinism that a conjecture of Roger Penrose, posed in the late 1970s, was formulated to protect.

A *Cauchy horizon* is the edge of the predicted world: the boundary beyond which the starting data no longer fix the shape of space and time. Certain exact solutions of Einstein's equations have one, deep inside a black hole, a region of space and time from inside which nothing, not even light, can get out: the solutions for a spinning black hole and for an electrically charged black hole (the Kerr and Reissner-Nordström solu-

tions). Beyond such an edge the starting data leave the geometry unsettled: where one continuation of the solution across the edge exists, others generally exist as well. In the terms of the last section, a Cauchy horizon marks the place where the uniqueness of the mathematics' development ends.

Picture a road on a map, running off the edge of the sheet. Anyone may draw the road onward on a fresh sheet, and many different continuations fit the road as far as the old sheet shows it. The predicted world is the old sheet, and the Cauchy horizon is its edge.

Penrose's conjecture is called *strong cosmic censorship*. A censor forbids what may not be shown, and the name pictures the conjecture as forbidding, in typical worlds, an edge that could be crossed. It asks whether such an edge arises from typical starting data (starting data with no special balance or perfect symmetry), or only from the exact, specially balanced solutions, and it holds that for typical starting data no continuation across the edge exists at a stated degree of smoothness. Its versions differ in how smooth a continuation must be to count, from fully smooth to merely unbroken, with creases allowed; physicists call that degree of smoothness the regularity.

That question belongs to the mathematics of gravity. It is a question about the typical behavior of the solutions of a set of equations, and what settles it is a proof within that field. Each field settles its own questions within itself, and so the degree of smoothness at which strong cosmic censorship holds or fails is left to the mathematics of gravity: the Immanent Metaphysics, as the books behind this one develop it, makes no prediction about it, and I make none. Chapter 15 reports what the mathematicians have proved about that degree of smoothness, under which assumptions, and which version the literature states as a conjecture.

This part of the book goes in a fixed order. Chapter 13 has already examined the record: why a finding, once laid down, is not undone, and why the reversals that laboratories perform, such as spin echoes, are lawful reversals and never faithful ones (a lawful reversal runs a motion back

as the laws allow; a faithful reversal would also unmake the records the motion left). I go on to say why whatever lies beyond a Cauchy horizon is geometric time, which returns no comparer to an earlier finding. Chapter 15 reports what the mathematics of gravity has found about time machines and Cauchy horizons. Chapter 16 sets the argument out as a chain of steps, each marked as what it is, and asks whether the phrase "time travel" names anything that could be done.

What Lies Beyond a Cauchy Horizon

The question of this part of the book is a different question from strong cosmic censorship. It is whether the phrase "time travel" names an idea. What the Cauchy horizon shows about that question is one result, argued in this book in four steps.

Suppose the predicted world grown from some starting data has a Cauchy horizon, and suppose it can be continued across that horizon at some degree of smoothness, any degree at all. Then the continuation is more geometric time: a further region of the solution's structure, which exists. If the continuation is not unique, there are several such regions, each a different geometry. No continuation, unique or not, makes an earlier finding of any comparer happen again.

First, a continuation of a shape of space and time is itself a shape of space and time, into which the original region fits exactly as one part, with every length and duration kept. Its clock times along paths, and its order of what can reach what, are carried by its geometry. This holds even for a continuation whose geometry is only unbroken, with creases, and not smooth; the Appendix gives the technical form of that case. So what the continuation adds is geometric time.

Second, a solution of Einstein's equations, together with each of its continuations, belongs to the mathematics of general relativity, which stands as structure: it exists. Third, if several continuations exist, the mathematics offers several geometries beyond the horizon, and the starting

data pick none of them. Several geometries are several structures; they are not an earlier act of comparison. Fourth, a finding is real: it is given in an act and laid down as a record, and its order among other findings is the order of the second axiom, which cannot be reversed. A continuation adds structure to the mathematics. It leaves the order of findings as it was, and it undoes no record. So no continuation makes an earlier finding happen again.

The result rests on four things: what a continuation of a shape of space and time is; what geometric time is; the difference between what exists and what is real; and the one-way order of findings. It holds at every degree of smoothness, and it holds whether or not strong cosmic censorship is true. Where strong cosmic censorship holds at a given smoothness, typical starting data admit no continuation at that smoothness, and the Cauchy horizon of a typical solution is the end of its development. Where strong cosmic censorship fails at that smoothness, typical starting data admit continuations, and the result says what those continuations are: more geometric time, possibly in several versions, and no return to a recorded finding. Either way, the argument of this part of the book holds, whatever the answer to strong cosmic censorship.

The same result holds for a time machine that someone builds: what lies beyond its horizon is more geometric time. A *time machine*, in the physicists' sense, is a region where paths that loop back in time appear in a world that had none before them. The edge of the region still settled by the world as it was before the machine was made is a Cauchy horizon, and here it is called a *chronology horizon*. What lies beyond it, if anything lies beyond it, is a region of geometric time in which looping paths exist. Whether a comparer traveling along such a loop could thereby return to an earlier finding is the question Chapter 16 asks of the first reading of "going back", and the answer given there is that it cannot.

Singularities and the Planck Scale

What happens deeper still inside a black hole, where the equations seem to break? The outer edge of a black hole, the surface from inside which nothing gets out, is its horizon. In the exact solutions, the description of space and time inside a black hole breaks down, and where it breaks down is called a *singularity*. The equations have two kinds of edge, and this section is here so that the two can be told apart: a Cauchy horizon is an edge where the equations stop settling one continuation, and a singularity is a place where the description of the setting itself stops holding. To say where a singularity is, one more scale is needed.

Chapter 12 named Planck's constant as the mark of the inner edge of the range in which comparisons can be made. Joined with the speed of light and Newton's constant of gravity, Planck's constant gives a length, a duration and a degree of bending, and together these make the *Planck scale*. The Planck length is smaller than the nucleus of an atom by about twenty powers of ten (a factor of one followed by twenty zeros), and it is formed from those three constants. Planck's constant sets no least length and no least duration: the Planck length marks a bound on what comparisons can find, and it is not a smallest step of space and time. The *Planck curvature* is bending so sharp that space and time curve appreciably within a single Planck length.

In The Gravitational Passage, my book that derives general relativity from the Immanent Metaphysics, I give an account of singularities in the terms of comparison. There the singularity is a breakdown of the description of the setting: the description that treats space and time as smooth and uneven, continuous and asymmetric. That description holds wherever space and time bend more gently than the Planck curvature, and it does not hold beyond the Planck curvature. The scale at which the description of the setting stops holding is what The Gravitational Passage calls the *modal boundary*. Its glossary defines it: "Modal boundary: the Planck scale, where the continuously asymmetric and the discontinuously symmetric descriptions exchange dominance; a structural transition, not a

barrier". In plain words: at the Planck scale the smooth, uneven description (general relativity's) and the jumpy, even description (quantum theory's) trade places as the one that does most of the work. The same book adds: "The boundary is not a wall that matter cannot cross, and it is not a regime of a third kind of physics". The arena itself is not broken into countable pieces there: gravitation is not quantized, at the Planck scale or anywhere, and what changes at the boundary is which of the two descriptions dominates.

The modal boundary is therefore a scale of bending, and not a surface in space and time. A Cauchy horizon is a surface in space and time. In the exact solutions for spinning and charged black holes, the bending on the Cauchy horizon is finite, and for a black hole with the mass of a star it is far below the Planck curvature. So a Cauchy horizon and a modal boundary are two different things, one a surface in space and time and the other a scale of bending, and this book does not call a Cauchy horizon a modal boundary.

On the insides of black holes, The Gravitational Passage states an expectation, and not a theorem: "For rotating and charged holes, the inner horizon of the exact solutions is unstable to the infall that forms real holes, and the generic end is expected to be of the same kind" as the end of the inside of a hole without spin or charge. That book describes that end as a moment, "the future end of every path inside the horizon", and not as a place. In plain words: a spinning or charged black hole has an inner horizon inside its outer one, and in the exact solutions that inner horizon is the Cauchy horizon. A real black hole forms from matter falling in, and that falling matter is expected to break up the inner horizon, so that the inside ends the way the inside of a plain black hole ends: in a last moment for every path within it, and not in a place where anything could stand. The expectation agrees in direction with the instability of the Cauchy horizon that the mathematicians report. The expectation leaves unsettled the degree of smoothness at which any continuation across the Cauchy horizon fails, and Chapter 15 reports what has been proved about that.

Where the Price of Being Inside Is Paid

Does the price an observer pays for being inside the universe bear on Cauchy horizons? Chapter 3 named that price, the price of being inside empty space, as one face of the Incommensuration Theorem. The price is paid at a horizon, and I have been speaking of horizons throughout, so a reader may fairly ask whether it is paid at a Cauchy horizon too. The answer is that the price is paid at exactly one horizon, and that horizon is the observer's own. In The Gravitational Passage I identify the *cosmological constant* with the share of the energy of empty space that an observer inside empty space can reach. The cosmological constant is a quantity in Einstein's equations that acts as an energy belonging to empty space itself, and that makes the expansion of the universe speed up. Physicists call the energy of empty space the vacuum energy.

Said in words, the formula is that the energy that the cosmological constant gives to each unit of volume of empty space is three, times a pure number of about 0.881, times the Planck energy density, divided by the count of the observer's horizon. The three is a whole number that book draws from the structure of measurement. The pure number, a little under nine tenths, is one minus three divided by eight times pi. The Planck energy density is the energy in each unit of volume at the Planck scale, formed from Newton's constant of gravity, which is measured. The count of the horizon is a quarter of the horizon's area, counted in squares one Planck length on a side. The Appendix gives the formula in its technical form, with what each of its parts is in that book.

The horizon in that formula is the observer's *particle horizon*: the edge of what the observer can see of the universe, the farthest distance from which light has had time to reach the observer. The three and the 0.881 are set by the structure, once that book's identification is taken; Newton's constant and the radius of the horizon are measured. That book states the formula as an identification, from which a value of the cosmological constant is computed. What the reachable share leaves out is what that book calls the *information deficit*, "the price of being inside".

The formula applies at one horizon: the particle horizon of the observer who measures. The Gravitational Passage states what follows for black holes: "the observer, outside the hole, pays no deficit at its horizon. For the exterior, a black hole's horizon carries no vacuum energy of its own, and the cosmological constant is set by one horizon, the particle horizon of the measurement". The same book states that a comparer inside a black hole keeps the count of its own particle horizon, so that for that comparer too the reachable share is unchanged. So the information deficit gives no account of the inner horizon of a black hole, and no account of any Cauchy horizon, and I draw no conclusion about Cauchy horizons from it.

What We Now Know

We now know which time people mean when they talk of going back to last Thursday: two times at once. Geometric time exists: last Thursday, as a region of the shape of space and time, stands as structure. Event time is real: last Thursday, as findings, was given in acts of comparison, and what remains of those acts is their records. No one time is found as both, because a time found as both would carry a pairing the Incommensuration Theorem forbids.

We also know what lies at the edges of Einstein's equations. Beyond a Cauchy horizon, the edge where the starting data stop settling the shape of space and time, lies more geometric time, and no continuation there, one or many, at any smoothness, makes an earlier finding happen again.

What we now know is the ground on which the book excludes time travel, which Chapter 16 makes exact. "Time travel" asks one time to be both: last Thursday as a place, to secure the destination, and last Thursday as a happening, to describe the arrival. The request is a category mistake before it is ever an engineering problem, and the edges of Einstein's equations, wherever they fall, offer more geometry and never an earlier finding.

Summary

The word "time" names three things in physics. Geometric time is the time a clock shows along its path and the order of what can reach what, carried by the shape of space and time; it is found smooth and uneven, and in the focus of time it goes with the omniscient modality. Event time is the order of findings for a comparer; it is found jumpy and even, its order is a this-before-that that cannot be reversed, and it goes with the immanent modality. Evolution time is the dial of quantum theory's rule of smooth change, smooth and exactly reversible within the mathematics, and it goes with the transcendent modality. The three are distinct, inseparable and not interchangeable, and no one time is all three with each perfected.

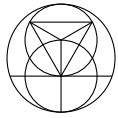
Geometric time exists, as structure; event time is real, as findings given in acts; neither is an illusion and neither does the other's work. A place is somewhere one can go, leave and return to, and a location in geometric time is met only once, in order, along a comparer's own path. Time travel, in the sense intended, takes a happening and treats it as a structure to be visited, and asks one time to be both geometric time and event time: a category mistake between two times that are each in good order.

A theory has two levels, its mathematics, which exists, and its findings, which are real, and the Incommensuration Theorem concerns findings; general relativity's symmetries, its backward sameness among them, live in its mathematics. The uniqueness of the predicted world is the determinism of that mathematics, and a Cauchy horizon marks where that uniqueness ends. Whatever lies beyond a Cauchy horizon, or beyond the chronology horizon of a built time machine, is more geometric time, which exists; no continuation makes an earlier finding happen again. This result covers every degree of smoothness and both answers to strong cosmic censorship; the smoothness question itself belongs to the mathematics of gravity, on which this book makes no prediction. A singularity is where the description of the setting stops holding, at the Planck scale; the modal boundary there is a change of structure, not a wall, and a scale of bending,

not a surface; gravitation is not quantized there or anywhere. The information deficit, the price of being inside, applies at the observer's particle horizon, and at no horizon of a black hole.

The chapter's statements are of several kinds, and each keeps its own standing. That no unified time exists is a claim from the structure of comparison itself: it is the Incommensuration Theorem applied to the characters of what is found of geometric time and of event time, as The Unquantized Arena proves it. The matching of the three times with the three modalities, and of general relativity with the description of context, are identifications, readings of known physics in the terms of comparison. The difference between what exists and what is real is a definition, quoted. The result about what lies beyond a Cauchy horizon is argued in this book. The uniqueness of the predicted world (Choquet-Bruhat and Geroch, 1969) and the hole argument (Earman and Norton, 1987) are results of the literature, reported with their authors. The accounts of the singularity, the modal boundary and the information deficit are reported from The Gravitational Passage, which states the formula for the cosmological constant as an identification. And the category mistake that the phrase "time travel" makes is the ground of a structural exclusion, which Chapter 16 states as a chain of steps.

But physicists have written down solutions of Einstein's equations in which paths through space and time loop back to their own beginnings: what exactly have they found, and what does it show?



Chapter 15

Time Machines on Paper

A Machine Drawn on Paper

Haven't physicists found solutions to Einstein's equations that allow time travel? They have found solutions with loops in them, and this chapter says exactly what those solutions are, what has been proved about them, and what is still stated as a conjecture.

Begin with a phrase that Chapter 14 used: "the largest prime number". The phrase has the grammar of a name, like "the tallest mountain", and it names no number, since every prime is exceeded by a larger prime. A phrase can be well formed and still have nothing answer to it.

Now picture something of a different kind: a drawing. A physicist writes down a whole shape of space and time, checks that it obeys Einstein's equations, and finds in it a path that a body could follow and that brings the body back to the very place and moment it set out from. The drawing is real mathematics, and every line of it can be checked. Call it a time machine on paper.

The two pictures raise two different questions. The drawing raises a question of geometry: which shapes do Einstein's equations allow, and which do they forbid? The phrase raises a question of meaning: when someone says "go back to last Thursday", does the phrase name anything

that could be done? In this chapter I answer the first question, the question of geometry, from what physicists and mathematicians have found. The next chapter answers the second.

Every result in this chapter belongs to the physics and mathematics it reports, and it rests on the proofs of the people who proved it. The argument of this book takes one principle from this chapter, the principle of self-consistency, and uses it at one step only, in the next chapter; it uses none of the other results reported here. The results are reported because they are what physics itself says about the machines that the phrase "time travel" imagines. They show what general relativity, Einstein's theory of gravity as the shape of space and time, does and does not permit when it is taken on its own terms, and they serve as checks on the conclusion of the next chapter.

Loops in Space and Time

Mark where a body is at every moment of its life, and join the marks. The line through space and time that results is the body's *worldline*. Each point on it is an event, a place at a moment. A body with mass always moves slower than light, so its worldline stays inside the light cones of its own events, the cones that Chapter 12 described. At each event, the forward cone holds every event that the event can affect. Physicists call a path of that kind, one a body with mass could follow, *timelike*.

A *closed timelike curve* is a timelike worldline that returns to its own starting event: a path through space and time that a body with mass could follow and arrive again at the very event from which it set out. I call such a path a time loop. Time loops belong to the mathematics of general relativity in this sense: some shapes of space and time that satisfy Einstein's equations contain them. Physicists call one whole shape of space and time a *spacetime*.

The mathematician Kurt Gödel gave the first example in 1949. Gödel's spacetime has a time loop through every point. To see what follows from that, one more term is needed. Physicists call a slice through a world a *Cauchy surface* when every path that a body or a flash of light could take, followed as far as it goes, crosses the slice exactly once. The conditions on such a slice then settle the whole world, past and future. Gödel's world has no Cauchy surface at all. Every point of it lies on a time loop, so every collection of its points, even a single point, is joined to itself by some time-like path (physicists say that no set in it is achronal). So the world has not even a part of a slice from which an edge of prediction could form. Chapter 14 called that edge the Cauchy horizon: the boundary of the region that the starting conditions settle, beyond which they no longer fix the shape of space and time. So Gödel's world settles a general question: a world can have time loops and no Cauchy horizon, and the general statement "every world with time loops has a Cauchy horizon" is false.

What Physicists Mean by a Time Machine

A *time machine*, in the sense the physics literature gives the words, is narrower than a time loop. It is a region in which time loops appear in a world that had none before them. Up to some time the world is orderly: everything in it is settled by the conditions on a starting slice (physicists call such a world, or such a part of one, *globally hyperbolic*). Loops appear only after that time, and the usual order of cause and effect is broken only after it. The edge of the region that the starting conditions still settle is then a Cauchy horizon of the starting slice, and in this setting physicists call it the *chronology horizon*.

The best known design for a time machine uses a wormhole: a tunnel through space joining two distant places, whose two ends are called its mouths. In 1988 Morris and Thorne described a wormhole that a body could pass through, held open by matter that breaks the energy conditions. The *energy conditions* are rules that ordinary matter obeys; the

simplest of them, the weak energy condition, says that the energy in any region, as any observer measures it, is never negative. In the same year Morris, Thorne and Yurtsever showed that moving one mouth of such a wormhole relative to the other, so that the two mouths age differently, produces time loops through the wormhole.

In 1992 Stephen Hawking asked how a time machine could form. A chronology horizon, like every Cauchy horizon, is a surface made of light paths, in the way a sheet can be pictured as made of threads. Hawking took as his model a time machine made in a finite region, by a finite apparatus, with no place where the bending of space and time grows without limit. He wrote that demand as a condition on the horizon: every light path that makes up the horizon, followed back into the past, enters a bounded region and stays within it (physicists call such a horizon compactly generated). A time machine with a compactly generated horizon is Hawking's model of a made time machine. Hawking's theorem is a theorem about that model: the theorem covers time machines of the model's form, and the theorem is not a theorem that every time machine has the model's form.

Hawking proved that, in his model, a time machine growing from a starting slice that runs on without end, like ordinary space, which goes on without end, needs matter that breaks the weak energy condition. Such a horizon cannot form from such a slice if the energy every observer measures is never negative. The summary at the head of his paper states the matching result for an averaged form of the rule: if the time loops develop from a starting slice that runs on without end, then the averaged form of the weak energy condition must fail on the horizon. Hawking then proposed, as a conjecture, that the laws of physics prevent time loops from appearing at all. He called it the *chronology protection conjecture*. A *conjecture*, in mathematics and physics, is a statement proposed for proof that has not been proved.

In 1997 Kay, Radzikowski and Wald proved a result about quantum fields in such worlds. A quantum field is quantum theory's description of matter or light spread through space. The simplest is a field with a single

number at each point and no direction, which interacts with nothing, itself included (physicists call it a free scalar field). Physicists work with states of a field that are well behaved at very short distances: states that, looked at closely enough around any point, behave as the field does in empty, flat space (physicists call them Hadamard states). Because Hadamard states behave, close up, as the field does in empty, flat space, physicists can compute from Hadamard states the field's energy and pressure point by point. Kay, Radzikowski and Wald proved that, for the free scalar field in a world with Hawking's kind of horizon, every state of that well-behaved kind, carried on beyond the region that the starting conditions settle, stops being well behaved at every point of a set on the horizon called its base set. In the picture of the sheet made of threads, the base set is where the threads, the light paths that make up the horizon, pile up when they are followed back into the past. At those points the field's energy and pressure cannot be defined by the usual procedure. The result concerns the points of the base set, not every point of the horizon.

A second line of work grants the loops and asks what follows. In 1990 Friedman, Morris, Novikov, Echeverria, Klinkhammer, Thorne and Yurtsever studied how fields in a world with time loops are fixed by their starting conditions. They worked under the *principle of self-consistency*: the only histories admitted are those that agree with themselves everywhere, including all the way around every time loop. Under that principle, whatever a body carries around a time loop must arrive back at the starting event in exactly the state in which it set out. The principle of self-consistency is the one principle of this chapter that I use in the argument of this book. The next chapter uses it when it considers a comparer (anything that makes a comparison and lays down a record of what it finds) carried around such a loop, laying down records as it goes.

The Edge Where Prediction Ends

Roger Penrose introduced the strong cosmic censorship conjecture in the late 1970s. It states a concern about whether Einstein's theory predicts. Give Einstein's equations suitable starting conditions on a slice through space at one moment, and they settle one unique world: the largest region that those conditions settle, with nothing left to choose. Physicists call it the maximal globally hyperbolic development of the starting conditions; Chapter 14 called it the predicted world, and so does this chapter. Physicists call this uniqueness the determinism of general relativity, in the sense of Laplace, whose all-knowing demon was the picture of Chapter 5: the present, given on a single slice, fixes the future and the past throughout the region that the starting conditions can reach.

In the terms of this book, that uniqueness belongs to the equations of general relativity. The equations are a law, and a law exists as structure: a pattern, laid out the same for every comparer. The uniqueness is not a property of what any comparison finds. Chapter 5 showed that no world is purely and completely deterministic. In the words of my book Measurement; The Steps from Comparison to Physics, the development of a state "is deterministic in the formalism, which exists as structure, and not in the findings, which are real", the formalism being the mathematics in which a theory is written. So the determinism at stake in this chapter is the determinism of the equations.

Certain exact solutions of Einstein's equations, shapes written out in full in a formula, have an edge to that uniqueness. Two families matter here. The Kerr family describes spinning black holes in empty space, and the Reissner-Nordström family describes charged black holes. A black hole is a region from which nothing escapes, not even light; its outer surface, the *event horizon*, is the surface from inside which nothing gets out. Inside each black hole of these two families lies a second surface, the inner horizon. It is a Cauchy horizon: a surface moving at the speed of light (physicists call such a surface null), beyond which the starting conditions no longer settle the world. Beyond it, the solution can be continued in

more than one way, each different from the others: different regions of space and time can be attached to the same predicted world, and every one of them obeys Einstein's equations. The Cauchy horizon is the edge of what the starting conditions settle.

Chapter 14 pictured this edge with a road on a map, running off the edge of the sheet: many different continuations fit the road as far as the old sheet shows it. The predicted world is the old sheet, and the Cauchy horizon is its edge.

Penrose's conjecture says that this edge is special to the exact solutions. On this view, the exact solutions that have Cauchy horizons are unstable: nudge their starting conditions in a typical way, and the Cauchy horizon is replaced by a *singularity*, a place where the shape of space and time breaks down, past which no continuation of a stated smoothness exists. The predicted world of typical starting conditions then cannot be continued at that smoothness. Physicists say "generic" where I write "typical", and a later section says exactly what the word has to mean. The conjecture grants that exact solutions have Cauchy horizons; what it says is that typical starting conditions do not produce Cauchy horizons that can be crossed.

Strong cosmic censorship, then, is Penrose's conjecture that the predicted world of typical starting conditions cannot be continued past its edge, at a stated smoothness. Strong cosmic censorship protects the uniqueness of the prediction: the determinism of the equations, and not a determinism of findings.

The argument of this book holds however the conjecture is settled, and it uses no result of strong cosmic censorship. Chapter 14 showed that whatever continuation a world has beyond a Cauchy horizon is further geometric time, more of the laid-out shape of space and time, and that no continuation, single or one of many, makes an earlier finding of any comparer occur again. The rest of this chapter reports what the literature has found about the conjecture, because that is what physics says about the edges at which a time machine would have to form.

The conjecture and Hawking's theorem meet at one point, and together they say what kind of matter a time machine of Hawking's kind would need. The meeting point is that a chronology horizon is a Cauchy horizon. The conjecture is posed for empty space and for kinds of matter that obey the energy conditions, while Hawking's theorem shows that a time machine of his kind needs matter that breaks the weak energy condition. Take first the kinds of matter for which the conjecture is posed. Within those kinds of matter the weak energy condition holds, so Hawking's theorem already rules out a time machine of his kind growing from a starting slice that runs on without end. The conjecture, stated at a given smoothness, adds that within those kinds of matter Cauchy horizons that can be crossed at that smoothness, chronology horizons included, arise only from starting conditions that are not typical. A time machine of the kind Hawking models therefore needs matter outside those kinds.

Three Standards of Smoothness

The next six sections report, for the reader who wants it, what the mathematics of general relativity has proved about the conjecture so far. The argument of the book uses none of it, and the reader who passes over them loses no step of the next chapter; the chapter's line resumes at "What the Equations Allow".

What the conjecture says depends on how smooth a continuation must be to count. The question "does the world end at its Cauchy horizon?" has a different answer for each standard of continuation one admits. Three standards are in use, and they differ. Go back to the road that runs off the edge of the map. One can ask three things of a drawn continuation. Is the road unbroken, with no gap? Is it smooth enough that the rules for driving on it can still be written down? Is it free of any bend so sharp that it would throw a car off the road?

The strictest version of the conjecture says that the shape of space and time cannot be continued past the edge even as an unbroken shape: no continuous continuation exists at all. It is the strongest claim of the three. If not even an unbroken continuation exists, then no smoother one exists either.

The middle version was stated by Chruściel in 1991 and by Christodoulou in 2009. It says that no continuation exists whose shape is unbroken and whose rates of change are square-integrable. The rates of change of a shape say how its measures of length and duration change from place to place. Square-integrable means this: the rates may grow very large at some points, but their squares, added up over any bounded region, come to a finite amount. A shape whose rates of change are square-integrable is just smooth enough for Einstein's equations to be written down in the loosest form of Einstein's equations. In that loosest form, Einstein's equations are checked on averages over small regions rather than point by point. That level is the roughest at which the equations can still be written at all. So, under this version, a world could carry on past its edge as an unbroken shape while Einstein's equations, even in their loosest form, cannot be written beyond it. I use the middle version when I speak of strong cosmic censorship without a qualifier, because it marks where the theory stops predicting: a continuation on which the equations cannot be written is a continuation that the equations do not govern.

The loosest version says that no twice-smooth continuation exists: none whose shape has well-defined rates of change, and well-defined rates of change of those rates, at every point. That no twice-smooth continuation exists is usually shown by the bending of space and time, its curvature, growing without limit at the edge. A body that crosses such an edge is stretched and squeezed without limit, because the stretching a body feels, its tide, is the difference in how neighboring freely falling paths accelerate, and the curvature sets that difference. This version is the weakest claim of the three: the bending can grow without limit while the shape itself still continues unbroken.

So the three stand in order. Strictest: no unbroken continuation, no road at all past the edge. Middle: no continuation on which the equations can be written, no road on which the rules of driving can be stated. Loosest: no continuation free of unlimited bending, no road without a bend that throws the car.

Blue Light at the Inner Horizon

What works against the Cauchy horizon is a process physicists call the blue-shift instability. Picture a body falling, inside a black hole, toward the inner horizon. Light and other radiation keep falling into the black hole after it, and that radiation reaches the body squeezed into shorter and shorter waves as the body nears the inner horizon. Squeezing a light wave shorter moves its color toward blue and raises its energy, which is why the effect is called a *blue-shift*. The blue-shift is like a recording played faster and faster, its pitch climbing without limit. The energy of that radiation, as the falling body measures it, grows without limit at the Cauchy horizon. The growth is no effect of how positions are labeled; it follows from the pattern of which events can affect which, inside the black hole near its inner horizon.

Whether the blue-shift destroys the Cauchy horizon is decided by a contest between two things: how fast the disturbances outside the black hole die away, and how fast the blue-shift magnifies what is left of them. The disturbances are small ripples, in the shape of space and time or in a field, on the black hole. The *cosmological constant* is a single number in Einstein's equations; when it is positive, space far away spreads apart ever faster. With no cosmological constant, ripples along the event horizon fade slowly, like one over a power of the elapsed time; physicists call this slow tail of fading Price's law. Near the inner horizon the blue-shift magnifies the ripples exponentially, multiplying them by the same factor in each equal interval of time, as doubling does. Doubling ten times multiplies by about a thousand, and doubling twenty times by about a million,

and growth of that kind outruns any fading by a fixed power of time. So exponential magnification is expected to win, and the rates of change of the ripples are expected to blow up at the Cauchy horizon, for every such slow rate of fading. With no cosmological constant, the contest is not close. A close contest of two rates arises only when the cosmological constant is positive, because there the fading is itself exponential.

In 1990 Poisson and Israel gave the first model of how the blue-shift acts back on the shape of the black hole's interior, and in 1991 Ori solved exactly a model of a round, charged black hole. The effect they found is called *mass inflation*. Picture a sphere centered on the black hole. Its size can be read off from its area: its area radius is the radius that a sphere of that area would have in flat space. From the area radius, and from how it changes from place to place, physicists compute a measure of the mass enclosed inside the sphere (the Misner-Sharp mass). At the Cauchy horizon that enclosed mass grows without limit, while the area radius stays finite and does not shrink to zero. In Ori's solution the Cauchy horizon is a singularity along a surface moving at the speed of light. There the enclosed mass grows without limit, the area radius stays finite and above zero, the measures of the shape itself stay finite, and a body crossing it is stretched by only a finite amount. A singularity of this kind is called a *weak null singularity*: weak because the shape itself holds together, null because it lies along a surface moving at the speed of light.

Dafermos proved the matching statements in 2003 and 2005, for a round, charged black hole containing a scalar field, a field with a single number at each point. For starting conditions that settle down to a charged black hole of the Reissner-Nordström family, the shape, the area radius and the field all continue unbroken up to the Cauchy horizon, and the area radius stays above zero on the Cauchy horizon near its future end point. If the field's fading along the event horizon is held from below, so that it does not die away too fast, then the enclosed mass blows up all along the Cauchy horizon, and the shape cannot be continued even once-smoothly, with well-defined rates of change. So, in the round case, the

Cauchy horizon is an edge across which the shape continues unbroken, the spheres keep a positive area, and the enclosed mass grows without limit.

These results are proved for round, charged black holes. For spinning black holes, which lack the round symmetry, the results for round black holes are the picture that guides the work, and not a theorem. In the next section I report what has been proved for spinning black holes.

What Has Been Proved

The proofs sort two ways: by how symmetric the world is, and by which standard of smoothness the statement uses. The Appendix gives each proof with its authors, its date and every assumption it makes; I say where the proofs stand.

For round, charged black holes holding a scalar field, the loosest version of strong cosmic censorship is proved: Luk and Oh proved in 2019 that, for typical starting conditions, the predicted world has no twice-smooth continuation into the future. In 2018 Sbierski proved that the plain black hole, neither spinning nor charged (the Schwarzschild black hole), cannot be continued at all, not even as an unbroken shape. The plain black hole has no Cauchy horizon inside it; it ends at a singularity that stretches across space like a last moment (physicists call it spacelike). The plain black hole is the contrast case: a world that ends in the strictest sense. It also shows that an edge of a predicted world can be something other than a Cauchy horizon.

For spinning black holes in empty space, the results stand at three levels. That the shape continues unbroken across a piece of the Cauchy horizon is a theorem of Dafermos and Luk, published in 2025, under its stated hypotheses inside the black hole. Joined to the stability of the outside of a spinning black hole, and to a second theorem that Dafermos and Luk have announced, it makes the strictest version of strong cosmic censorship false for spinning black holes in empty space near the exact spin-

ning solution. That the shape cannot continue with bounded rates of change is proved under stated assumptions on the starting conditions, in two preprints of 2026, one by Gurriaran and one by Luk and Sbierski, with a theorem of Sbierski of 2026 for the step from bending that grows without limit. (A preprint is a paper posted for other scientists to read before a journal's referees have checked it.) That the shape cannot continue with square-integrable rates of change, so that the equations cannot be written past the edge, is the middle version itself, which the literature states as a conjecture.

In the picture of the road: it is proved that the road runs on past the edge of the sheet without a gap, at least for a stretch; it is proved, under stated assumptions, that it cannot run on without kinks that grow sharper without limit; and whether the rules of driving can still be written on it past the edge is the middle version, stated as a conjecture. Every continuation with bounded rates of change is also one with square-integrable rates of change, since a bounded amount, squared and added up over a bounded region, comes to a finite total. The reverse fails: a shape with no continuation of bounded rates may still have one whose rates are square-integrable. So the results on bounded rates narrow the distance between what is proved and the middle version, and the middle version itself remains stated as a conjecture.

When the Cosmological Constant Is Positive

The contest changes when the cosmological constant is positive. Space far from the black hole then spreads apart ever faster, and the black hole's world gains a second edge, far outside the black hole: the cosmological horizon. A disturbed black hole rings like a struck bell, in a set of tones each of which dies away at its own rate. With a positive cosmological constant, ripples along the event horizon fade exponentially, at a rate set by the slowest-fading tone. Near the inner horizon the blue-shift magnifies what is left, exponentially too, at a rate set by a single number belonging

to the inner horizon (physicists call it the surface gravity of the inner horizon). The contest is then between two exponential rates, and it is decided by their ratio: the fading rate divided by the magnifying rate.

In 2017 Hintz and Vasy proved how smooth small ripples on such black holes are at the Cauchy horizon, and from their theorem the dividing line for the middle version is a ratio of one half. The dividing line holds in the linear analysis: the analysis of small ripples, treated as so small that they do not change the shape they move on. Where the ratio is below one half, the middle version is expected to hold; where it is above one half, the middle version is endangered. The number one half is a pure number, a ratio with no units. Its standing: it is set by the arithmetic of the proved theorem about small ripples; it is a dividing line drawn from that theorem, and it is not a theorem about the full equations.

The direction of the dividing line runs against a first guess. Fast fading outside is what endangers the conjecture, and slow fading is what protects it. Think of an amplifier and a loudspeaker. The amplifier can wreck the loudspeaker only if enough signal reaches it; if the signal dies away fast before it arrives, too little is left to amplify, and the loudspeaker survives. The blue-shift is the amplifier, and it is what destroys the Cauchy horizon. If the ripples fade too fast before they reach the inner horizon, too little remains for the blue-shift to magnify, and the Cauchy horizon survives, smooth enough to be crossed by continuations on which the equations can still be written in their loosest form. So fast fading lets the Cauchy horizon survive, and slow fading feeds the blue-shift that destroys it; the tempting sentence "fading overcomes the blue-shift, so strong cosmic censorship holds" states the relation backward.

For spinning black holes with a positive cosmological constant, Dias, Eperon, Reall and Santos found in 2018, in the linear analysis, that the ratio stays below one half at every spin below the greatest. The spinning case with a positive cosmological constant is therefore expected to respect the middle version.

The charged case differs. Take black holes with a positive cosmological constant and nearly the greatest charge their mass allows, rippled by a smooth field with no charge and no mass. Cardoso, Costa, Destounis, Hintz and Jansen found in 2018 that for these black holes the ratio is above one half. For that kind of ripple on that kind of black hole, the middle version is violated in the linear analysis, and the black holes on which the violation occurs fill a whole region of the possible black holes, rather than a vanishingly thin set. So no argument that sets aside vanishingly thin sets can remove the violation.

The literature discusses three replies to this violation: that charged or massive fields change the rate of fading; that rough starting conditions, rather than smooth ones, bring the blow-up back; and that the energy of quantum fields grows without limit at the Cauchy horizon of these charged black holes, whatever the rate at which the ripples fade. Each reply treats a different side of the violation, and the literature accepts none of them as settling it. A further position is held in the literature: that a conjecture whose statement must set aside a whole region of smooth counterexamples, in the theory without quantum effects, may be the wrong statement of the conjecture. I record that position, and leave the field to decide it. The Appendix gives the three replies with their authors, and a physical argument for strong cosmic censorship from the generalized second law, with its standing.

What "Typical" Means

The word "typical" carries content. The conjecture says that, for all starting conditions outside a negligible set, blow-up does occur, and not only that it can occur. A starting condition gives a value at every point of a slice, so the collection of all starting conditions has endlessly many independent directions in which a condition can be varied. For a collection that large, mathematics has more than one meaning of "negligible". Three meanings frame the question.

Picture a sheet of paper with a pattern of fine lines on it, such that however small a patch of the paper you pick, there is a smaller patch inside it that the lines miss. A set like that is called *nowhere dense*: it is thick nowhere. Now lay one such pattern over another, and another, in a list that never ends but can be counted off, first, second, third. A set made of countably many nowhere dense pieces is called *meager*: thin in the sense of position. Meager, thin in the sense of position, is the first meaning of "negligible".

The second meaning is thin in the sense of chance. Model the uncertainty in the starting conditions by spreading a probability over them in the shape of a bell curve, most likely near the conditions of interest and less likely further away. A set that this probability gives zero chance of being hit has *measure zero*. Hunt, Sauer and Yorke defined in 1992 a variant of this meaning, called prevalence: another way of saying "almost every" in the sense of chance, built for collections with endlessly many directions. Prevalence and thinness of position are different notions: a set can be almost everything in the sense of prevalence and still meager, or almost everything in the sense of position and still negligible in the sense of prevalence.

The third meaning comes from describing starting conditions. Describe a condition to finite precision, then ever finer, as one writes a number to more and more decimal places. Call a description decisive when every starting condition that fits it sets off the blow-up. The question is whether every description, however it begins, can be refined into a decisive one. Asking whether every description, however it begins, can be refined into a decisive one is the same as asking whether the set where the middle version fails is nowhere dense. The third meaning implies the first, and the reverse fails: a nowhere dense set is meager, but a meager set need not be nowhere dense. The second meaning is independent of the first. None of the three proves that blow-up is typical; each states exactly what "typical" claims, in its own sense of negligible.

The failure set is the set of starting conditions whose predicted world can be continued across the Cauchy horizon with the equations still written: the set where the middle version fails. In the sense of position, mathematicians can rank how complicated the failure set is. Granting a technical assumption, the rank is this: the failure set is at most of the simplest kind beyond the plainly definable sets (mathematicians call such a set analytic, the first level of what they call the projective hierarchy). The set of starting conditions outside it is at most of the matching kind (coanalytic). The rank shows that the failure set is either thin in the sense of position or fills almost all of some region in that same sense. Which of the two holds is a well-posed question in the standard axioms of mathematics, one whose answer stays the same however those axioms are consistently filled out, and one that can be posed with no assumption beyond them. The rank poses the question, and it leaves the question unanswered. Knowing how complicated the failure set is does not tell whether it is thin, so an argument that concludes "the failure set is thin" from its rank assumes its conclusion. The Appendix gives the rank in its technical form, with the results of set theory on which it rests and their authors. The statement "the failure set is thin" is the middle version with "typical" read as "outside a thin set", and the literature states it as a conjecture.

In the sense of chance, the charged black holes of the last section bear directly on the question. For them, the condition for violation in the linear analysis is that the ratio exceed one half, and the black holes that meet it fill a region. An argument that a set has zero chance because it needs infinitely many separate conditions to hold at once fails for a set defined by a condition that fills a region, as this one does. Whether the matching set of starting conditions has a positive chance depends on how starting conditions map to the black holes they settle into, and that is not settled. Claims that blow-up is typical in the sense of chance, with a positive cosmological constant, must therefore be limited to kinds of conditions and black holes for which the ratio is known to be below one half, such as the spinning black holes in the linear analysis of Dias, Eperon, Reall and Santos.

In the sense of descriptions, the refining form is a strong form of the conjecture, with "typical" read as "outside a nowhere dense set". As a restatement of the conjecture the refining form is exact. An argument for it that sets out from "every nudge, however small, to the exact spinning black hole sets off the instability" assumes the statement it sets out to prove, and so, offered as a proof, it is circular.

All three restatements ask one question, in three senses of negligible: does mass inflation, the enclosed mass growing without limit at the Cauchy horizon, occur for typical starting conditions, at the middle standard of smoothness? The literature states as a conjecture that it does. In words: for typical starting conditions, for spinning black holes in empty space and for round, charged black holes holding a scalar field, the enclosed mass grows without limit at the Cauchy horizon, and the predicted world cannot be continued across it with square-integrable rates of change. So the equations cannot be written past it. This statement is the middle version of strong cosmic censorship, joined to mass inflation. What is proved of it is what the section on proofs set out. For round, charged black holes, the loosest version and the mechanism are proved (Luk and Oh, 2019; Dafermos, 2005). For spinning ones, unbroken continuation across a piece of the Cauchy horizon is proved (Dafermos and Luk, 2025), together with failure of continuation with bounded rates of change, under stated assumptions (Gurriaran, 2026, and Luk and Sbierski, 2026, both preprints, with Sbierski's theorem of 2026 for the step from the blow-up of the bending). I state the conjecture as the literature states it.

What Kind of Problem It Is

The literature has stated strong cosmic censorship as a conjecture since Penrose posed it, and its versions, its partial results and the mathematical obstacles in its way are many. It is a statement about the typical behavior of solutions of Einstein's equations inside black holes. The methods it calls for are the methods of mathematical relativity, the branch of mathematics

that studies Einstein's equations: the analysis of equations whose solutions move as waves and act back on themselves, the study of a black hole's tones, and the geometry of surfaces moving at the speed of light. No result shows, or even suggests, that any of its versions cannot be decided in the systems of mathematics in which mathematical relativity is done.

The books on which this book draws leave the unsolved problems of each field to that field, and they say nothing about strong cosmic censorship in particular. My book *The Six Cells of Mathematics* states the division plainly. Which particular statements cannot be decided in a given system of mathematics is found out by proof theory, the branch of mathematics that studies proofs, and the unsolved problems of each field remain that field's work. I follow that division: I report where the conjecture stands, and leave the settling of it to mathematical relativity.

What has been proved, and what the literature states as a conjecture, can both be named. Both concern how ripples behave at the Cauchy horizon of a spinning black hole under the full equations, with no symmetry assumed: the behavior that mass inflation shows in the round, charged models. The proofs reach that behavior under stated assumptions on the starting conditions, at the standard of bounded rates of change. The statement of that behavior for typical starting conditions, at the middle standard, is the middle version, which the literature states as a conjecture.

What the conjecture bears on is a matter apart from the argument of this book: whether general relativity, taken on its own terms, lets typical starting conditions produce Cauchy horizons that its equations can cross. I state that matter exactly, because it bears on whether a time machine of the kind physicists describe could form from typical starting conditions, with the kinds of matter for which the conjecture is posed.

What the Equations Allow

We now know what general relativity offers on paper. It offers geometry: shapes of space and time, some with time loops in them, from Gödel's world of 1949 to the wormholes of 1988. It offers theorems about those shapes: Hawking's, that a made time machine of his kind, growing from a slice that runs on without end, needs matter that breaks the weak energy condition; and that of Kay, Radzikowski and Wald, that a quantum field's well-behaved states stop being well behaved at the base set of such a horizon. General relativity offers one principle for worlds with loops, self-consistency. And it offers a conjecture, Penrose's, about whether typical starting conditions ever make an edge of prediction that can be crossed, with what is proved of it and what is still stated as a conjecture set out exactly. All of it is geometry: every loop, every horizon and every continuation past a horizon is a shape laid out in geometric time. Whether any of it describes a trip, a going back to an earlier occasion as it was, is a different question.

The difference between a shape laid out in geometric time and a trip back to an earlier occasion is what decides the book's question. A time machine on paper is a shape that obeys Einstein's equations, and the phrase "time travel" asks for something done. The next chapter answers what the phrase asks for.

Summary

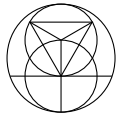
A time machine on paper is a shape of space and time that obeys Einstein's equations and holds time loops. Gödel's world of 1949 has a loop through every point and no Cauchy horizon at all. A time machine in the physicist's sense is a region where loops appear in a world that had none, bounded by a chronology horizon, which is a Cauchy horizon. A made time machine of Hawking's kind, growing from a starting slice that runs on without end, needs matter that breaks the weak energy condition. The

book's argument takes from this literature one principle only, self-consistency. Strong cosmic censorship, Penrose's conjecture, says that typical starting conditions do not produce Cauchy horizons that can be crossed; it protects the determinism of the equations, not a determinism of findings. Its strictest version is false for spinning black holes in empty space once two further steps are in place; its loosest version is proved for round, charged black holes; its middle version, which says that no continuation exists on which the equations can be written, is stated as a conjecture. For spinning black holes, the nearest results are proved under stated assumptions, at the standard of bounded rates of change, which asks more of a continuation than the middle standard does. With a positive cosmological constant, the question turns, in the linear analysis, on whether a ratio of rates lies below or above one half. Every result in this chapter describes shapes laid out in geometric time.

The chapter's statements are of several kinds, and each keeps its own standing. Most are results of the literature of mathematics and physics, each stated with its authors and carrying their warrant: Gödel's world (1949); the wormholes of Morris, Thorne and Yurtsever (1988); Hawking's theorem (1992); the theorem of Kay, Radzikowski and Wald (1997); the principle of self-consistency of Friedman and his coauthors (1990), the one result the next chapter uses; mass inflation (Poisson and Israel, 1990; Ori, 1991; Dafermos, 2003 and 2005); the proofs of Luk and Oh (2019), of Sbierski (2018, 2023 and 2026), of Luk and Sbierski (2016) and of Dafermos and Luk (2025); the preprints of 2025 and 2026, each named as a preprint; and the results of the linear analysis with a positive cosmological constant (Hintz and Vasy, 2017; Cardoso and his coauthors, 2018; Dias, Eperon, Reall and Santos, 2018). The three meanings of "typical" are restatements, in exact terms, of what the conjecture claims, and the rank of the failure set is a stated fact of mathematics, under a stated assumption. Two statements are conjectures, each reported as the literature states it: Hawking's chronology protection conjecture, and Penrose's strong cosmic censorship, whose middle version, joined to mass inflation, the literature

states as a conjecture. The one reading this book adds is that the uniqueness the conjecture protects is the determinism of the equations, which exists as structure, and not a determinism of findings.

What exactly would "going back" have to be, and can anything be it?



Chapter 16

Why Time Travel Names Nothing

Very Hard, or Nothing There?

Is time travel only very hard, the way a journey to another star is very hard? Or is there nothing there to do?

Picture a door with a small brass plate on it, and on the plate the words “Last Thursday”. The stories ask us to imagine that the door opens onto last Thursday itself: the same morning, the same rain on the same window, the same talk at breakfast, and you there among it all. An engineer who meets that door asks how to build it, and how much energy it would take. In this chapter I ask a question that comes before the engineer’s question: what would be on the other side of the door?

Ask three people what “going back to last Thursday” would mean, and you may get three answers. The first unrolls a map of space and time, the kind Chapter 14 described, on which days are laid out the way cities are laid out on a map of Europe, and she points to a spot on the map that lies earlier than the spot where she stands. Going back, for her, is reaching that earlier spot. The second holds up his notebook of findings, page after page in the order he wrote them, and he says that going back would mean the pages coming to stand in another order: a page written today coming before a page written last Thursday, or a page already written becoming unwritten. The third says that she means more than a spot on a map, and more than a notebook in a new order. Going back to last Thursday means

being at last Thursday: the very spot on the map, and the very day as it happened, with everything in it happening as it happened, and you present among it. She means that day itself, and not a copy of it or a day like it somewhere else.

Hold a smaller picture beside the door. A counter on a laboratory bench clicks each time an atom near it decays, and a paper tape records every click. Anyone may unroll the tape and read the record of what happened ten findings ago. The reading is real, and it happens now. The click that the tape records happened once, and the tape is what remains of that click.

I show that the third answer is what the phrase "time travel" means, that the first answer describes something else, and that the second and the third ask for what nothing satisfies. "Time travel", in the sense in which the phrase is used, names no concept: the phrase picks out no thing at all, and so there is nothing for an engineer to build. It is not that the engineering is hard.

The Chain

I set out the argument as a *chain*: a row of steps in which each step rests only on premises named in advance and on steps already made, so that a reader can check each link on its own and see the conclusion held by all the links together. The chain ends in one conclusion: the phrase "time travel", in the sense in which it is used, names no concept. Each step is marked as the kind of step it is, so that the reader can see which steps this chapter argues and which it takes from the books that prove them. Some steps are carried over from my other books, where they are proved or argued, and a step carried over stands on the argument given for it in the book it comes from; each time, I name the book. Some steps are definitions, which fix what a word will mean. The rest are argued here, and for each of those I say what it rests on.

Three words from the first chapter of this book carry the whole argument, so here they are again. A comparison is the act of finding whether two things are the same or different, in some respect. A comparer is anything that makes a comparison (an instrument, a detector, an interaction), and the word 'comparer' never means a mind. A finding is what one act of comparison yields: one outcome, laid down as a record. A traveler who carries a clock is a comparer, and each reading of the clock is a finding.

The chain takes very little from physics. From everything that physics has found about time machines, it uses one principle only, the rule of self-consistency for paths that loop back in time, which we meet below, and it uses that rule only in examining one of the three readings of the phrase. It uses no result about strong cosmic censorship, no rule about energy, no value of any constant of nature, and no account of how smooth or rough any horizon is. Strong cosmic censorship, which Chapter 15 described, is the conjecture that for typical starting conditions space and time cannot be continued in any well-behaved way past a Cauchy horizon, the edge beyond which what is given at the start no longer fixes what happens. So the chain's conclusion holds whether strong cosmic censorship is true or false, and in whichever formulation the conjecture is settled, however smooth that formulation requires a continuation past the horizon to be. Whatever lies beyond a Cauchy horizon, if anything lies beyond it, is more geometric time, as Chapter 14 showed. The other results of physics about time machines are checks on the conclusion, and we come to them near the end of the chapter.

The Four Premises

The first premise is the Incommensuration Theorem, which Chapter 3 set out. The theorem is stated over one act of comparison. That one act returns two verdicts: one on the content, what is compared, and one on the context, the setting in which it is compared; and each verdict is "same" or "different". The four ways of putting the two verdicts together are the four

of Chapter 2: continuity (the same content in the same setting, so that nothing jumps), discontinuity (different content in the same setting: a jump, a click), symmetry (the same content across different settings), and asymmetry (different content across different settings). I call each of the four a character of what a comparison finds.

The theorem says that any comparison, perfected at its foundation, is one of the two survivors of Chapter 3: it is continuous and asymmetric, smooth and going somewhere, or it is discontinuous and symmetric, the same everywhere and countable. In the theorem's own words: "Symmetry and continuity cannot both be applied, simultaneously and fundamentally, to any event or comparison; and neither can asymmetry and discontinuity. Any comparison, perfected at its foundation, is either continuously asymmetric or discontinuously symmetric". "Perfected at its foundation" means exact at the root of the comparison, with no exception at any scale. The theorem is proved from comparison alone, with no added assumption, in my book Incommensuration. A finite form of it is also stated and proved in the language of category theory, the mathematics of structure-preserving maps, in my book Category Theory and the Immanent Metaphysics. The Incommensuration Theorem is a structural fact about the nature of knowing. In the chain it is a theorem, carried over from Incommensuration, which proves it.

The two pairings the theorem forbids, then, are continuity with symmetry and asymmetry with discontinuity, and the theorem governs what one comparison finds when that finding is perfected. The smooth symmetries that the mathematics of a physical theory carries belong to the theory's *formalism*, its equations and the rules for working them, and not to what the theory finds. The theorem leaves them in place, as Chapter 14 showed in setting out the two levels of a theory, its mathematics and its findings.

The second premise is the three axioms. Chapter 4 named the three modalities they speak of: the immanent, which is process, what is happening here and now; the omniscient, which is structure, what holds the same for everyone, laid out as a pattern; and the transcendent, which is possibil-

ity. Incommensuration derives the axioms from the theorem together with the six intrinsics of comparison (sameness and difference, content and context, the subjective and the objective). The theorem comes first, proved from comparison alone, and the axioms are its internal structure, made explicit. In the words of *An Immanent Metaphysics*:

> **Axiom I.** The immanent is more fundamental than the omniscient and/or the transcendent. The omniscient and the transcendent are in conjugate relationship to one another.
>
> **Axiom II.** A class of the transcendent will precede an instance of the immanent. A class of the immanent will precede an instance of the omniscient. A class of the omniscient will precede an instance of the transcendent.
>
> **Axiom III.** The (classes/instances) of the immanent, omniscient, and transcendent are distinct, inseparable, and non-interchangeable.

Said plainly: process is more basic than structure or possibility, and structure and possibility each can be perfected only at the cost of the other. A range of possibility comes before an act, an act comes before its record, and a record comes before the next range of possibility. And the three are different from one another, none of them occurs without the other two, and none can stand in for another.

Two points about the axioms carry weight in the chain. The first point is that the priority Axiom I states is a priority of ground, and not an order in clock time: in the words of Incommensuration, it is "a claim about priority, not about sequence within process". The second point is that the word "precede" belongs to Axiom II, and Axiom II is the axiom that gives a direction. Of the cycle it describes, my book Measurement says that the direction "is fixed and cannot be reversed". In the chain, the axioms are quoted, and carried over from Incommensuration with their derivation from the theorem.

The third premise is a matching of the three times of Chapter 14 with the three modalities. When the subject is time, all three axioms apply, and I identify the three times as follows. *Geometric time* is the time laid out in the shape of space and time: the time a clock reads along each path, and the order of cause among events, both carried by the metric, the rule that gives distances and durations in the geometry. Geometric time is identified with the omniscient, and what is found of it is continuous and asymmetric. *Event time*, the order of findings for a comparer, each finding one outcome laid down as a record, is identified with the immanent, and what is found of it is discontinuous and symmetric. *Evolution time*, the variable along which quantum theory's description of a system changes smoothly, is identified with the transcendent, and it is continuous in the formalism. The one-way order of findings, the order of Axiom II, is identified with the direction of time. That direction is the direction of event time, and it is a different thing from the marking of future and past that the geometry carries at each place (physicists call that marking the time orientation of the metric). In the chain, the matching of the three times with the three modalities is an identification, carried over from The Unquantized Arena, which makes it, and which also shows that the three times are distinct, inseparable and not interchangeable.

The fourth premise is the difference between existing and being real, from Chapter 4, in the words of The Unquantized Arena: "To say that something exists is to say that it stands as structure: timeless, the same in every instance. To say that something is real is to say that it is given in an actual act of comparison: temporal, occurring now, with the future still open. A law exists; a measurement is real". My book The Basis of Physics applies it to time: "Geometric time is the dimension, and it exists; event time is the process, and it is real". In the chain, the difference between existing and being real is a definition, quoted.

No Unified Time

The first step after the premises concerns a *unified time*: "one time that is geometric time, event time and evolution time together, with all three perfected in one formalism, none of them external and none primitive". "None external" means that no one of the three times is left outside the theory; "none primitive" means that no one of the three times is taken as given without an account. The step says: "No unified time exists. A time found both as geometric time and as event time would carry at least one of the two forbidden pairings, and, where all four characters are asked of the same directions, both". The same directions, here, are the same directions of space and time: the case in which one and the same direction is asked to carry all four characters. In the chain, the step of no unified time is a result carried over from The Unquantized Arena, which proves it. In that book the step rests on the theorem and on the characters given to what is found of geometric time and of event time.

The argument for the step is short. What is found of geometric time is continuous and asymmetric, and what is found of event time is discontinuous and symmetric. A time found as both, with each character perfected in one comparison, would carry continuity together with symmetry, or discontinuity together with asymmetry, or both of those pairings where all four characters are asked of the same directions. Continuity with symmetry and discontinuity with asymmetry are the two pairings that the Incommensuration Theorem forbids. Evolution time does not change the verdict, because its continuity lies on the same side as the continuity of the geometry: what evolution time brings is continuity, the character that geometric time already brings, and so it adds no character that would change which pairings the time would carry.

The second sentence of the step is the one the chain uses: it does not need evolution time, and it does not need a single formalism. It needs only that one time be found, in one comparison, both as geometric time and as event time, each character perfected. The Unquantized Arena adds that "the forbidden pairing lies in the demand for one time, and not inside

any of the three". Each of the three times is in good order on its own terms. The fault appears only when one time is asked to be two of them at once, each perfected.

The One-Way Order of Findings

The next step concerns the counter's tape. The order in which a comparer's findings stand is the order of Axiom II. It is "a precedence and not a clock reading, and it cannot be reversed": no later finding comes to precede an earlier one, and no finding, once laid down as a record, is undone as a finding. A finding is discrete, one click, and its record is not undone; the smooth change of quantum theory's description is continuous, and in its mathematics it can be run in reverse. In the chain, this step is carried over from The Unquantized Arena and my book The Quantum Passage, which prove it.

The same axiom settles the question of signals sent into the past, which physicists call retrocausal signaling. A signal, as Chapter 12 used the word, is something that passes from one event to another and lets a record at the second event depend on an act at the first. As The Quantum Passage states it, retrocausal signaling would require reversing the one-way cycle of Axiom II, from structure to possibility, from possibility to act, from act to a new record of structure; and the axiom forbids that reversal.

The step is about findings. It is not about quantum states, and it is not about physical traces, the marks that a finding leaves in matter. A quantum state's smooth change can be run backward in its mathematics. A laboratory can also reverse some physical traces in part, as a spin echo does (the spin echo of Chapter 13, in which a spreading crowd of tiny magnets is made to gather again). Chapter 13 showed that neither of these acts on the order of findings: a reversal is itself a further act, and its own findings come after the findings it is said to undo.

What the Phrase Asks

The phrase "to travel to an earlier time" is said of a traveler, and a traveler, in the terms of this book, is a comparer departing from an event at which it has made findings. The phrase does not say which of the three times it means. Evolution time is not a candidate: it is the variable of a description that ranges over possibilities, not a time in which anything is found, and no one proposes to travel to an earlier value of a variable. That leaves geometric time, event time, or both.

The definition of the three readings is my own, made in this book. Let a traveler depart from an event, the departure, at which the traveler has made findings and laid down their records. "The traveler travels to an earlier time" has exactly three readings, the three answers of the people at the door.

The first reading is the geometric reading, the answer of the woman with the map. The traveler's path through space and time, its *worldline* (Chapter 15), continued into the future of the departure, reaches an event that lies earlier than the departure in geometric time. "Earlier" here means that the event lies in the departure's causal past and not in its causal future. Chapter 12 named these: at each event, the causal future is everything the event can affect (its forward light cone), and the causal past is everything that can affect the event (its backward light cone). The event reached may also be the departure itself, or an event that lies in both the causal past and the causal future of the departure, reached along a *closed timelike curve* (a time loop, in Chapter 15's plain name). A closed timelike curve is a path that a clock can follow and that returns to its own starting event, as Chapter 15 described.

The second reading is the event reading, the answer of the man with the notebook. The traveler's findings come to stand in an order other than the order in which they were made: a finding made after another comes to precede it, or a finding already laid down is undone as a finding.

The third reading is the joint reading, the answer of the third person at the door. The traveler arrives at an earlier *occasion*, a moment as it was given, with the findings made in it, and arrives at it as the occasion it was. The occasion is a location in geometric time, earlier than the departure, to which the traveler goes; and at the same time it is the occasion of event time that it was, with its findings given again and the traveler present among them. "As the occasion it was" means the same occasion exactly, and not a copy of it or a near likeness of it elsewhere.

The joint reading is the sense in which the phrase is used: going back to last Thursday means being at last Thursday, and not at some other event that resembles it.

The Geometric Reading

Under the geometric reading, what occurs, wherever the geometry allows it at all, is a later finding of the traveler, and not an arrival at an earlier occasion as the occasion it was. The chain argues this in four parts, and I give each part with its reason.

The first part: no path running forward from the departure reaches an event that lies in the departure's causal past and not in its causal future. The reason is a definition. An event reached from the departure by a path running forward in time, at or below the speed of light, lies, by definition, in the causal future of the departure. So an event outside the causal future is reached by no such path, whether or not the *spacetime* (space and time taken together, with their shape) contains loops. That no path running forward reaches an event outside the causal future is a fact of the geometry of relativity. It holds even for a geometry that is only continuous rather than smooth, when the paths are taken to be of the kind mathematicians call locally Lipschitz: paths that may turn corners, but never change direction infinitely fast.

The second part: the traveler's path may reach an event that lies in both the causal past and the causal future of the departure, which can happen only where closed loops pass through both events. Then the arrival is a later finding of the traveler, not an earlier one. A path from the departure to such an event, followed by a path from that event back to the departure, is a closed loop. Geometric time places such an event both before and after the departure, so geometric time does not place such an event earlier rather than later. The order of the traveler's findings is settled by precedence, the order of Axiom II, and not by the geometry. The traveler carries to the arrival the records it laid down at the departure, and its findings at the arrival are made from a range of possibilities that those records shaped. Under Axiom II, each of those findings therefore follows the findings at the departure, in the order of the traveler's findings.

The traveler's clock agrees that the findings at the arrival come after the findings at the departure. The time a clock reads along its own path (physicists call it proper time) increases along every stretch of a path a clock can follow, so the traveler's clock shows a later reading at the arrival than at the departure. The solution of Einstein's equations, a structure that exists, contains the traveler's path through the arrival event, and it contains no version of that event without that path. The event, as the structure lays it out, already has the traveler's arrival in it, and there is no version of it to reach in which the traveler is absent. The traveler's findings at the arrival are real, made in acts, and later than its findings at the departure.

The third part: the traveler's path may return to the departure itself along a closed loop, and then the arrival is a later finding of the traveler, and not the occasion of its departure given again. In the third part I bring in the chain's one principle from physics. In 1990 John Friedman, Michael Morris, Igor Novikov, Fernando Echeverria, Gunnar Klinkhammer, Kip Thorne and Ulvi Yurtsever set out the principle of self-consistency for spacetimes with closed loops. The principle allows only those histories in which the physical state at each event is single, one state and not two. A history in which the traveler returns to its departure therefore requires

that every physical trace of the records the traveler laid down along the loop, the state of its clock among them, be restored, on the way round, to the state it had when the traveler departed. Such a restoration acts on traces, not on findings. The restoration is itself a further act, and its findings follow, in the order of the traveler's findings, the findings it is said to undo. So the arrival at the departure is a later finding of the traveler. The findings the traveler made along the loop were made, and the occasion of its departure is not given again.

The fourth part: where a spacetime is continued beyond a Cauchy horizon, the first three parts apply to the region beyond as to any region. The reason is the result of Chapter 14: the region beyond a Cauchy horizon, if the spacetime goes on there at all, is more geometric time, and the first three parts use nothing but the order of what can reach what in the region where the traveler moves.

The geometric reading is argued in this book. It rests on the order of cause in the geometry of relativity, on the matching of the three times, on Axiom II, on the one-way order of findings, on the result that a reversal is a further act, on the result about what lies beyond a Cauchy horizon, and on the principle of self-consistency, carried over from the physics literature.

The third part says what happens to whatever goes around a closed loop and returns: it arrives at a later finding of its own, and its departure is not given again. It does not say that no physical system can go around such a loop. A system that keeps no record of any finding made on the way is not a traveler in the sense of the phrase, because it brings nothing from the later occasion to the earlier one.

The Event Reading and the Joint Reading

The event reading asks that a later finding of the traveler come to precede an earlier one, or that a finding already laid down be undone as a finding. The order of findings is the order of Axiom II, which runs one way and

cannot be reversed, and a finding's record is not undone as a finding. So the event reading asks for something that does not occur. A return to an earlier event, taken as the event it was, would require the order of findings to run backward, and that is the very reversal the one-way order excludes. The reversals a laboratory performs do not supply it, because they act on states and on traces, not on findings. The event reading is argued in this book; it rests on the one-way order of findings, on Axiom II, and on the result that a reversal is a further act.

The joint reading asks that one time be found, in one comparison, both as geometric time and as event time, each character perfected. By the step of no unified time, such a time would carry at least one of the pairings the theorem forbids, and so nothing satisfies the joint reading. The chain argues this in four steps.

Step one is the comparison. Under the joint reading, the traveler arrives at the earlier occasion. The arrival is an act of comparison: in it, the traveler finds where it is and what is happening there. The joint reading asks two things of what the traveler finds in that one act.

Step two is the two demands. First, the occasion is found as a location in geometric time, one definite place among the light cones, standing in the laid-out structure of space and time; in the terms of the fourth premise, the occasion exists. What is found of geometric time is continuous and asymmetric. Second, the occasion is found as the occasion of event time that it was, its findings given with the same definite outcomes, reading the same as before; in the terms of the fourth premise, the occasion is real, given in an actual act. What is found of event time is discontinuous and symmetric.

Step three is perfection. The joint reading asks for the same occasion exactly, and not for a copy of it or a near likeness of it. Picture the near misses. A copy of the kitchen, built tomorrow, with actors saying last Thursday's words, is a copy. A morning somewhere else with the same rain and the same talk is a likeness. A morning a hair's breadth away from last Thursday's, in space or in time, is another place. Each of these differs

from last Thursday somewhere, in its place or in its findings, however slightly, and the third person at the door refuses every one of them. So the demand allows no difference at all, at any scale. In the Incommensuration Theorem, "perfected at its foundation" means exact at the root of the comparison, with no exception at any scale. The joint reading's demand, which allows no difference at any scale, is a demand for perfection in that sense. A demand for exact sameness in both respects is a demand that each character hold exactly in the act of arrival: the continuity and asymmetry of the place, and the discontinuity and symmetry of the occasion. For the four characters to be perfected at the foundation of one comparison is for each character to hold exactly in the act of arrival. Anything less than exact sameness in either respect gives an occasion that is not the one the phrase names.

Step four is the verdict. Steps one to three ask one time to be found, in one comparison, both as geometric time and as event time, each character perfected. By the second sentence of the step of no unified time, such a time would carry at least one of the two pairings the theorem forbids. So no arrival satisfies the joint reading. The joint reading is argued in this book. It rests on the difference between existing and being real, on the matching of the three times, on the definition of the three readings, on the step of no unified time, and on the theorem.

The Conclusion

In this book, a phrase *names no concept* when every account of what would satisfy it either describes something other than what the phrase is used to mean, or asks for a set of conditions, all at once, that nothing satisfies. "The largest prime number" names no number in the second way. It asks for a prime that no prime exceeds, and every prime is exceeded by a larger prime. "Time travel" fails in both ways, reading by reading: its geometric reading describes something else, and its event and joint readings ask for what nothing satisfies.

The conclusion of the chain is that "time travel", in the sense in which the phrase is used, names no concept. Its geometric reading, where the geometry allows it, describes a later finding of the traveler, and not an arrival at an earlier occasion as it was. Its event reading asks for the one reversal that the order of findings excludes. Its joint reading, which is the sense in which the phrase is used, asks that one time be found as both geometric time and event time, each character perfected, and the step of no unified time excludes that. The conclusion is argued in this book, from the verdicts on the three readings.

The mistake the phrase makes is a *category mistake*, in the plain sense of the term: it takes something of one kind and treats it as though it were of another kind. The use of the phrase rests on a confusion between two of the three times, and the phrase names no physical process. In geometric time, time is laid out as structure. It makes sense to ask what the geometry is like at an earlier location in the order of cause, in the same way that it makes sense to ask what lies at another place in space. The solution of Einstein's equations gives the answer, and the region asked about exists: it stands as structure, and it can be described from any vantage. In event time, time is the order of findings. A finding is given in an act of comparison. The finding is real, happening when the finding happens, and what remains of the finding later is the record of the finding. So it makes no sense to ask how to reach a finding that came earlier, because a finding is not a location. One can read the record of what happened ten findings ago. One cannot be present at that finding again.

The time travel enthusiast takes a feature of geometric time, that every event has a location to which a question can be put, and gives it to event time, expecting that an earlier finding can be reached the way a location is reached. The phrase, used this way, takes an occasion of event time, which is real, given once in an act of comparison, and treats it as though it existed, standing as structure that can be visited. The move fails for two reasons. The two times are distinct, and neither can stand in for the other; and one time found with the characters of both, each perfected, would carry at

least one of the pairings that the Incommensuration Theorem forbids. Neither of the two times is faulty. The fault is the demand that one time be both.

The region of the geometry in which last Thursday lies exists, as structure. The findings of last Thursday were real when they were given, and what is real of them now is their records. You cannot travel to last Thursday, in the sense the phrase intends, because last Thursday, as the occasion it was, is not a place. Neither the region nor the records is an occasion that one can attend a second time. That the findings of last Thursday are over is not a limitation of technology. Their being over is what it means for a finding to have been made.

The conclusion has the compulsion of the steps it rests on. The Incommensuration Theorem is a theorem, proved from comparison alone, and it compels as a theorem does: it holds of every comparison, with no condition. The step of no unified time applies the theorem to the characters that The Unquantized Arena gives to what is found of geometric time and of event time, and the third premise binds those two times to their modalities. The conclusion therefore holds as far as those characters and that identification hold: if geometric time and event time were not two objects of different kinds, with the characters stated, the step of no unified time would have nothing to forbid. The conclusion's condition, that geometric time and event time have the characters stated and that the identification holds, is a condition on applying the theorem to the two times; the theorem itself holds with no condition. The Unquantized Arena gives the reasons the two times are of different kinds: "A time function is a smooth map from a manifold to the real numbers. Event time is a discrete ordered set of records". Said plainly: a time function, the form geometric time takes in the mathematics, is a smooth rule that gives every event of space and time a number on the unbroken number line. Event time is a row of separate records, each in its place in the order. Geometric time and event time are two objects of different kinds, with the characters stated, and that difference is what gives the step of no unified time something to forbid.

What the Chain Uses

Here, in one place, is the whole of what the chain takes from physics, so that the reader can see how little the conclusion asks of it. The chain uses general relativity, Einstein's theory of gravity as the shape of space and time, for three things only: for the definition of geometric time, for two facts of its geometry, and for one principle reported from the physics literature. The first fact is that the time a clock reads increases along every stretch of a path a clock can follow. The second fact is that a path from an event into that event's causal past, followed by a path from there back to the event, is a closed loop. The principle is the principle of self-consistency of Friedman and his co-authors, used in the third part of the geometric reading. The chain also uses the result of Chapter 14 about what lies beyond a Cauchy horizon, which rests on none of the results named below.

That list is complete. Everything else that physics has found about time machines lies outside the chain: Stephen Hawking's theorem of 1992 on whether the horizon of a time machine built in a limited region can form, the result of Bernard Kay, Marek Radzikowski and Robert Wald of 1997 on quantum fields at such a horizon, every rule about energy, and the strong cosmic censorship conjecture in all its forms, with every result that bears on it. The chain does not use the information deficit, which Chapter 14 placed at the edge of what an observer can see of the universe (the particle horizon) and at no horizon of a black hole. So the information deficit gives no account of any Cauchy horizon. It uses no value of any constant of nature.

Three Questions a Reader May Ask

Is the conclusion true by definition? The conclusion follows from a theorem, and is not true by definition. The definitions of the three times are definitions of three mathematical objects of different kinds, and the three readings are definitions of what the phrase could mean. The step that ex-

cludes the joint reading is the step of no unified time, which applies the theorem to the characters of what is found of geometric time and of event time. The theorem is proved from comparison alone, and the characters are stated with their grounds in The Unquantized Arena; neither is a definition of this book.

In 1949 Kurt Gödel found solutions of Einstein's equations with closed loops in time through every event, and argued from them that the passing of time is not objective. Does that show that time is a place after all? Gödel's solution exists, in the sense of the fourth premise: it is a structure, laid out whole. Gödel's solution shows a fact about geometric time in Gödel's geometry. It does not show that event time is a part of geometric time. The Basis of Physics puts the point in one sentence: "The block universe takes geometric time for the whole, and drops event time". (The block universe is the picture of all of space and time as one finished block, with every moment laid out in it at once.) The geometric reading states what a traveler does on such a loop.

Some ways of writing quantum theory are symmetric in time, and some ways of reading it allow influences from the future. Do they reopen the question? A way of writing quantum theory is a formalism, and its symmetries belong to the formalism, which exists as structure. The order of findings is not a feature of any formalism. The Quantum Passage states the limit on signals: a later choice of test is a new stance, which follows the finding and cannot reach back. A stance is the standpoint from which a comparison is made. Choosing tomorrow which test to make on a record is taking a new standpoint, and that choice comes after the finding it would test, as possibility follows the record in the cycle of Axiom II. So the choice can shape what is found next, and it cannot reach back to what was already found.

Does the argument say that the past is unreal? It says that a past occasion was real, given in an act, and that it is not given again. Its records are real now, each in the acts that read it, and the geometry in which the occasion is placed exists. Nothing in the argument removes either the records or the geometry.

Two Arrows, One Direction

The chain rests on the one-way order of findings, and a reader who knows the usual account of the arrow of time may ask whether that order is only an accident of how the universe began. The answer comes from what Chapter 13 explained: the arrow of time, the one-way direction of time, has two forms, and the chain needs only one of them. The usual explanation of the arrow is drawn from heat. The *entropy* of a closed-off system, the measure of how many ways its parts can be arranged while the whole looks the same, increases, and it increases because the universe began in a state of low entropy (the "Past Hypothesis"). On that account the arrow is an accident of the start: had the start been different, the arrow could have pointed the other way.

I find two arrows, and not one. The first arrow is structural, and my book Incommensuration states what the structural arrow is: "The non-entropic arrow of time is the structural precedence of Axiom II applied to the physical domain. The direction of time is the direction in which comparison proceeds". This arrow is present in every single comparison: the range of what is possible comes first, then the act, then the record, and the record shapes the next range. It needs no beginning of low entropy.

The second arrow is the entropic arrow, and Measurement states it: "It is a statistical arrow: it is found in systems of very many parts, and it needs a beginning of low entropy from which the growth proceeds". The entropic arrow is real, and it points the same way as the structural arrow wherever both are found. The reason is that a record is itself a change that cannot be undone, and making it raises the entropy of the instrument that holds it, "So each act of recording joins the two arrows in one event".

The structural arrow is therefore the root of the direction of time, and the entropic arrow is how that direction shows itself at the scale of eggs and engines. Why the early universe had low entropy is a question about the start, which the structural arrow leaves to the account of the start. The

Appendix gives, in the words of The Basis of Physics, what the structural reading says of that beginning. I need no more than the structural arrow, and I claim no more.

The Four Barriers

The whole account sets four barriers in the way of the phrase, and they are of different kinds. The first three are structural: they reach the phrase itself, and the chain rests on them. The fourth is physical: it reaches the designs for a time machine that physics has examined, and it is a check on the conclusion, not a premise of the chain.

The first barrier is that there is no unified time: no one time is geometric time, event time and evolution time together, with geometric time and event time each perfected as found and evolution time perfected in the formalism. The barrier rests on the Incommensuration Theorem and on the characters of what is found of geometric time and of event time, and the argument of the joint reading applies it to the phrase. The second barrier is that the order of findings runs one way, in the order of Axiom II. Axiom II is derived from the Incommensuration Theorem, and the same axiom forbids sending signals into the past. The argument of the event reading applies the barrier to a trip backward in event time. The barrier belongs to Axiom II, and not to Axiom I, whose priority of the immanent is a priority of ground and not a sequence in time.

The third barrier is that the record is not undone. In the words of The Basis of Physics: “Symmetry of the laws guarantees lawful reversal. It does not guarantee faithful reversal”. Chapter 13 drew the line between the two: a *lawful reversal* is a reversed motion that obeys the laws, and a *faithful reversal* is one that restores the original state in every respect. A reversal that a laboratory performs, a spin echo for example, runs part of a change backward where the drift is held in only a few of the ways the sys-

tem can change, and it leaves the records of the laboratory itself in place. Chapter 13 set out why no reversal, however lawful, returns a comparer to a finding it has already made.

The fourth barrier is that physics guards its own Cauchy horizons. This barrier is made of results reported from the literature of general relativity, and of one conjecture of that literature, all of which Chapter 15 set out. Stephen Hawking, in 1992, described a time machine built in a limited region of space. In his description, such a machine is a spacetime whose Cauchy horizon is compactly generated: its edge is formed by light paths that, followed back into the past, stay within a bounded region. He proved that such a horizon cannot form from a slice of space at one moment that runs on without end, if the weak energy condition holds. The weak energy condition is the rule that the energy any observer measures is never negative. His own summary of the paper states the result as a requirement: on the horizon, the averaged form of that energy rule must fail. Bernard Kay, Marek Radzikowski and Robert Wald proved a result in 1997 about the free scalar field of Chapter 15 (physicists also call it the Klein-Gordon field). They proved that every well-behaved state of the field (every Hadamard state, in Chapter 15's gloss) breaks down at every point of a core part of such a horizon (its base set). There the field's measure of how its values at two nearby points go together, from which its energy is computed, is singular.

Chapter 15 reported what has been proved of strong cosmic censorship, and under which assumptions, and the form in which the literature states it as a conjecture. The conjecture is posed for space empty of matter and for kinds of matter that obey rules about energy. Each formulation fixes how smooth a continuation past the horizon must be to count. If the conjecture holds in a given formulation, then, within those kinds of matter, typical starting conditions do not produce Cauchy horizons that can be crossed in that formulation, the horizons of time machines among them. And if a Cauchy horizon could be crossed, what lies beyond it

would be geometry, and not a return. None of these results is a premise of the chain. They are checks: they show that physics, in its own terms, does not supply the door that the phrase imagines.

The Reach of the Argument

The argument concerns the phrase "travel to an earlier time" in the sense its users intend: arriving at an earlier occasion as the occasion it was. Its conclusion is that the phrase, in that sense, names no concept. The conclusion says that nothing would satisfy the phrase in its intended reading, and the geometric reading states what a spacetime with closed loops in time does allow. The argument is about the phrase, and not about engineering: it does not say that a time machine is too costly, needs too much energy, or lies beyond present skill.

The argument leaves every reversal that physics performs in good order: spin echoes, the fact that the laws of mechanics and of quantum theory read the same with time reversed, and the reversibility of evolution time in the mathematics of quantum theory. Each of them is a lawful reversal carried out within the one-way order of findings; none of them reverses that order. Chapter 13 told Josef Loschmidt's objection of 1876 to Ludwig Boltzmann (reverse every velocity in a gas, and the laws run its history backward); the objection turns on this same difference between lawful and faithful reversal.

So the conclusion can be said in the plainest words. You cannot go back. The reason is not that the engineering is hard, not that the energy needed is extreme, and not that exotic matter does not exist. You cannot go back because "back", said of time, borrows the direction of a location in geometric time and applies it to the order of findings, where the reverse direction is not there to be had. Geometric time is laid out as structure, and it exists. Event time is the order of acts, and it is real. One time found as both is the time that the step of no unified time, from the Incommensuration Theorem, excludes.

The question that matters is "why did we ever think the phrase meant something?", and not "how do we build a time machine?" The answer is that we confused the symmetry of the laws with the course of the world. The laws of general relativity and of quantum theory's smooth change read the same with time reversed. Geometric time, as found, is not symmetric. And the world happens. Its findings follow one another in one order, and no equation can reverse that order, because every equation is written, solved and checked within event time, in an order of findings that the equation does not hold and cannot run backward. The door with the brass plate does not open.

What We Now Know

We now know that "time travel", in the sense in which the phrase is used, names no concept. Going back to last Thursday means being at last Thursday, the very place and the very occasion; that asks one time to be found as geometric time and as event time at once, each perfected; and the Incommensuration Theorem excludes exactly that demand. Of the other two readings, the geometric reading describes a later finding of the traveler, and the event reading asks the one-way order of findings to run the other way.

What we now know matters for the book's question, what cannot be, and why. Chapter 12 found that a signal faster than light asks for a comparison beyond the edge of comparison, and so names nothing. Chapter 13 found that the record follows the act it records, always, so that no reversal is faithful. I make the second exclusion in time, after Chapter 13's exclusion of a faithful reversal. The trip that would have to undo the order in which a record follows the act it records is a phrase with no thing behind it, and not a trip that is far away. In the next part of the book I ask how far physics itself can reach.

Summary

The chain runs from four premises: the Incommensuration Theorem, a structural fact about the nature of knowing; the three axioms, with the one-way order of Axiom II; the identification of geometric time with structure, of event time with process, and of evolution time with possibility; and the difference between what exists and what is real. It carries over two results: no unified time exists, and the order of findings cannot be reversed. The sentence of the first result that the chain uses does not need evolution time, and it does not need a single formalism. It needs only one time found, in one comparison, both as geometric time and as event time, each character perfected. Retrocausal signaling would require reversing the one-way cycle of Axiom II, which the axiom forbids; a later choice of test is a new stance, and it cannot reach back.

"Going back" has exactly three readings. The geometric reading describes a later finding of the traveler; the event reading asks for a reversal that Axiom II forbids; the joint reading, the sense in which the phrase is used, asks one time to be both geometric time and event time, each character perfected, and nothing satisfies it. So the phrase "time travel", in the sense in which the phrase is used, names no concept, and the reason is not that the engineering is hard. The conclusion has the compulsion of the steps it rests on. The theorem compels as a theorem does, with no condition. The conclusion holds as far as the characters of the two times and their identification with the modalities hold, so that the condition lies on the application of the theorem to the two times, and not on the theorem. The conclusion is about the phrase as its users intend it. The conclusion leaves in place every reversal that physics performs, the fact that past occasions were real, and the geometry in which the past is placed.

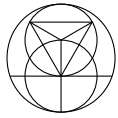
Every step of the chain is marked with the kind of step it is, and a step carried over from another of my books stands on the argument given for it there. Carried over are the theorem (from Incommensuration), the three axioms (quoted, and derived from the theorem in Incommensuration), the identification of the three times, the definition of existing and being

real, the steps of no unified time and of the one-way order of findings (from The Unquantized Arena, with The Quantum Passage for the second), and the difference between lawful and faithful reversal (from The Basis of Physics). Argued in this book are the result about what lies beyond a Cauchy horizon, the result that a reversal is itself a further act, the definition of the three readings, the verdicts on the three readings, and the conclusion. The principle of self-consistency is reported from the physics literature, with its authors; the strong cosmic censorship conjecture is a conjecture of that literature, reported as the literature states it, and not a premise of the chain; and the exclusion of time travel is a structural exclusion, drawn from the structure of comparison itself.

If the two great theories each describe one kind of time, why can they not be joined into one theory?

PART FIVE

THE EDGES OF PHYSICS



Chapter 17

Two Theories That Cannot Be One

A Closed Room

Physicists have tried for more than ninety years to make gravity and the quantum into one theory. Why has nobody managed it? Is the problem only very hard, waiting for a cleverer person or a better idea? The answer of this chapter is that the request, in its strong form, asks for one description that is perfect in the two ways that, as Chapter 3 showed, no description can be perfect at once. The answer is also that the two theories are already joined, in the only way two such theories can be.

Begin where general relativity began, in a closed room. In 1907 Albert Einstein imagined a person shut in a room with no windows. The floor presses up against the person's feet. By no experiment made at a single point can the person tell whether the room stands on the surface of a planet or is being driven through empty space by a rocket beneath it. Now let the room fall freely. The person feels nothing at all, and at a single point cannot tell whether the room is falling toward a mass or drifting far from every mass, with no force acting on it. At a point, standing on a planet and being pushed by a rocket are the same for every experiment that can be made there, and so are falling under gravity and moving with no force at all.

The *principle of equivalence* is that, at a single point, the acceleration of gravity cannot be told apart from acceleration of any other kind. It is the founding physical idea of general relativity, and it is what allowed gravity to be read as the geometry of space and time rather than as a force acting within them. The closed room has more to teach than it first seems to, and I come back to the room at the end of the chapter.

Two Theories Where Physics Wants One

For about a century physics has had two basic theories where it wants one. General relativity, complete in its essentials in Einstein's work of 1915, describes gravitation as the geometry of space and time: a smooth spread of events, with no gaps, whose curving is fixed by where mass and energy are, and whose curving in turn fixes how mass and energy move. I call space and time together the *arena*: the setting of every measurement, in which every length is a length and every duration is a duration. In general relativity, the shape of the arena is gravity.

Quantum theory, complete in its essentials a decade later, describes matter and the other forces as states among possibilities. Those states change by rules under which certain things stay exactly the same, and when they are measured they give separate, countable outcomes, each with definite odds.

Each theory has been confirmed so precisely that there is no serious doubt it is correct within its range. Each predicts things the other cannot address. And where their ranges are forced to overlap, they contradict one another. Three such places are usually named: the Planck scale, far below anything yet measured, where both theories claim the same object; the singularities at the centers of black holes, where the curving that general relativity computes grows without limit; and the earliest moments of the universe.

The standard response has been to seek a *unification*: a single theory of which both would be limiting cases, each following from it when some quantity is taken very large or very small, as Newton's mechanics follows from relativity when speeds are small. String theory, loop quantum gravity, causal set theory, twistor theory, asymptotic safety, emergent gravity: the list is long, and the literature of each is deep. Twistor theory is Roger Penrose's program of rebuilding space and time from the paths of light; emergent gravity is the family of programs in which gravity arises from something more basic; the other four are described later in this chapter. Roger Penrose wrote in 2004 that joining the two theories will require a revolution in ideas comparable to the revolutions that produced them, and that neither theory can be regarded as more basic than the other. Chris Isham set out in 1995, with great care, what any theory of quantum gravity would have to do to be a theory at all. Carlo Rovelli (in 2004) and Lee Smolin (in 2001), each committed to a program, have both acknowledged in writing how far the programs are from what they were meant to reach. I name these four so that the reader hears how the search stands in the words of people who work inside it.

The usual framing of this situation is that unification is hard. I claim something different. Unification in the strong sense is not possible, and the reason is a fact about comparison rather than a fact about physics; physics, being measurement, cannot escape it. The claim is among the strongest I make about physics, and a claim of impossibility owes the reader an exact statement of what is being called impossible, so the steps below make that statement with care.

Two Tools and Three Scales

The steps use only what earlier chapters built. Every comparison has six parts it cannot do without: sameness and difference, content and context, the subjective and the objective (Chapter 1). Sameness and difference, applied to content (what a comparison finds) and to context (the setting

that makes one finding distinct from another), give four compounds (Chapter 2). Continuity is sameness of content within sameness of context: hold the setting fixed, and nothing jumps. Discontinuity is difference of content within sameness of context: a jump, a click, a step. Symmetry is sameness of content across difference of context: change the setting, and the content stays the same. Asymmetry is difference of content across difference of context: change the setting, and the content changes with it.

The Incommensuration Theorem, which Chapter 3 set out, is a structural fact about the nature of knowing: symmetry and continuity cannot both be applied, at once and absolutely, to one comparison, and neither can asymmetry and discontinuity. Two pairings survive: continuity with asymmetry, smooth and going somewhere; and symmetry with discontinuity, the same everywhere and countable. The words "at once and absolutely" are part of the theorem. In any finite domain, symmetry and continuity occur together to a finite degree; what they cannot do is both be perfected at once.

Chapter 2 also showed that every law is a law of symmetry or a law of continuity, and Chapter 4 named the three modalities. The two kinds of law line up with the modalities in a fixed way. This first line-up matches tools, the two kinds of law, with the modalities, and it holds in what I call the focus of the two laws, the frame in which the two kinds of law are what is being considered. Symmetry is the tool of the omniscient, the modality of structure, because sameness across different settings is what a shared, repeatable pattern is: what any observer, anywhere, will find again. Continuity is the tool of the transcendent, the modality of possibility, because sameness within one setting is what the connected, unbroken whole is: the field of possibility from which a thing does not depart while it remains itself. And interaction belongs to the immanent, the modality of process, because interaction is the here and now in which content meets context at all.

There is a second line-up, of scales rather than tools, and the argument turns on it. One rule tells how two line-ups sit together. Which modality a thing is matched with depends on what one is considering, and is no label

fixed to the thing; Chapter 18 states this rule in full. Chapter 14 matched the three times with the modalities, in the focus of time; I now match tools, then scales, then kinds of acceleration, and each line-up says what it is matching. So the tool a theory uses and the domain it describes can be matched with different modalities, and the next sections show that they must be.

The second line-up matches scales, the domains that theories describe, with the modalities. The very large, which physicists call the *macroscopic*, is the omniscient domain: the layout of the whole, the pattern seen from outside, the universe as a fact that anyone can share. The very small, the *microscopic*, is the transcendent domain: the region of what has not yet shown itself, of possibility and probability, of the single event that is only itself and exceeds any one pattern laid over it. Between them is the *mesoscopic*, our own scale, the scale at which every measurement is made (in this book the word always names the observer's own scale); it is the immanent domain, the domain of interaction.

Each Theory Reaches Across

Now look at how the two theories are actually built. General relativity is built on continuity. Its basic object is a smooth spread of events, a geometry with no gaps, no jumps and no separate pieces. Its field equations, Einstein's rules connecting the curving of the arena to the matter and energy in it, are rules about how things change smoothly from point to point, and they need smoothness for their solutions to exist. Even the principle of equivalence, the lesson of the closed room, is a statement of continuity: at every point, falling freely and moving with no force are one, so the geometry joins the physics at each place without a seam. Continuity is built into what general relativity finds. The theory also stays the same under every smooth relabeling of places and times (physicists say, under every change of coordinates). But that sameness belongs to its *form-*

alism, the mathematics in which the theory is written, and it is not something the theory finds in the world. And continuity is the tool of the transcendent.

Yet general relativity describes the very large: the overall shape of space and time, the motions of planets and galaxies, the curving of the whole universe, the horizons of black holes, the expansion of the whole. The very large is the omniscient domain.

Quantum theory is built on symmetry. What it finds stays exactly the same across settings. Swap two electrons, and nothing that can be found changes. States are sorted by whole-number labels, which physicists call quantum numbers. Charges fix which interactions there are, and each charge is conserved without exception wherever the symmetry behind it is exact. Conservation follows from symmetry, by the theorem Emmy Noether proved in 1918, and conservation is the backbone of every quantum calculation. The smooth change of the quantum wave lives in the theory's formalism, in a continuous space of possible states. Symmetry is built into what quantum theory finds. And symmetry is the tool of the omniscient.

Yet quantum theory describes the very small: the changing possibilities of what has not yet shown itself, the superposition of possible states (their combination before any one of them is found), the weights that set the odds of each outcome, the single decay, the single click, the event that is only itself. The very small is the transcendent domain.

Each theory's building principle belongs to the opposite modality from the domain it describes. General relativity uses a tool of the transcendent to describe a domain of the omniscient. Quantum theory uses a tool of the omniscient to describe a domain of the transcendent. I call this the *cross-modal reach*: each theory reaches, with its tool, across into a domain of the other modality. By a theory's *tool* I mean the character it perfects in what it finds. Its reach into the other domain is dense and never exhaustive: as close as one likes, everywhere, and never the whole.

One might have expected the domain of pattern to be described with the tool of pattern, and the domain of possibility with the tool of possibility. Physics developed the other way round, both times, by different people for different reasons, with no thought of modalities at all. The pattern is clean enough that I do not take it for an accident, and the next step gives its reason.

Why the Reach Is Required

Chapter 3 settled the ten possible pairings of the four compounds by three considerations, and the first was that applying a thing to itself says nothing new. Symmetry of symmetry is symmetry. Continuity of continuity is continuity. The duplicate says nothing the single term did not say, and in the same way a tool of one modality applied to a domain of the same modality says nothing new, at the scale of a whole theory.

Suppose a theory used symmetry to describe a domain whose character is symmetry. It would be a structural description of structure: a map of a map, a pattern of a pattern. It would say of the domain only what the domain already is, because the tool and the domain would be one concept applied twice. Suppose instead a theory used continuity on a domain whose character is continuity. It would be connectedness describing connectedness, participation in participation, process of process, and it too would say nothing.

So a theory that tells us anything about its domain, anything the domain does not already say of itself, must take its building principle from a different modality than its domain. The tool must reach across. That requirement is neither a preference nor a habit of history; it is what it takes for a description to describe at all.

The same point can be reached from the axioms. The first axiom says that the omniscient and the transcendent are in conjugate relationship to one another. Two things are *conjugate*, as Chapter 3 put it, when each can be perfected only at the expense of the other, and neither comes before the

other in definition. A theory that perfected both members of a conjugate pair in what it finds would perfect two conjugates at once. Said in the names of the modalities, that is the exclusion this chapter states below: no single theory finds symmetry and continuity both perfected. Because neither comes first, each is reached by way of the other. To learn the shape of the omniscient, one reaches into it with the instrument of the transcendent; to learn the shape of the transcendent, one reaches into it with the instrument of the omniscient. The two theories cross, and they must: crossing is how conjugates give an account of one another.

A third route starts from the theorem itself. Symmetry and continuity cannot be joined absolutely, so no single theory can reach both domains by finding both characters perfected. Each theory perfects one, and the domain it describes informatively is the domain of the other. In each case the character of the tool and the character of the domain differ, for the reason already given: if they were the same, the theory would be empty.

The line-up of the two laws stands as it was. Symmetry is still the tool of the omniscient, and continuity still the tool of the transcendent; those assignments hold at the level of the six parts of comparison, where they are made. The reach adds a finer point: a theory's tool need not share the modality of its domain, and in fact must not, if the theory is to say anything.

The Two Survivors, Made Physical

So far only symmetry and continuity have been used, one at a time. The theorem is about pairs, and the theories, looked at closely, are pairs too.

General relativity is continuity with asymmetry. The building principle of general relativity is continuity: the smooth geometry, the unbroken field, the join of falling freely to moving with no force, made without a seam. What that principle produces, in every solution, is asymmetry, difference from place to place. The curving breaks the evenness of flat space, so that no two regions of a curved geometry are alike: mass

here is not mass there, and the geometry says so. The light cones of Chapter 12 (at each event, the forward cone of what it can affect and the backward cone of what can affect it) are directed, with a forward half and a backward half; the horizons of black holes are one-way; collapse cannot be undone. Continuity and asymmetry are both present in every feature of the theory, and the continuity produces the asymmetry, by way of the field equations. General relativity is the first survivor of the theorem, made physical.

Quantum theory is symmetry with discontinuity. Its building principle is symmetry: sameness under the smooth changes of its wave, under the exchange of identical particles, and under the transformations that define the forces. What that principle produces, whenever the theory touches the world, is discontinuity, a difference within one setting. Energy comes in levels; a measurement gives one outcome of many; what is found comes in separate steps, whatever account is given of how the outcome is reached. Symmetry and discontinuity are both present in every feature of the theory, and the symmetry produces the steps, by way of quantization: as Chapter 2 put it, the separate steps are the shadow of the symmetry. Quantum theory is the second survivor, made physical.

A reader who knows some physics will ask a fair question here. General relativity is full of symmetry: it stays the same under every smooth relabeling of places and times, and its physics at each point is the same for every observer in steady motion. Quantum theory is full of continuity: its wave changes smoothly, its space of states has no gaps, and the weights of the outcomes vary without a break. If each theory has both, how can either be read as one survivor?

The answer is that a physical theory has two levels: its formalism, which exists, and its findings, which are real. In the words of Chapter 4, to exist is to stand as structure, the same for every comparer, and to be real is to be in an actual act of comparison. The theorem constrains which character a theory perfects in what it finds; the other character, exact in its own right, lives in the formalism. General relativity finds a smooth geometry, perfect without exception, and writes it in mathematics whose symmetry

is exact and belongs to the description, not to the world found. Quantum theory finds exact labels, charges and quantum numbers, perfect without exception, and writes them in mathematics whose smooth change is exact at its own level. Every finding is one definite outcome, whatever account is given of how it is reached. Each theory perfects one compound in what it finds and the other in its formalism, and that split, at two levels, is exactly what the theorem permits.

A smooth symmetry shows the same split in miniature. Turn an experiment through any angle, and its laws stay the same. The sameness is found, as a conserved amount (for turning, the amount of turning motion a system carries). The smoothness, the fact that the angle can be any angle at all, lives in the formalism, as the smooth dial of the symmetry.

Chapter 3 found the same two survivors inside quantum theory alone: the wave as continuity with asymmetry, the particle as symmetry with discontinuity, related by the measurement that shows them. Chapter 3's reading, of the wave and the particle inside quantum theory, and this chapter's reading, of each theory as one survivor, agree. The wave side is the continuity that quantum theory carries in its formalism; the particle side is the symmetry that it perfects in what it finds. Niels Bohr's complementarity, in his paper of 1928, was the recognition that the two descriptions cannot both be given at once, and that each captures what the other must leave out. I read Bohr's recognition as a sighting of the general result in one place: the two survivors, seen from inside one theory. Between the two theories the same complementarity appears again, at the scale of the whole of physics. General relativity and quantum theory are mirror images across space, time and possibility. General relativity works out space and time, and leaves out possibility; quantum theory works out possibility and how it changes in time, and leaves out the actual.

What One Theory Would Have to Be

A unified theory of quantum gravity, in the sense the phrase usually carries, would be one theory, more basic than either, with general relativity and quantum theory both as its limiting cases. To give back general relativity in its limit, it would need to find a continuous geometry, with continuity perfected in what it finds. To give back quantum theory in its limit, it would need to find exact unchanging labels, with symmetry perfected in what it finds. It would need to perfect both at once, in what one and the same comparison finds.

The two surviving pairings are exact duals: each is defined by lacking what the other is defined by. Continuity with asymmetry is the pairing without symmetry; symmetry with discontinuity is the pairing without continuity. To find both, perfected, in one theory is to join symmetry with continuity absolutely, and that is the pairing the theorem forbids. Chapter 3 gave the reason in a few lines: a content that stays the same whether or not its context changes does not depend on its context at all, and then the line between content and context, which every comparison needs, has nothing to hold it apart.

My claim is that the long failure to unify the two theories is a structural exclusion. It is the same exclusion that Chapter 5 found in determinism's demand for a perfect causal law, and that Chapter 10 found in the demand, at the heart of Bell's theorem, for lawfulness and locality together: the Incommensuration Theorem, met at the largest scale on which physics works. So the failure is no failure of mathematical cleverness, of physical insight, or of guidance from experiment.

The claim is a claim of impossibility, and a claim of impossibility differs in kind from a claim of difficulty. A claim of difficulty says the problem is hard and might be solved by a cleverer person. A claim of impossibility says that what is asked for is not a coherent thing to ask for. Kurt Gödel's theorems of 1931 show a request of that kind in mathematics: a fixed set of rules of proof for the arithmetic of whole numbers that can prove or disprove every statement it can express, and never proves a

contradiction. The two kinds of claim call for different responses. To a difficulty, one responds with more effort. To an impossibility, one responds by asking what the coherent version of the request is, and pursuing that instead.

What the Theorem Leaves in Place

The theorem excludes one thing, and it is as important to name what it leaves in place. A consistent theory that agrees with general relativity at large scales and with quantum field theory (quantum theory applied to fields such as light) at small ones is allowed, provided it does not find both survivors perfected in one comparison. Gravity treated at low energies as an *effective field theory*, a description that holds within a stated range (here, gentle curving and small waves) and states its own range, is allowed too. A small wave of the arena can be described as if it were content, and that description then has quanta. They are quanta of the description, and never quanta of the arena.

A picture of a crystal shows what a quantum of a description is. A crystal is a lattice of atoms held in place by the forces between them. Strike it, and a wave of motion runs through it: each atom swings a little about its place and passes the swing to its neighbors, and that wave is sound. In quantum theory each way the lattice can vibrate takes up energy only in equal steps, and one step is called a *phonon*: the smallest packet of vibration in a crystal, a quantum of the description of sound, and not a separate thing in the crystal. The crystal is made of atoms, and nobody concludes from phonons that the crystal is made of phonons.

The phonon is the nearest sense in which a quantum of the arena's small waves can be made intelligible, and the picture carries its own limit. A crystal is a *held structure*: its atoms are held in place by the forces between them. The arena is not a held structure; it has no parts that forces hold in place, as the atoms of a crystal are held. So even the phonon's sense does not make such a quantum a quantum of the arena. In this book

the word *graviton* names, at most, a quantum of a description, the description of small waves of the arena's geometry, and never a quantum of the arena. Gravitation is not quantized. What the theorem excludes is one theory whose findings perfect continuity and symmetry together.

Joined Through the Act of Measurement

The coherent version of the request is available, and the framework says what it is. The first axiom, quoted in Chapter 4, reads: "The immanent is more fundamental than the omniscient and/or the transcendent. The omniscient and the transcendent are in conjugate relationship to one another". The two theories are the omniscient and the transcendent, reaching across into one another. What relates them is the immanent: the interaction, the measurement, in which the geometry and the quantum state are both brought to bear on a single event. No third formalism stands above both.

The same point can be put in a grid, which Chapter 19 sets out in full; here it is in words, enough for this chapter. The six measures of physics fall into two rows and three columns. One row holds what is found, the content: pattern (which, as Chapter 1 said, takes in mass), force and probability. The other row holds the setting, the context: space, time and possibility. The three columns are the three modalities, and each column pairs one content with one setting. The omniscient column holds mass in space; the immanent column, the middle one, holds force in time; and the transcendent column holds probability in possibility. The two theories share that middle column, and they share it differently: one treats force in time smoothly, the other in steps. A unification would have to connect the omniscient column to the transcendent column without passing through the immanent column, and the first axiom says no such connection exists. The connection runs through the interaction, and not directly.

So the two theories are already unified, in the only way they can be: as the two conjugate sides of one domain of comparison, met in the act of measurement. The unification of the two theories is a product, the two theories set side by side and joined through the act of measurement, and not a merger. I think it is what Penrose's revolution in ideas would have to turn out to be.

The first axiom applies to the three scales as well. Our own scale, where every measurement is made, is more basic than the very small and the very large, because every finding about either is made by an interaction at our scale. The very large and the very small are the conjugate pair. Both theories are reached from our scale: every measurement that tests either is made with instruments of human size. The wave of quantum theory is defined over the space of all arrangements a system of Newton's mechanics can have, and the geometry of general relativity extends the geometry of the space and time of Newton's mechanics. The mathematics of both (rules of smooth change, the principles that pick out the actual motion among all possible ones, the mathematics of symmetry) was first worked out for that older mechanics. Newton's mechanics was found first, at our scale, and the order of discovery fits the first axiom's placing of the immanent as the ground: the scale at which comparison is made is the scale from which all other physics is reached.

The one formalism in which the two theories are written together shows exactly this joining: the two set side by side and joined through the act of measurement, and not merged. In the spectral geometry of Ali Chamseddine and Alain Connes (1997), a geometry is described by the set of frequencies at which it can vibrate, as a drum can be known by the tones it gives when struck. There the arena and the content are two factors of one product, set side by side; each character holds over its own factor, so that the joining is made through scope, each character keeping to its own part; and the result is no single geometry perfecting both characters. The theory's central rule, the one that sets the tones of the product, acts on each factor in its own character, which is this chapter's claim that what relates the two theories is the act. And the one step of the theory that

forces a single common measure on both factors is exactly the step at which the relations the theory gives at its highest energy disagree with two measured values. The two values are the weak mixing angle, which sets how two of the forces share out between them, and the mass of the Higgs particle, the particle whose field gives the other particles of matter their masses. Incommensurable means having no common measure, and the step that imposes a common measure on the two factors is the step at which the theory disagrees with what is measured. My book Non-Commutative Geometry and the Immanent Metaphysics works this out, and the Appendix gives the technical form.

How Each Program Trades

The theorem forbids one thing: symmetry and continuity perfected together in what a single theory finds. It leaves room for approaches to quantum gravity that do not ask for that thing, and it says that every program will be found to have made a trade between the two survivors. A program makes a *partial trade-off* when it gives up some symmetry for continuity, or some continuity for symmetry; the giving up is the signature of the theorem at work. A trade is of the kind the theorem permits when the program keeps the arena of space and time out of the quantum state, the quantum description of what is found. A program that puts the arena itself into the quantum state treats the setting of physics as if it were a content found within it, and it requires one of the two pairings the theorem excludes. So every program makes a trade, and a trade is of the permitted kind only when the program keeps the arena out of the quantum state. The four programs that I read below all put the arena into the quantum state. This way of reading the programs does not depend on how well they succeed technically, and I offer it as a reading.

String theory, which describes matter and gravity as the vibrations of tiny strings, keeps the tool of symmetry whole and lets the geometry become secondary. It keeps the full apparatus of quantum theory and ex-

tends it, with further symmetries and correspondences (physicists call them dualities) beyond those the older theory had. What it gives up is the continuity built into general relativity: the smooth geometry is no longer the primary object, but a background on which the symmetric structure is defined, or something to be recovered from that structure in a limit. Gravity itself is treated as content, one way of vibrating among the ways a string can vibrate. On this reading, string theory has given up some continuity to keep symmetry perfect.

Loop quantum gravity makes the opposite trade. It keeps general relativity's geometry: there is no background, the geometry is the primary object, and general relativity's character, relational and whole first, is kept. What it gives up is quantum theory's exact symmetry over the whole of space: the symmetries of flat space are to be recovered only in a limit, and the recovery has been the program's recurring difficulty. The geometry is quantized, so that areas and volumes come in separate values: the steps that symmetry produces have been applied to the setting itself. The program takes these stepped values as what a measurement of area or volume would find, a least step of space, and so it carries a forbidden pairing, because context has no least step. The theorem's verdict falls on what a program requires its findings to be, and the program's founding commitment, the step it takes on the arena at its start, fixes that requirement before any calculation is made. The calculations made under that commitment stand or fall as physics. On this reading, loop quantum gravity has given up some symmetry to keep continuity first.

The other programs can be read the same way. Causal set theory breaks the smooth whole into separate points outright and seeks to recover smoothness statistically. A lattice and a causal set each have a least step, and a least step of space or of time is what the framework expects will not be found. Asymptotic safety keeps the smooth whole and asks the apparatus of symmetry to behave at very short distances in a way it has not been made to behave anywhere else. The apparatus it asks to behave so is the geometry treated as a quantum field, and so the program carries a

quantum state of the geometry past the Planck scale, as the foundation of the arena. The arena has no such state, since gravitation has no quanta, and past the Planck scale no other description takes over.

Each program gives up something from one survivor to gain something from the other. That giving up is the structural signature of the theorem at work, and I expect it to be found in every proposal that is ever made. String theory, loop quantum gravity, causal set theory and asymptotic safety make their sacrifice on the arena itself: each treats the setting of physics as if it were content, and each requires one of the two pairings the theorem excludes. Each of the four programs therefore asks for a merger the theorem excludes, and none of the four makes a partial trade-off of the kind the theorem permits.

The framework gives one question to put to any proposed unification: where does it give up symmetry for continuity, or continuity for symmetry? If the answer is nowhere, and the proposal claims to find both perfected, the framework expects the proposal to be incomplete or inconsistent in a place that has not yet been examined. If the answer is a definite place, and the proposal keeps the arena out of the quantum state, the proposal is a partial unification of the kind the theorem permits, like the kinds of theory the theorem leaves in place, named earlier in this chapter. Whether it is a good one is a question of physics, not of metaphysics. If the answer is a definite place, and the proposal makes its sacrifice by putting the arena itself into the quantum state, the proposal treats a context as a content, and it requires one of the two pairings the theorem excludes.

Gravity Is the Setting

One more reading bears on where the programs have most often looked: it concerns what gravity is. In the grid of the six measures, gravity is in no cell. Gravity is the shape of the context row: the curving of space and time by which mass and force are related to their setting. Chapter 12 put it directly: space is the pattern of connection among causes, and its curving is

gravitation. On this reading, gravity is context rather than content. Gravity is the form of the domain itself, the geometry of the paths along which signals and bodies can go, and no occupant of the domain.

One standard argument for a particle of gravity meets this reading at a black hole. (The word graviton, as said above, names in this book at most a quantum of a description.) The argument asks: if gravity were carried by graviton particles passing between bodies, how would those particles get out of a black hole, from inside whose horizon nothing escapes, to carry its pull to the space beyond? The standard reply is that the outside field was set up before the horizon formed, and that the gravitons in question are "virtual": entries in the calculation that need not travel outward as real particles do (physicists say they are off-shell). The reply is consistent within its own method of calculation, which works by small corrections laid over a smooth background; but it requires virtual particles to differ from real ones in exactly the way needed to avoid the problem. The Immanent Metaphysics observes: if the gravitational field is not being signaled outward, then the field is a feature of the geometric setting, and no consequence of particles passed back and forth. That conclusion agrees with the claim made here.

The framework keeps three relations apart: opposition, conjugation and merger. Opposition is like adding: when one side goes up, the other goes down by the same amount, and either side can reach zero. Conjugation is like multiplying: neither side can reach zero, and neither can be removed in favor of the other. The first axiom is sometimes written as a short formula: the immanent equals the omniscient times the transcendent. The formula reads the immanent as the conjugate relationship of the omniscient and the transcendent, and it multiplies no amounts. At the small-scale limit of a domain, conjugation takes the form of strict opposition; at the large-scale limit, how real either side is rises in direct proportion to how real the other is. A merger would erase the difference between the two entirely.

General relativity and quantum theory are conjugates. They are distinct, because each perfects a different pairing; inseparable, because every finding has a content and a context; and not interchangeable, because neither can do the other's work. They do not oppose each other, and they cannot be merged. They are two descriptions of the same domain, each capturing what the other must leave out.

The principle of equivalence shows, from inside physics, that gravity is context and not content. Every body, whatever it is made of, falls in the same way. That every body falls in the same way is sameness of content across difference of context, symmetry of the plainest kind. That symmetry makes gravity belong to all contents alike rather than to any one content. What belongs to all contents alike is the setting that the contents share, and not a content. The last part of this chapter shows, from what a body feels, that gravity is the setting. Of the three kinds of acceleration a body can undergo, only gravity's is one that no local push can produce, because it is the shape of the setting and not an action within it.

Gravity, being context, has no particle, since every particle is a kind of content. It has no label that comes in steps, no charge, and no carrier that a shield could stop. The quanta that a description of its small waves has are, as the crystal and its phonons showed earlier in this chapter, quanta of that description, and never quanta of the arena. The arena itself, the background of every comparison, stays context. To look for a particle of gravity at all is to treat a context as a content, which is the same confusion that the demand for unification makes at the level of theories.

Physics has its own versions of this thought: Ted Jacobson's reading, in 1995, of Einstein's field equations as an equation of state (a rule of the kind that relates the pressure, temperature and volume of a gas), and Erik Verlinde's reading, in 2011, of gravity as a force arising from entropy, the measure of disorder that Chapter 13 met. I take those as evidence that the reading is not foreign to physics.

The reading stands without the grid of the six measures. Anything that plays the part of context acts on everything, admits no shield, pulls on each body through the same mass by which the body resists being moved, and carries no charge that comes in steps. Gravity has all four marks. The structure of comparison gives the law of continuity to everything that is of context, because context has no least step; and the arena of space and time is context. And in the spectral geometry of the product, gravity comes from the arena's factor, and the theory's fields include no spin-two field, the kind of field that a particle of gravity would need.

Two Times, and a Floor Beneath Both

Chris Isham, and Jeremy Butterfield with Isham (in 1995 and 1999), identified what has come to be called the problem of time in quantum gravity. General relativity treats time as a dimension of the smooth whole, with no preferred way of slicing it into moments and no clock outside it. Quantum theory treats time as the variable against which the state changes, outside the system and shared by all of it. Any attempt to quantize the geometry has to decide which of these time is, and neither answer satisfies the other theory.

The problem of time is the cross-modal reach applied to time. Chapter 14 found three kinds of time: geometric time, laid out in the shape of space and time; evolution time, the parameter of quantum theory's smooth change; and event time, the order of findings for a comparer. This line-up matches the times of the two theories, and the time of measurement, with the modalities, and it agrees with the matching Chapter 14 made in the focus of time. General relativity's time is omniscient: geometric, structural, reversible in its equations, the same for every observer who reads the geometry. Quantum theory's time is transcendent: a parameter of change, sweeping through possibilities that have not yet resolved. The moment of measurement, the time of the interaction, is immanent, and in

that moment the two meet and disagree. A reader may ask why event time, whose findings are discontinuous and symmetric, as what quantum theory finds is, is not quantum theory's own time. The answer is the split at two levels that this chapter has already shown: what quantum theory finds, each click with its record, is laid down in event time, at the moment of measurement; the time in quantum theory's formalism, the dial of its rule of smooth change, is evolution time. Chapter 13 showed that the arrow of time belongs to the immanent, to process, and to neither the omniscient nor the transcendent alone. The problem of time is what one gets by asking the two conjugate times to be one time without going through the third.

John Wheeler's program of 1989, under the slogan "it from bit" (the bit of information comes before the thing), and the approaches built on information that followed it, have looked for something more basic than either space and time or the quantum state. The Root Tautology of Chapter 1, that comparison, interaction and relation are one act, is a starting point of that kind: comparison is the most elementary act of information. What the cross-modal reach adds is that information is not a neutral floor. Information carries the modalities in it, and a program that sets information as the common floor beneath both theories will find, I expect, that the floor has the same two conjugate sides the theories had, and will have to choose which one to perfect.

Noether's theorem, finally, shows the reach inside Newton's physics. Every smooth symmetry of a motion yields a conserved amount. The symmetry is omniscient; the motion, force in time, is immanent; and the theorem is the tool of the omniscient reaching into the domain of the immanent to produce the unchanging amounts that make the motion intelligible. That the reach appears there too is a sign that it belongs to more than the two great theories: it is the general condition under which any tool describes any domain.

Back in the Closed Room

Return now to Einstein's room. The principle of equivalence is stated for a point, and it is exact only there. Over a region, the cases come apart. Let the falling room be large. Objects released at its two sides drift toward one another as they fall, because each is headed for the center of the planet, and the center is one point. Objects released one above the other drift apart, because the lower one is nearer the center and falls faster. The room in the rocket finds no such drift: objects released anywhere in it keep their spacing exactly. And a room on a spinning platform, which is also a room whose floor presses, finds that objects released in it move outward from the axis, faster the farther out they are, and curve as they go.

So the principle of equivalence, which says the cases cannot be told apart at a point, also says how they are told apart over a region: by comparison. One body at one place feels only its own state. Two bodies, or one body at two times, show the kind of acceleration they are in by what happens between them. The principle of equivalence, read over a region, has the shape of claim that I make throughout this book. What is a fact at a single point becomes structure only when a comparison is made across a difference of context. The three kinds of acceleration are the three things such a comparison can find.

The Three Kinds

Name the three by what they do to the *free paths* of neighboring bodies, the paths that bodies take when nothing pushes them, because those paths are what an observer in any of the three actually finds.

The first kind is *convergent*: the acceleration of gravity, the fall toward a center of mass. Every free body in the region is headed for the same point, and the paths of neighboring bodies close on one another as they go. The observer feels nothing: free fall is weightless, with no floor, no push and no sense of motion. What the observer finds, by comparison, is the *tide*:

the slow drawing together of things released side by side, and the slow drawing apart of things released one above the other, both of them signs of a single center toward which all the paths run. In empty space the tide squeezes in two directions and stretches in one, with no change of volume at first. The distortion it produces (physicists call it shear) then draws the paths together, and where there is matter among them, the matter draws them together directly. Both are the convergent kind.

The second kind is *parallel*: the acceleration of a push, as in a rocket, a lift, or the floor of any room driven in a straight line. Every free body in the region, once released, keeps its direction and its spacing from every other; the paths stay parallel, and the room speeds up past them. The observer feels weight: the floor pressing upward, the same everywhere in the room, the same at every height. What the observer finds by comparison is nothing further. There is no tide: released objects fall to the floor together, and land as far apart as they began.

The third kind is *divergent*: the acceleration of rotation, as on a spinning platform, in a centrifuge, or in any *frame* (the room, platform or vessel from which an observer watches) that turns about its own axis. Every free body, once released, moves outward from the axis and never returns to it; the paths, seen from the platform, run apart from a common line, and the farther from the axis a body is, the faster it is thrown. The observer feels an outward pull that grows with distance from the center, and a sideways push on anything that moves, so that a thrown ball curves. What the observer finds by comparison is that the pull is not the same everywhere: it is zero at the axis and grows without limit outward. That alone tells it apart from the push of the rocket, which is the same at every place in the room.

These three kinds are three states of experience, and I insist on the word "experience". They are three states a body can be in, each with its own feel, and each with its own signature when a comparison is made, and a pilot can tell the three apart without an equation.

Three, and No Fourth

The claim that there are three kinds, and not four, needs an argument, and the argument is short. Release a family of free bodies together in a small region, and watch them from a frame that is accelerated in some way. Left to themselves, the bodies follow the straightest paths available in the domain; the frame, being accelerated, sees those paths as curved. The family of curves either has a common *locus*, a point or a line that all the paths share, or it has none. If it has one, the locus lies ahead of the paths and the paths run together toward it, or the locus lies behind the paths and the paths run apart from it. If the family of curves has no common locus, the paths run parallel. Ahead and behind are said along the paths themselves: ahead is where the paths are going, and behind is where they are coming from, as the paths on a spinning platform come from its axis, a line in space, and run away from it. Ahead, behind or nowhere: those are all the possibilities, because a shared locus for a family of paths is in their future, in their past, or absent, just as a quantity that must be less than, greater than or equal to a given value has three cases and no fourth.

Any actual acceleration over a region is some mixture of the three. A body falling toward a planet, on a platform that is spinning, while a rocket pushes it sideways, is in all three at once, and an observer's comparisons will find the signatures of all three laid over one another. The mixture is the three kinds, inseparable in one instance, and no fourth kind; every partition the metaphysics makes is to be expected to show its parts inseparable in each instance. The third axiom says that the three modalities, in their classes and their instances, are distinct, inseparable and not interchangeable. Distinct: the three signatures differ and can be told apart. Inseparable: every actual case involves all three to some degree, if only the degree zero. Not interchangeable: none of the three can be made from the others, and that last point matters most, as the chapter shows below.

The claim has an exact scope. The ways in which a family of free paths can relate to a common locus are three, and the three kinds of acceleration are those three relations as a body experiences them. The claim counts

kinds of relation between free paths and a locus; it makes no statement that every mathematical description of relative motion has exactly three pieces, or that any formula has exactly three terms. Physics has its own way of taking relative motion apart: the equation Amal Kumar Raychaudhuri published in 1955, a formula for how a cloud of free paths swells or shrinks. Its terms answer to the three kinds. With no curving, distortion or rotation, the paths keep their spacing: the parallel kind. Matter, and the distortion of the tide, draw them together: the convergent kind. Rotation spreads the paths apart: the divergent kind.

The Three Kinds and the Three Modalities

Chapter 4 set out the three modalities. The omniscient is the structural, the repeatable, the shareable, the view in the third person: existence, pattern, law. The immanent is the happening, the relating, the interacting, the here and now: interaction, change, reality. The transcendent is the valuing, the creating, the potential, the generating: choice, possibility, principle. As Chapter 7 showed, causality is omniscient, change is immanent, and choice is transcendent. The three kinds of acceleration sort into these three without strain. This line-up matches kinds of acceleration, as a body undergoes them, with the modalities. By the rule given earlier in this chapter, what is considered decides the match. So the fall can go with the transcendent here while general relativity's domain, the very large, goes with the omniscient in the line-up of scales. A kind of acceleration that a body undergoes and a domain that a theory describes are two different things considered.

Read in this triple, the first axiom says that acceleration in a straight line (immanent) is more basic than either gravitation (transcendent) or rotation (omniscient), and that gravitation and rotation are the conjugate pair: the many to the one and the one to the many, the paths running to-

gether toward a center and the paths running apart from an axis. I make this reading of the first axiom as an identification, and the reading is not a theorem of mechanics.

Acceleration is the quantity in which all three kinds are expressed. In each kind, the force a body feels is in proportion to the mass by which the body resists being moved, so each kind is fully described by an acceleration, the force for each unit of that mass. Newton's field of gravity is the acceleration that masses produce in a free test body, and the curving of space and time is observed as the acceleration of neighboring free bodies relative to one another. A body moving round a circle accelerates toward the center by an amount equal to its speed multiplied by itself, divided by the radius of the circle: the acceleration is fixed by the speed and by how sharply the path curves. Newton's second law, force equals mass times acceleration, puts acceleration at the center. So the triple fits the first axiom: the immanent idea, acceleration as such, whose plainest case is the push, is the idea in whose terms both the transcendent kind (gravitation) and the omniscient kind (rotation) are stated.

The convergent kind is transcendent. For a test body, one small enough that its own mass does not disturb its setting, the fall is fixed entirely by what is outside the body: by the mass and energy of the whole surrounding region, which no push applied to the body alters, and which no description confined to the body's own place can contain. The fall is the same for every body, whatever the body is made of. It is the acceleration in which the body does nothing at all: free fall is the body simply existing in a domain whose shape is given. It is the many to the one: every path in the region runs to one center, the same for all of them. And it is context, potential and surrounding: the setting within which both paths and interactions occur.

The parallel kind is immanent. It is the push: a force applied at a point of contact, felt as weight, here and now. It exists only while the interaction continues; the moment the rocket's engine stops, it is gone. It is exchange: the rocket speeds up by throwing mass out behind it, and its acceleration is the other side of that exchange. It is one to one: a body and what pushes

it, in direct relation, with the free paths around them staying as they were. And it is change, the altering of a state by an interaction, which the Immanent Metaphysics assigns to the immanent. The grid of the six measures puts force in time in the immanent column, and the parallel kind is force in time and nothing else.

The divergent kind is omniscient. The divergent kind is the relation of a frame to the frame's own axis, and is structural. The divergent kind is fixed by the knowable shape of the motion, the curving of each path about the axis, the pattern of the whole motion. It needs no outside mass and no continued push: a platform once spun keeps spinning, and the outward pull stays without anything further being done. It is the one to the many: from a single line, paths spread out in every direction, and the effect grows without limit outward from the source. It is the kind that Isaac Newton, in a note in his Principia of 1687, found to be detectable from inside, without any comparison to anything outside: the water in a spinning bucket climbs the bucket's wall, and the climb shows the spin. Ernst Mach objected in 1883 that the distant stars might play a part in rotation after all, and that question has never been fully settled; but what a spinning body feels, and what the climbing water shows, is not in doubt. Rotation is the acceleration that knows itself.

The partition into three kinds stands on its own, on the three cases of ahead, behind and nowhere; the matching with the modalities is a reading laid over it. What the modalities add is a reason the three should be these three: the same three that shape comparison, interaction and measurement everywhere else in this book. To my mind, the reading is confirmed by how little has to be forced. Each kind lands where its modality places it: the fall, and not the push, is the one fixed by the world around the body; the spin, and not the fall, is the one that spreads outward from its source; and the push is the one made here and now, at a point of contact. The sorting could have come out otherwise, and the Appendix records which results would have counted against the matching of the three kinds of acceleration with the three modalities.

The One Kind No Push Can Make

The point that matters most is the one the third axiom calls being not interchangeable. None of the three kinds can be made from the others, and one of them cannot be made at all by anything a body does.

Two senses of the word "curvature" must be kept apart here. The curvature of a path is a property of the path: a body moving round a circle has its acceleration toward the center, speed times speed divided by radius, in flat space and time just as in curved. The curvature of space and time is a property of the arena: general relativity puts it in place of the force of gravity, and what shows it is the acceleration of neighboring free bodies relative to one another, the tide. Rotation and its pull toward the center belong to the first sense, gravitation to the second, and general relativity does not identify the two.

The parallel kind is made by a push: any body can produce it for itself, given something to push against. The divergent kind is made by a spin: any body can produce it for itself, given something to spin against, and once produced it needs no further action. The convergent kind cannot be produced by a body at all. No push, applied anywhere, to anything, makes the free paths of neighboring bodies run together toward a common center. A push moves the thing pushed; it does nothing to the paths of free bodies that are not pushed, and convergence is a fact about those paths.

One could push each body of a family separately and aim them all at one point. Pushing each body separately toward one point is many parallel accelerations imitating a convergent acceleration. The imitation fails the moment a body is added that was not pushed: the added body does not converge. The convergent kind belongs to the free paths themselves, to the shape of the domain in which they lie. No push applied to one body reshapes the free paths of the others; what shapes the setting is the content in it.

That the convergent kind belongs to the shape of the domain, which no push applied to one body reshapes, is what it means to say that gravity is context rather than content. In the grid of the six measures, a push is content acting on content: force applied to mass. A spin is a frame's relation to its own axis, the shape of a motion held in a state. The fall is in no cell of the grid. The fall is the shape of the context row, the curving of the space and time in which mass and force are related to their setting. The fall is what all contents share alike, and what all contents share alike is the setting, and not a content.

General relativity says the same in its own terms. In that theory, gravity is not a force. Free fall is motion with no force acting; what is called the force of gravity is the floor pushing up, which is the parallel kind, felt because the floor stops the fall. The running together of free paths is the curving of the geometry, and the curving is fixed by where mass and energy are, which is to say by what is outside any given body. I read this as physics arriving, from its own side, at the same placing: gravity is the form of the domain, and no occupant of it. What follows was set out earlier in the chapter: gravity, being context, has no particle and no quanta, and the quanta that a description of its small waves has are quanta of that description only. The three kinds of acceleration support, without the grid, the reading that gravity, being context, has no particle and no quanta. The convergent kind is the one no local action produces, and a thing that no local action produces is hard to read as something that goes from one place to another.

The room has one thing left to say. The principle of equivalence says that at a point the fall and the push are one. On this reading, that is the statement that at a point, without comparison, context and content cannot be told apart: the shape of the setting and an act within it feel the same to a body that has nothing to compare. Over a region they come apart, because a region allows comparison across a difference of context, and a difference of context is where the context shows itself.

What We Now Know

We now know what the long search has been asking for, and why it has not been given. Each basic theory of physics describes its domain with the tool that belongs to the modality of the other theory's domain. It must, because a tool applied to a domain of its own kind says nothing, and the two modalities are conjugate, each giving an account of the other. The reach is what makes each theory powerful. A formalism that was both theories at once would perfect symmetry and continuity together, and that is the one pairing comparison cannot make. The failure to unify the two theories is no problem awaiting a solution. It is a structural exclusion, and the coherent form of the request it refuses is already met. Both theories are right, each in its range, for a reason that also says why their ranges cannot be made one. Among the programs that keep the arena out of the quantum state, which is nearer the truth is a question for physics, and I leave it to physics.

We also know what gravity is, on this reading: the shape of the setting, context and not content, with no particle and no quanta of its own. At a point all acceleration is one; over a region it shows itself in three kinds, the fall, the push and the rotation, and no fourth; and the fall is the one kind that no push can make, because it is the shape of the setting and not an action within it.

The two theories are conjugate sides of a single domain, related through the interaction in which both are brought to bear on a single event, and through no third theory above them. Each captures what the other must leave out, and each, with its tool, reaches across the divide between the modalities to describe a domain whose modality its tool does not share. They are not rival descriptions of one reality. The reach is the power. That the two reaches cannot be combined into a single grasp is the limit. And the limit is not ours; it is comparison's.

For the question of this book, what cannot be and why, the chapter adds the widest case. Determinism asked for symmetry and continuity perfected together in the causes of the world; Bell's theorem, with the ex-

periments that confirm it, showed that lawfulness and locality cannot both be had absolutely; here the same exclusion meets the whole of physics. Like the faster than light signal of Chapter 12 and the trip back to last Thursday of Chapter 16, one perfected theory of everything physical is excluded by structure, and by no shortfall of skill or effort. What can be is what physics already has: two theories joined side by side, each right in its range, met in every act of measurement.

Summary

General relativity and quantum theory are the two survivors of the Incommensuration Theorem, made physical: general relativity is continuity with asymmetry, quantum theory is symmetry with discontinuity. Each perfects one character in what it finds, and carries the other, exactly, in its formalism. Each reaches, with the tool of one modality, into the domain of the other, as close as one likes and never the whole, and the reach is required, because a tool applied to its own kind says nothing. The two theories are joined side by side, through the act of measurement: a product of conjugates, and never a merger. They cannot be unified in a single theory whose findings perfect both continuity and symmetry.

Gravity is the shape of the arena, context and not content, and it has no particle. Gravitation is not quantized. A graviton is at most a quantum of a description of the arena's small waves, never a quantum of the arena, as a phonon is a quantum of the description of sound in a crystal, with the stated limit that a crystal is a held structure and the arena is not. String theory, loop quantum gravity, causal set theory and asymptotic safety each make their trade on the arena itself, putting the arena into the quantum state, and so each of the four asks for a merger the theorem excludes.

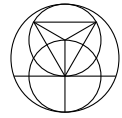
At a point, all acceleration is one. Over a region it is three: the free paths of neighboring bodies run together, run parallel, or run apart. The fall, the push and the rotation are the three, and there is no fourth, be-

cause a family of paths has its common locus ahead, nowhere or behind. The fall is transcendent: set by what lies beyond the body, the context in which every path lies, the same for everything, the many to the one. The push is immanent: an exchange at a point of contact, felt as weight, here and now. The rotation is omniscient: a frame's relation to its own axis, structural, knowable from the shape of the motion, lasting without further action, the one to the many. They are distinct, inseparable and not interchangeable, as the third axiom says the modalities are, and the one that cannot be made is the one that is the shape of the setting.

The chapter's statements are of three kinds, and each keeps its own standing. The exclusion of one perfected theory is a structural exclusion, a claim from the structure of comparison itself: it says what physics will not find, and it says why. The reading of each theory as one survivor, of the programs as trades, of the problem of time as the reach applied to time, and of complementarity and Noether's theorem as instances of one structure, are readings of known results. That there are exactly three kinds of acceleration is a structural claim, resting on the three ways a family of free paths can relate to a common locus. The matching of the three kinds to the three modalities is a reading of known physics, and it gives the structural reason the principle of equivalence has the form it has, and why the three cases come apart over a region as they do. That gravity is context, and will not be made a force among forces, is a structural exclusion, and it does not rest on the grid alone. General relativity is right about gravity, and the chapter has said what it is right about.

On my reading, gravity is the form of the domain in which every content and every act finds its path, and gravity is neither a content among contents nor an act within the domain. The principle of equivalence, which began as a thought about a closed room, says what the metaphysics says about content and context everywhere: alone, at a point, the two cannot be told apart, and comparison, across a difference of context, is where the difference appears.

Each theory reaches into the domain of the other modality as close as one likes, everywhere, and never covers the whole of that domain. Is there, then, a limit to what physics can ever find out, or only a limit to what it has found so far?



Chapter 18

The Three Horizons

A Photograph, a Glance, and What Neither Holds

Is there a limit to what physics can ever find out, or only a limit to what it has found so far? Most of us carry the second picture. Science, as it is ordinarily practiced, sorts things into two kinds: the known and the unknown. The unknown is taken to be the not yet known. It is a frontier, and the whole method of inquiry is a way of pushing the frontier back. On that picture, what we know today and what we will know tomorrow differ only in effort, in instruments and in time, and never in kind.

Hold three images side by side. The first image is a photograph. Everything in it is there at once, and nothing stops you from looking at any part of it. The known is like the photograph: complete, available, open to inspection. The second image is the act of looking around you, now. You see a great deal, always from where you stand, and never all of it at once. The unknown is like the act of looking around you: the process of coming to find out, perception in the here and now, which gives a view of things without giving the whole. The third image is the one no image can hold: what overflows every capture, what has no fixed shape or pattern by which a description could take it in. What overflows every capture is the unknowable.

The *unknowable* is a third kind, beside the known and the unknown: what cannot be known, even in principle, by any observer, at any time, by any method, because comparison itself does not reach it. Chapter 1 showed that every finding is a comparison, the act of finding whether two things are the same or different, in some respect; so what comparison cannot reach, nothing can find. The unknowable is no gap in present knowledge waiting for a future discovery. There is a place that reasoning does not get past, and where no amount of experience will ever yield what is sought. That place is where the unknown stops and the unknowable begins.

The line between the unknown and the unknowable is the line between ordinary knowing and what I call the *Great Mystery*, the name for the place where the unknown and the unknowable converge. The Great Mystery has three forms, and they are the three horizons of this chapter: the Mystery of the Horizon, at the limit of the very large; the Mystery of the Infinitesimal, at the limit of the very small; and the Mystery of the Subconscious, at the limit of the middle scale, the scale of the self, which this chapter calls the surface of unconsciousness. The Mystery of the Subconscious, the surface of unconsciousness, the horizon of the self, and, for short, the unconscious are four names for one edge: the first names it as a form of the Great Mystery, the second as the surface past which awareness fails, the third as one of the three horizons, and the fourth is the everyday word for the same moving edge.

The three kinds of knowledge line up with the three modalities of Chapter 4: the immanent (process, what is happening here and now), the omniscient (structure, what holds the same for everyone, laid out as a pattern), and the transcendent (possibility, what makes things possible and new). The known has the character of the omniscient, complete like the photograph. The unknown has the character of the immanent, the finding out that is happening now. The unknowable has the character of the transcendent, what exceeds. The match does real work, and it is no decoration: each kind of knowledge is what it is because it has the character of its

modality. The omniscient is the mode of structure and completeness; the immanent is the mode of process and becoming; the transcendent is the mode of excess and creation.

The matching of the three kinds of knowledge with the three modalities is the first of several matchings with the modalities in this chapter and the next two chapters, and one rule holds for all of those matchings. Which modality a thing is matched with depends on what one is considering, and it is no label fixed to the thing. Think of one woman who is a mother to her son, a daughter to her own mother, and a teacher to her class: she is one person, and which name fits her depends on who is looking, and in what relation. So each time a matching changes in these chapters, the text says what is now in view. Here what is in view is the three kinds of knowledge.

The unknowable has a shape. There are exactly three places where the knowable gives way to it, and two of them are the edges of physics. In this chapter I describe the three places, show why there are three, and follow the two physical edges into what physics has found at those two edges. Behind all of this chapter stands a general point, that a theory can never run without a break into the thing that the theory describes. I argue that point in my book Unlocking Reality, and I take that point up again in Chapter 20. Everything the argument of this chapter uses, this book has already shown, or this chapter shows.

The Horizon of the Very Large

Picture a traveler crossing a wide plain toward the line where the land meets the sky. The line never comes nearer. Each mile brings into view land that was hidden, and each mile shows a new line beyond it. The traveler passes every particular horizon, and no traveler ever reaches a place where horizons stop. I use the word *horizon* in exactly that sense: an edge

of what can be known from here, which moves back as it is approached, so that any one horizon can be passed while the kind of edge is always there.

The first horizon is at the very large. It is the edge beyond which perception does not reach, however far one travels and however sensitive one's instruments become. In physics this edge is set by the speed of light, a little under three hundred thousand kilometers each second. That figure is measured, and it serves here only to picture the edge; the argument of this chapter uses the edge, and not its value. No signal, no information and no cause crosses from one event to another faster than light.

Chapter 12 drew the picture. From any event, any one happening at one place and one moment, a cone opens forward into the future, holding every event it can affect, and a cone opens backward into the past, holding every event that can affect it. Whatever lies inside the cones can be known from here, in principle. Whatever lies outside both, in what physicists call the *absolute elsewhere* (the events that no cause joins to this one, in either direction), cannot. Chapter 12 also showed that the speed of light is, at root, the speed of no one thing: it is the outer edge of the domain of comparison, the whole range within which comparisons can be made. And space is nothing other than the pattern of which events can affect which. Here one further feature of the edge matters.

The horizon is an edge that moves; it is not a wall, and it is not a thing one can walk up to and touch. Any particular horizon is passed by moving toward it, or by waiting: what lay outside the cones of this event enters the past cone of a later one. The elsewhere of yesterday is the past of today. And a new horizon always appears. No event, however late, has a cone that contains everything. Our universe grows, and its growth is speeding up: Adam Riess and colleagues measured the speeding up in 1998, and Saul Perlmutter and colleagues in 1999. In a universe like that, some events will never enter the past cone of a given observer, at any time, and the edge of those events is called the *cosmological event horizon*. It is a

different edge from the cosmological horizon of Chapter 12, the surface that bounds what has reached us so far: the cosmological event horizon bounds the events whose light never will reach us.

So two things must be kept apart. I call any one particular horizon an *instance*, and the kind that every instance belongs to the *class*. The horizon as an instance is temporary; the horizon as a class is permanent; and the edge of the knowable is the class, not any instance.

The horizon as a class is the first form of the Great Mystery. From any position there is always something past the horizon one has, something that exceeds what can be known from where one stands, and what it is changes as one moves. For the horizon as a class, the edge of the knowable itself, the word "beyond" has no meaning. The limit the Incommensuration Theorem draws there is, in a phrase from my own account of that theorem, "the curvature of the field itself, rather than a wall at the edge of an otherwise open field". Think of the curve of the earth, which makes the traveler's horizon, and not of a fence at the end of a flat meadow. Moving changes the instance, and never the class.

The Horizon of the Very Small

Now picture a particle, an electron say, a speck far too small to see, loose in an empty box. The particle is somewhere in the box, so the particle has a position. It is moving somehow, so it has a momentum: its amount of motion, its mass times its speed, together with the direction it moves in. We are a little unsure where it is, and a little unsure how fast it moves, and which way.

The second horizon is at the very small. It is the finest scale at which measurement is possible, and it is the home ground of quantum mechanics. Physics expresses it through Planck's constant, the small fixed amount that sets how big the packets of quantum theory are, and through the uncertainty relation that carries Werner Heisenberg's name, from 1927. The energy of one packet of light, for instance, is its frequency, the number of

times each second its wave rises and falls, times Planck's constant. Planck's constant is measured in units of *action*: a momentum times a distance, or an energy times a time. A constant in those units cannot by itself set a length or a duration. Planck's constant sets no least length and no least duration; how fine a comparison can be made depends on the energy brought to it.

Two things are easily run together here, and I keep them apart. Planck's constant is a number, and its number expresses a unit. The microscopic limit is that same constant read as the place where an edge lies. The difference matters because this chapter uses the constant in the second way: as the place of an edge, and never as a smallest size of anything.

The *uncertainty relation* says that position and momentum cannot both be pinned down perfectly at once. Earle Kennard proved the exact bound in 1927: the spread in a particle's position, times the spread in its momentum, is never less than half of the reduced Planck constant. (The reduced constant is Planck's constant divided by twice pi, pi being about three and one seventh, the number of times a circle's width fits around its edge.) The bound of half the reduced Planck constant is set by the structure of quantum theory. In that sense the relation limits how much can be known about a physical system at once.

Each spread is the width of a distribution, and it has no hard edge: the spread of places where the particle could be found, or the spread of momenta it could be found to have. Max Born's rule of 1926, the rule that gives the chance of each outcome from the state, gives those widths to the state itself; they belong to the particle's state, not to the instrument. The states for which the two widths multiplied together are as small as they can be have the shape of a bell curve, which mathematicians call a Gaussian. The point of the relation is that these widths can be given a concrete number. Narrow the one and the other widens; their product cannot fall below an amount fixed by the constant. So far this is a number and a rule about a number. The relation becomes a limit, an edge of the domain, only when one asks what happens as the position is narrowed without end.

Physicists themselves take the step from number to edge in the physics literature, and I borrow the step from that literature, as a reading of that literature. Pin down where a particle is to a smaller and smaller region, and the spread in its momentum grows, and momentum carries energy of motion. Confine any measurable mass to a small enough region, and the spread in its energy becomes enormous. Energy weighs, as mass does: that is Einstein's equivalence of mass and energy. So there is a scale at which that energy is enough to curve the space in which the measurement is being made so severely that the region closes on itself. The measurement would create a black hole larger than the region it set out to probe. The scale at which the region closes on itself is a bound on how sharply a comparison made from inside the world can place a body. The bound is a bound on findings, and not a least step of space and time.

That scale is the *Planck scale*. Every body has two sizes. Its gravitational size is the size at which its own gravity would close space around it; its quantum size is the size over which quantum theory spreads it. The Planck scale is where the two sizes are equal: where the description of the setting, the context, and the description of what is found, the content, change places as to which one dominates. The *Planck length*, the length of that scale, is built from Planck's constant, the speed of light and Newton's constant (the number that sets how strongly one mass pulls on another). I quote its size from the physics literature: of the order of ten to the minus thirty-five meters, which, written as a decimal, is a point, then thirty-four zeros, then a one.

Light from distant galaxies already shows that the Planck length is no least length of the arena, the space and time that together are the setting of every measurement (Chapter 17). Gamma-ray bursts are brief, enormous flashes from distant galaxies. In 2009 A. A. Abdo and colleagues reported that light from such a burst, of very different colors, arrived together after its long crossing. That places any grain of space that would slow light in proportion to its energy below the Planck length.

What the microscopic limit marks is a bound on joint knowledge, and not a least length: below the action that Planck's constant marks, no measurement distinguishes jointly what a pair of conjugate quantities would need to fix. (Chapter 3 called two quantities *conjugate* when each can be perfected only at the cost of the other, neither prior; position and momentum are such a pair.) The limit is a structural feature of the physical domain, and it is no limitation of instruments.

Like the large horizon, the small one moves back. What was once too small to observe comes within reach as instruments improve, and over the history of physics the reachable scale has come down by many powers of ten. But there is always a scale below which observation stops. The instance of the microscopic limit moves; the class of the microscopic limit persists.

The two ways of reading Planck's constant, as a number and as the place of an edge, also differ in modality, and that matters when the two horizons are set side by side. Planck's constant, as a number written in a theory, is known structure, a fixed and shareable pattern, and read that way it is omniscient; the number itself expresses only the unit in which action is measured. Used to say where the edge of the domain lies in scale, the same number takes on a transcendent part, set against the other constant that also marks an edge, the speed of light. It is one number read in two settings, and the second reading is the one this chapter is about.

The Horizon of the Self

Shine two laser beams so that they cross in a dark room. Through the electric and magnetic field alone, the waves pass right through one another, unchanged. Light carries no electric charge, and light does not see light. (Light does act on light, through fleeting pairs of charged particles, but so slightly that no lamp or laser makes the effect large enough to see; the ATLAS team at the Large Hadron Collider found evidence of it in 2017, and observed it in 2019.) So in every lamp and laser of ordinary life, one beam

passes through another unseen. In the same way, the act of comparing is not one of the things it compares. The self stands where the beam stands: it does the comparing, and its own comparing is not among the things it finds. The third horizon comes from exactly this fact: the self's own comparing is not among the things the self finds.

The third horizon belongs to the knower, and not to the world. It lies at neither the largest scale nor the smallest but at the middle scale, which physicists call *mesoscopic*: the scale of the self. It is the *surface of unconsciousness*. As one looks inward, seeking to know oneself, there is a point at which awareness fails; one becomes, with respect to some part of oneself, unconscious. The process of coming to know oneself is never complete. I cannot perceive the perceiver; I cannot perceive perception; I can perceive only what is the object of perception. Of the three (the perceiver, the perceived and the perceiving), perception can perceive only one. That is why the relationship itself is the thing we must study.

The gap is structural, and it has nothing to do with how diligently one looks inward. A self that knew itself completely would have to be both the knower and the known: the same as itself, as pattern, which is symmetry, and unbroken with itself, as process, which is continuity, both perfectly at once. The Incommensuration Theorem of Chapter 3 forbids exactly that: symmetry and continuity cannot both be applied, at once and absolutely, to one comparison, and the theorem is a structural fact about the nature of knowing.

As awareness approaches the surface, the surface recedes, in exactly the way the horizon recedes before the traveler. There will always be aspects of oneself still to discover. Self-knowledge deepens, and it never becomes total. The unconscious is a moving edge, a class rather than an instance; it is not a fixed region of the mind, a territory that could in principle be fully mapped.

This third horizon belongs in a book about physics for two reasons. First, it has every structural feature the other two have, and the next section shows that these shared features follow from the structure of compar-

ison. Second, physics is done by selves. Chapter 1 showed that measurement is interaction. An interaction has two sides, and one side of every measurement is a self that does not fully see itself. So the edge of what a self can know of itself is an edge of what any measurement, by any physicist, can report.

What the Three Share, and Why There Are Three

The three horizons are the edges at which the three limits on knowing of Chapter 9 stand: the light cone at the horizon of the very large, the uncertainty relation at the horizon of the very small, and the limit in time at the horizon of the self. The three horizons have three features in common.

All three move as one approaches them. The cone of a later event holds what the cone of an earlier one could not; better instruments resolve what could not be resolved before; deeper attention brings to light what was unconscious. Each is an edge that moves, and none of them is a wall or a place.

All three are permanent as classes. The kind of limit lasts, even as every particular instance of it dissolves. No later event has a cone containing everything; no instrument reaches a scale with no scale beneath it; no self completes itself. The difference between class and instance does real work here, because the *absolute boundary* between the unknown and the unknowable is set by classes. The *relative boundary*, the one a given observer meets from a given place at a given time, is set by instances. The relative boundary can always be crossed. The absolute boundary is never crossed, because crossing any instance of it brings one to another instance of the same class. Now what is in view changes, from the kinds of knowledge to the kinds of boundary between them, and the boundaries take modalities of their own. The absolute boundary has the character of the omniscient, perfect and unchanging; the relative boundary has the character of the transcendent, moving and exceeding; and the act of approaching a boundary, which in itself is neither absolute nor relative, has the charac-

ter of the immanent. The unknowable that lies past the absolute boundary keeps its own character, the transcendent; what is omniscient here is the boundary itself, which never moves.

The rule that which modality a thing is matched with depends on what one is considering applies once more when a third thing is put in view. The third thing is each modality taken as a point of view, and what that point of view holds fixed. From the omniscient point of view the frame of reference (the setting against which everything is placed) is held absolute; from the transcendent point of view, the material; from the immanent point of view, the relation itself. That matching concerns what each point of view holds fixed, and the one just made concerns the boundaries of knowledge; they match two different things, so the two stand together without conflict.

All three horizons mark the place where the unknown gives way to the unknowable. Short of the horizon lies what is not yet known and could be. At the horizon lies what could not be known from here, by any means, and the words "from here" are what make it a horizon rather than a wall.

Why three? Two of the horizons are objective, the large and the small; one is subjective, the surface of unconsciousness. That follows a theorem of tripling found all through the account of comparison, grounded in the *Law of Three*: a relation is always a relation of at least three. The theorem says that for every way in which something is subjective there are two ways in which it is objective, and for every way in which something is objective there are two ways in which it is subjective. The two halves say different things, and the count runs both ways: a subjective side comes with two objective sides, and an objective side comes with two subjective sides. The horizons show the first half: one subjective horizon, the self's, comes with two objective ones, the large and the small. That pattern of two and one appears at every level, and the three horizons are one case of it.

The principle comes from the six intrinsics of comparison (Chapter 1): sameness and difference, content and context, the subjective and the objective. Every content has a context, and every sameness has a difference;

and three is the smallest number of distinct parts that can meet these needs while none of them reduces to the others. Two things can be the same or different, but their sameness or difference needs a third from which to be seen. The subjective and the objective are the pair that supplies the seeing.

The third axiom, in the words of *An Immanent Metaphysics*, reads: "The (classes/instances) of the immanent, omniscient, and transcendent are distinct, inseparable, and non-interchangeable". Read aloud, the slash means "classes, and also instances": the classes of the three modalities are distinct, inseparable and non-interchangeable, and so are their instances. Because the three are inseparable, none of them occurs without the other two; so wherever one modality appears as an edge, the other two must appear as edges too. Now what is in view is the edges themselves, each by where it lies, and each horizon has one modality: the horizon of the very large is omniscient, the limit of the very small is transcendent, and the surface of unconsciousness, at the scale of the self, is immanent. Each of the three is a place where the unknowable begins, and the unknowable keeps its own character, the transcendent; what has one modality each is the edge, by where it lies. The count is exactly three. A fourth horizon would need a fourth modality, and there are three. And there are not only two horizons: with only two, one of the three modalities would have no edge. A modality with no edge would be one in which the unknowable did not appear, which is to say one in which everything is knowable, and the Incommensuration Theorem rules that out.

The count of three has a use beyond the horizons: it shows why absolute claims about knowledge trip over themselves, because each such claim works with two boxes where knowledge has three kinds. Take the saying "everything is permissible and nothing is true". As a working tool for exploring perception and loosening beliefs, it has power. Applied as a flat claim, it undermines itself. Ask whether the saying is itself absolutely true: if it is, then at least one thing is true, and the saying refutes itself. "Everything is knowable" runs into the unknowable. "Nothing is certain" runs into the certainty of the claim itself. Each of these traps comes from

working with only two boxes: true and false, absolute and relative. The way out is a third kind of idea, one that falls under neither truth nor falsity, a different kind of concept altogether: the immanent act of finding out. Only an account that holds the known, the unknown and the unknowable as three distinct kinds escapes these traps, and the axioms have their strength in that third kind.

Two Constants, Two Gateways

Now set the two physical horizons side by side. A *domain*, as Chapter 12 used the word, is the range within which one kind of comparison can be made; this universe is one domain, the range of physical comparison, and Chapter 19 gives the word its full meaning. To speak of the two horizons as the edges of any domain, and not of this universe alone, I give the constants that mark them general names: the *C constant*, for the one at the outer edge, which in our universe is the speed of light, and the *H constant*, for the one at the inner edge, which in our universe is Planck's constant. (The letters, read "see" and "aitch", are the ones physicists use for the two constants.) Chapter 19 develops that generality; here what matters is how the two stand to each other.

The C constant marks a limit from the outside. It says which positions are too far away for anything about them to be known from here. The C constant relates space to time by one ratio, the same throughout the domain, while space and time stay distinct. So each comparer has a horizon in space and time, and the history it pieces together from inside finds a beginning. The reason is the one ratio: whatever reaches a comparer from farther away left its source longer ago, so looking outward is looking back, and the history pieced together from what arrives runs back to an earliest time from which anything can have arrived. Past that horizon nothing has arrived, and nothing can. What the C constant bounds is the pattern of connection, which events can affect which, and that pattern is structure,

the same for every comparer, which is the mark of the omniscient. The C constant is therefore the omniscient limit of the domain: the boundary on the side of context, the farthest reach of the pattern of connection.

The H constant marks a limit everywhere on the inside. The H constant is present at every point, at the smallest scale of every measurement made anywhere, and the H constant is not far away at all. It says which distinctions are too fine for anything about them to be known. Below it, definite description is shut off, and what could be outruns what is, which is the mark of the transcendent. The H constant is therefore the transcendent limit of the domain: the boundary on the side of content, the finest grain of what can be told apart.

So the two are two different kinds of edge, and more than two numbers. One is reached by going outward; the other is reached by going inward, and from any point. Between them lies everything that can be known as content within the domain, which is to say everything that is physics. And the two edges, different in kind, have this in common: each is a passage from what could be known into what definitely cannot be known. The microscopic limit is a gateway of that kind, and so is the speed of light. Each is a place where the physics stops; neither is a place where the physics merely becomes hard.

The Incommensuration Theorem says what to expect at each gateway, and the two expectations differ. At the microscopic limit, where symmetry is asked for absolutely, where one demands the finest and most repeatable distinction, the theorem says one will find discontinuity: the graininess of the quantum in what is found, the separate exact value a quantity is allowed to take (which physicists call an eigenvalue), the outcome that cannot be refined further. At the macroscopic limit, where continuity is asked for absolutely, where one demands that the pattern of connection run on without a seam, the theorem says one will find asymmetry: the light cone that opens one way, the horizon crossed from one side only, the past that is fixed and the elsewhere that is not. The unknowable in physics sits at the extremes of scale for exactly this reason: the extremes are where per-

fection is demanded, and where perfection is demanded the theorem says a complete description is impossible. The framework says where these limits are, and it stays on this side of them.

Physics measures six kinds of quantity, the six measures that Chapter 1 set out as two rows of three, and that Chapter 19 sets out in full. Three measure the setting, and they are the context measures: space, time and possibility. The C constant bounds them, and at the large edge they must be continuous. Three measure what is found in the setting, and they are the content measures: pattern, force and probability. The H constant bounds them, and at the small edge they must be quantized, coming in whole units. The units of the two constants show the same placing as these bounds: the C constant on the side of context, and the H constant where content joins context. The speed of light is a ratio of two context measures, a distance over a time. Planck's constant is a content measure times a context measure: a momentum times a distance, or an energy times a time. So the C constant stands on the side of context, and the H constant joins content to context at the scale where content can no longer be divided. The uncertainty relation is the statement that a content and its context cannot both be fixed there beyond the bound the constant sets. I place the C constant on the side of context, and the H constant where content joins context. That placing, written in the units of the two constants, is my reading of what the two constants are constants of. In the grid of Chapter 19, which sets each measure beside its modality, Planck's constant is the bridge between the immanent and the transcendent, between time and possibility: it turns a rate in time into an energy, and an amount of action divided by it counts the states a system can be in, which is a measure of possibility. (Planck's constant is the room that one state takes up, counted in action, so a total amount of action divided by it counts how many states fit.)

Two kinds of word are at work here, and they point in two directions. "The side of content" and "the side of context" name the two rows of that grid: content above, what is found; context below, the setting. "Immanent" and "transcendent" name its columns. A row and a column are two

directions in one grid, and a box lies in both. So Planck's constant can bound the content row at its small edge, join a content to a context by its units, and stand, among the columns, between time in the immanent column and possibility in the transcendent, all at once and without conflict.

Linked Across a Distance

Write down everything about two towns: every street, every house, every person, the weather in each. You have still not written down how far apart the two towns are. The distance belongs to the pair, to the relation between them, and it sits in neither town.

A domain bounded by the two constants shows a property of that kind in one of the most discussed results in physics. Chapter 10 described *entangled* particles: two particles whose outcomes, measured far apart, agree in a pattern no separate instructions could give them. When two systems are prepared together and then separated, a measurement on one is found to agree with a measurement on the other more strongly than any account in terms of properties each carries separately allows. Albert Einstein, Boris Podolsky and Nathan Rosen described that situation in 1935 to argue that quantum mechanics was incomplete. John Bell's theorem of 1964, and the experiments that followed it, turned it into a result about the world rather than about the theory. Chapter 10 read that result as a case of the Incommensuration Theorem. Here one more side of it matters.

The result is commonly called a paradox, because it seems to break the rule that what happens here depends only on what is nearby: the measurement here seems to settle something there, across a separation no signal can cross. Within this framework it is a mark of what I call *domain curvature*, and no paradox at all. The two systems share a feature of the setting in which they were made together. That feature does not travel

between them, because it was never located in either of them. It belongs to the domain, and the domain is curved. The two systems send each other nothing.

I use the word curvature in a sense of my own. In this chapter three uses of the word 'curvature' now stand side by side. The first use is the curve of the earth, which was a picture of the traveler's horizon; the second use is the bending of space and time by which physics describes gravity; and the third use is domain curvature. The bending by which physics describes gravity is a different thing from domain curvature; domain curvature is a property of any domain as such. A domain is *flat* if everything about it can be given by giving what is at each point. A domain is *curved* if it has features that no point-by-point list captures: features of the relations among points, and of no point, as the distance between the towns belongs to neither town. Domain curvature, then, is the property by which no domain can be fully described place by place.

Chapter 12 showed that space is the pattern of which events can affect which, and that the light cone comes first. So the domain is a structure of relations before it is a set of places, and a structure of relations can carry features located nowhere in particular. Entanglement is such a feature. The agreements it produces exceed what *local causality* (the account in which every cause acts only on what is next to it) can account for, because local causality is the flat approximation to a domain that is not flat.

This reading fits both edges. The outer edge forbids signaling across the separation, and entanglement sends no signal: the results at one end cannot be steered from the other. The inner edge forbids fixing any content finer than the grain of the domain, and entanglement fixes nothing: each outcome, taken alone, is as random as anything at the small edge must be. What is shared is context, and context is not carried by content. The linking is then a direct expression of a structural principle, with nothing mysterious left to explain away: no finite domain is flat, and curvature shows itself as agreement beyond what any local account contains. The domain is a structure of relations before it is a set of places, as Chapter 12 showed for this one. My account of entanglement is a reading of known

physics, the second of the three kinds of claim that I name in the Preface. In that reading I say why the mathematics of quantum theory has the shape it has. The reading leaves the mathematics of quantum theory, and every prediction made with the mathematics, unchanged.

Where the Two Edges Meet

Picture a small cone of light drawn at every point of space near a black hole, each cone showing the directions in which a cause from that point can go. Far away, the cones stand upright. Nearer in, every cone leans toward the center, and the nearer, the more it leans, until at one surface every forward direction of every cone points inward.

The two physical horizons meet in one object, the black hole, and here the framework does more than read a known result: it says what physics will not find. Recall the step from Planck's constant to the Planck scale. Physics described the small edge in the terms of a black hole: the region one tries to measure closes on itself when the energy needed to measure it is enough to curve space around it. So physics itself describes the smallest scale of the domain in the language of the largest. A black hole is the object in which both edges appear together. The large horizon, the one-way surface, appears as the *event horizon*, the surface from inside which no signal gets out. The small limit, the point at which space and time cease to apply, appears as the *singularity*, the point at the center where, as general relativity describes it, the quantities of the theory become infinite.

The theorem forbids two pairings applied absolutely: symmetry with continuity, and asymmetry with discontinuity. The second forbidden pairing, asymmetry with discontinuity, is the pairing that matters here. Imagine a sharp wall that lets things through in one direction and never in the other, with nothing at all crossing back. That wall would be perfect asymmetry joined to perfect discontinuity, and the theorem says there is no such thing.

The event horizon, as general relativity describes it, is the leaning of the cones, and not that wall. Where the arena differs enough from place to place, the cone of every event tips inward, and the asymmetry shows as a one-way boundary. A falling body meets nothing local there; the geometry is smooth across the horizon, and the horizon's one-way character belongs to the whole pattern of what can reach what. The event horizon, smooth and one-way, is continuity with asymmetry, carried as far as that pairing goes, and the Incommensuration Theorem permits that pairing. Stephen Hawking showed in 1975 that a black hole gives off a faint glow. That radiation arises at and outside the horizon, and in the calculation that finds it, no signal sent from inside the horizon reaches the outside; so the horizon, as a boundary of the arena, stays one-way.

The singularity is another matter. As general relativity describes it, it is a point at which the quantities of the theory become infinite, a perfect discontinuity, and a point toward which everything falls and from which nothing rises, a perfect asymmetry. The singularity, as general relativity describes it, is a case of the second forbidden pairing, asymmetry with discontinuity. The theorem forbids perfect asymmetry joined to perfect discontinuity in any comparison, and the singularity, as described, is exactly that joining; so what is there cannot be as described. The framework is more general than general relativity, in this sense: general relativity is one of the two survivors of the theorem, applied to one domain, while the theorem applies to comparison as such. So the framework expects the singularity to be an idealization that will not be found in nature as described. There the smooth description reaches the edge of its own domain, and asking for the geometry at the singularity applies the idea of a smooth geometry where it does not belong. A perfectly broken and perfectly one-way point is incoherent as structure, by the same argument that makes pure nothingness incoherent. The singularity is one of the standing reasons physicists give for wanting a theory beyond what general relativity alone supplies. The claim that the singularity, as described, will not be found in nature is a structural exclusion, the third kind of claim the Preface names: it says what physics will not find.

What takes the singularity's place, and what a black hole is when described by a theory that honors both survivors in the way Chapter 17 described, is for physics to find. In that way, each theory applies the tool of one modality to the domain of the other, and the two are joined side by side and never merged. The framework names the structure of the problem, and finding its solution is the work of physics.

Where the New Comes In

Think again of the single radioactive atom of Chapter 8, which decays when it does, with no earlier state of the atom, or of anything else, fixing the moment. Something arrives in that click that was not there before. Where does it come in?

Creation, in this book, is the side of every event by which something new comes to be. It is real, but it does not exist: in the words of Chapter 4, it is in the actual act of comparison, here and now, and it does not stand as structure, the same for every comparer. It is an aspect of all events rather than an event, and it has no source, no location and no measure. Each of the following marks of creation says what creation is unlike, and each mark follows from what creation is. In the next paragraph I give the one positive mark of creation that physics can meet. Creation is potential, and actuality is pattern, so creation is not actual. Creation is the context from which content emerges, and observation yields content, so creation cannot be observed. Creation is transcendent, on the side of continuity, and conservation is a law of symmetry, so creation is not conserved. (Continuity is the transcendent's own character. At the small edge physics brings symmetry, the tool of the other modality, and so it finds discontinuity there; Chapter 19 walks through how the two fit.) The transcendent is asymmetric, and repetition is symmetric, so creation is not repeatable. And creation is cooperative across every boundary of self, which is why it has no source.

An Immanent Metaphysics gives one positive mark of creation that physics can meet. Creation, although it is not a process, typically involves the microscopic boundary of a domain. New possibility enters through the scale of the very small, but at no particular position in time or space. Said exactly: a domain has a boundary at its small end; what is transcendent to the domain reaches the domain across that boundary; and the reaching, seen from inside, has the look of the very small. The boundary is the place of entrance. The entrance itself has no place. So creation is not located at the small scale, since it has no location, and it is no tiny process, since it is no process at all.

Why the small edge, and not the large one? The large edge, marked by the speed of light, has the character of the omniscient: it is a limit from the outside, the edge of the structural description, beyond which the pattern of the domain ceases to apply. The small edge, marked by Planck's constant, has the character of the transcendent: it is a limit everywhere within, the edge below which definite description is shut off. Creation has the character of the transcendent, and what is transcendent enters through the transcendent boundary. The small edge is the surface through which creation comes in, and it is everywhere, because the very small is everywhere; it is no wall that keeps creation out.

An Immanent Metaphysics states the physical consequence directly. Quantum *indeterminacy*, the fact that a single quantum event's outcome is fixed by nothing before it, is the physical expression of the creative aspect of reality at the small edge. Single quantum events show almost no organization. Vast numbers of quantum events, taken together, show the order that makes up the everyday world. New possibility enters through the small and shows itself through the many. Three steps make that plain: what indeterminacy is; why it must look as it does from inside the domain; and how the many turn it into order.

First, what it is. Chapter 8 showed that there is a residual in what happens, a part that no prior state accounts for. The reason is that there is nothing prior for the residual to follow from, and not merely that the prior state is unknown. The residual is what anything arriving through the

small edge must look like from inside the domain. Here it gets its name. Seen from the omniscient side, the side of pattern and prediction, what comes through the edge is hard randomness: what no shorter description can hold, what nothing accounts for. Seen from the transcendent side, the side of possibility and the new, what comes through is creation. They are one thing under two descriptions, and the descriptions differ because the modalities they are made from differ. Randomness is the name the world gives to what the self would call novelty.

Second, why it must look like indeterminacy. The domain is a domain of comparison, and a comparison needs a content that is the same or different, within a context or across contexts. What is new has no earlier content to be compared against. It cannot be predicted, since prediction carries a known pattern forward, and the new is exactly what the pattern does not contain. It cannot be observed as a finished object, since observing turns context into content, and what has been turned is actual, no longer potential. So at the boundary where the new comes in, what one finds is an event whose outcome nothing within the domain fixes. An event whose outcome nothing within the domain fixes is what 'indeterminate' means. The indeterminacy of the quantum event is the mark of something arriving that the domain did not contain; it is no failure of the domain to settle anything.

Two familiar readings of quantum indeterminacy stand beside this one. The first holds that it is ignorance: there is a definite value, and we do not know it. Chapter 8 showed why, after the experiments that closed the loopholes, that reading asks more of the world than the alternative does. The second holds that the outcome is a bare random draw, with no further character. My reading agrees that there is no hidden value, and agrees that no prior state fixes the outcome, and it holds that what arrives has a character: the character of the transcendent. It is potential becoming actual at the one place in a domain where potential can become actual without having been actual before: the boundary.

Order From the Many

Chapter 8 drew the picture already. No one atom's moment of decay is fixed, and yet a large batch of the same atoms keeps its half life, the time in which half of them decay. The order is in the batch, and in no one atom.

An Immanent Metaphysics says that the strength with which creation shows itself is measured by how organized and coherent large numbers of tiny events are, taken together, and that any small, finite handful of such events shows less of possibility than the whole crowd does. That has an exact form in the second axiom, which Chapter 4 quoted: "A class of the transcendent will precede an instance of the immanent. A class of the immanent will precede an instance of the omniscient. A class of the omniscient will precede an instance of the transcendent". In plain words: many possibilities come before one act; many acts come before one lasting structure; many structures come before one new possibility.

Read into physics, following that one-way cycle from the omniscient to the transcendent to the immanent and around again: many large structures, taken together as the known world, come before one meeting with what exceeds them; many small possibilities come before one actual interaction; and many actual interactions come before one large structure. The large object, the thing that exists, that lasts, that can be measured again and again, is an instance of the omniscient. The axiom says a class of the immanent comes before it, a multitude of interactions; and a class of the transcendent comes before each of those, a multitude of possibilities. The large is built on the small as a pattern is built on a crowd of events, and not as a house stands on its foundation.

Because the large is built on the small as a pattern is built on a crowd of events, the single quantum event shows almost no organization and the many quantum events show order. A single event is one instance drawn from a class of possibilities; there is nothing yet for it to be coherent with. A large number of events, interacting with one another, make up the class of the immanent from which an omniscient instance, a structure, can arise. Coherence belongs to the many, and not to the one. And here the reading

runs against a common hunch. The order of the everyday world is creation showing itself, and not creation held down. The strength of creation is measured by how organized the crowd is, not by how disorganized the single event is. A world of quantum events with no connection among them would not be a highly creative world: it would be a world into which possibility had entered, and of which nothing had been made. Creation enters through the small, and it shows itself in what the many small events together become.

Physics gives the numbers of this order: the odds of quantum events, the form of Born's rule, and how fast the coherence of many events builds into everyday order. It has a full account of them under the name *decoherence*: the way a system's many interactions with its surroundings turn its quantum alternatives into the settled order of everyday things. The reading says why there is an order of the many at all, and why the single event does not show it: because the large is an instance of the omniscient, which the second axiom places after a class of the immanent, which it places after a class of the transcendent. The reading gives the reason there is an order of the many at all, and physics gives the numbers of that order.

Why Position and Momentum Cannot Both Be Sharp

Clap once in a quiet room. The clap is sharp in time, but ask what note it was and there is no answer: a sharp click holds a smear of every pitch. Hum a steady note, and the pitch is sharp, but the note has no sharp moment, because the note lasts. A sound can be blurred in both, or sharp in one, and never sharp in both at once.

The uncertainty relation has the same form, and it is a case of the Incommensuration Theorem. Recall the theorem's two forbidden pairings: applied absolutely and together, symmetry cannot be joined to continuity, and asymmetry cannot be joined to discontinuity. Within any finite domain the members of each pair occur together, each to a finite degree; what they cannot do is both be perfected at once. Heisenberg's relation

says that the spread in a particle's position, times the spread in its momentum, cannot fall below an amount fixed by Planck's constant. Narrow the one and the other widens; neither can be made exact while the other stays finite. In Chapter 8 I read Heisenberg's relation as a limit of access, and earlier in this chapter the number in Heisenberg's relation marked where the small edge lies. Here the relation itself is the subject: the fact that position and momentum in particular cannot both be perfected is a case of the first forbidden pairing. Three routes lead to that reading, and they agree.

The first route is the trajectory. A body with a definite position and a definite momentum at every instant has a *trajectory*, an exact path. A trajectory given at every instant with no gap, under one law that holds everywhere, is continuity and symmetry perfected together in what is found, and that is the first forbidden pairing. Quantum theory keeps what it finds exact, in whole steps with exact labels, and gives up the trajectory. The giving up shows as the bound on how sharply position and momentum can be known together, and Planck's constant sets the size of that bound.

The second route pairs the two survivors of Chapter 3. Physics ties momentum to symmetry more directly than any other quantity. Emmy Noether proved in 1918 that momentum is conserved because the laws are the same when everything is shifted from one place to another. A state of exactly definite momentum is the same at every position, apart from the timing of its crests; physicists call it a *plane wave*. That state is symmetry of momentum with discontinuity of position: it has given up being anywhere in particular. A state confined to one place is continuity of position, a sameness of the spatial setting, paired with asymmetry of momentum, a difference in its motion. The relation pairs the two survivors, and never the forbidden pair.

The third route is the form of the relation. Position and momentum are a conjugate pair under the Fourier transform, the mathematics that rewrites a description by place as a description by waves of different lengths, as a chord can be rewritten as the notes that make it. That form is

the pairing of the omniscient and the transcendent in the first axiom, which reads, in the words of *An Immanent Metaphysics*: "The immanent is more fundamental than the omniscient and/or the transcendent. The omniscient and the transcendent are in conjugate relationship to one another". The first axiom itself follows as a theorem from the Incommensuration Theorem, so the uncertainty relation is an outcome of that theorem. Separately, a state that is the same across every place cannot be confined to one place: that would be symmetry against asymmetry. The second of Chapter 3's considerations already excludes it: a character and its opposite exclude each other when both are made perfect.

The reading gives the reason for the pairing, which is no accident: position and momentum cannot both be known exactly together because a trajectory would perfect continuity and symmetry at once, and the theorem forbids perfecting that pair. Physics gives the rest: that there is a bound, the size of the bound, and the mathematical form of the relation. The size of Planck's constant is a number of physics, and its number expresses a unit. The reading extends as the physics extends: every conjugate pair, energy and time, or angular momentum (the amount of turning) and angle, pairs a sameness across settings with an exactness within one, and each carries the same bound for the same structural reason.

What the theorem permits teaches as much as what it forbids. The two pairings that can hold absolutely are continuity with asymmetry and symmetry with discontinuity. A *wavepacket*, a wave bunched in one region, with a definite enough position and a phase that varies from point to point, is continuity with asymmetry. The plane wave, or the single countable quantum of one definite mode (one definite pattern of vibration), the same everywhere and countable in its steps, is symmetry with discontinuity. These two pairings are the wave and the particle of Chapter 3, the two faces of quantum matter. The uncertainty relation says that physics can have either survivor perfected, and never the forbidden pair, and that between the two survivors runs a smooth range of finite trade-offs. The

smooth range of finite trade-offs between the two survivors is exactly the Incommensuration Theorem's own condition: within a finite domain, symmetry and continuity both hold, each to a finite degree.

The Diamond and the Coal

The diamond and the coal are both carbon. The coal is simpler in structure and lets no light through; the diamond is more organized and lets nearly all of it through. The diamond, not the coal, is the image of a reality that lets in the new.

An Immanent Metaphysics adds a mark that says which physical systems show creation most. To be clear is to allow creation. How strongly creation shows varies with clarity and transparency, with emptiness, with possibility, and it varies inversely with simplicity. That needs a distinction often missed. A *simple* system has few parts and few relations, and so a simple system is easy to describe completely. A *clear* system is one that passes things on: what enters it can pass through and be made into something.

Read into physics: a simple, closed, fully described system, such as one atom held alone in a trap with only two states it can be in, is the system in which physics comes nearest to exact prediction, and it is also the system in which creation shows least. Possibility enters it, since it is a quantum system and has a small edge like any other, but there is nothing for what enters to become. A complex, open, interconnected system, with many parts in many relations, in exchange with its surroundings, is the system in which physics predicts least exactly and in which the new shows most. The framework says these are one fact and not two: the system that is hardest to describe is the system that passes on the most, and what it passes on is what came in through the small.

That complex systems are hard to predict is a commonplace of physics, and it needs no metaphysics to observe. What the framework adds is the reason, that clarity and creation vary together, and I offer that reason as an interpretation. It is a reading at the level of interpretation, the second kind of claim the Preface names, and it changes no prediction of physics.

How Physics Meets the Three Mysteries

Physics meets the three mysteries in three ways.

The first way is the shape of inquiry itself. Seen from the impersonal, objective point of view, the known has the character of the omniscient, the unknown of the immanent, and the unknowable of the transcendent. The one-way cycle of the second axiom then says: many things known come before one meeting with the unknowable; many meetings with the unknowable come before one fresh engagement with the unknown; many engagements with the unknown come before one new thing known. Knowledge piles up, and at some depth the piled-up knowledge shows an edge that no further piling will cross; meeting that edge sends inquiry back to the unknown, to what is not yet measured, with new engagement; and from that engagement new knowledge comes. The uncertainty relation was such a meeting. A great body of known physics, mechanics, optics, the sharp colors that each kind of atom gives out, came before the recognition that something could not, by its structure, be pinned down; and that recognition redirected inquiry instead of ending it.

From one's own, first-person point of view, the roles shift. The known becomes immanent: what I am in touch with, here and now. The unknown becomes omniscient: the vast territory I have not explored, which is there whether or not I explore it. The unknowable stays transcendent. What the three are stays the same; the known and the unknown exchange roles with the point of view, and the unknowable keeps its place. That exchange of two roles is a flip, carried by the first axiom, and not a step around the cycle. Run through the second axiom, it would turn the cycle

backward, so the cycle of the second axiom is read from the impersonal point of view only. This exchange is one case of the rule set out at the start of this chapter, which the framework holds throughout: which modality a thing is matched with depends on what one is considering, and is no label fixed to the thing. For physics the practical consequence is small and real. What physics as a discipline knows, the omniscient body of results, is for any one physicist mostly unexplored territory, and what any one physicist knows is the part of that territory they are in touch with. The discipline and the person do not have the same known.

The second way concerns what the small edge is, as a mystery. *An Immanent Metaphysics* calls it the mystery of the infinitesimal: the scale at which knowledge of a definite state or quality becomes impossible, and it names the uncertainty limit as its physical expression. Below that scale, structure itself shuts out knowledge of structure, so that knowledge of structure is more than missing there. The scale below which structure shuts out knowledge of structure is the *hard boundary*, the boundary between the unknown and the unknowable. The hard boundary is hard for the reason given above: what comes through the hard boundary is what the domain did not contain, and no amount of what the domain contains will reach what comes through. The *soft boundary*, between the known and the unknown, is the one physics moves every year. The hard boundary is the one it does not move; as instruments improve, the instance recedes, and the class stays where it is.

The third way concerns what the hard boundary does to the one who meets it. *An Immanent Metaphysics* says that mystery is neither conserved nor diminished by any degree of knowing, understanding, reason or science. A physicist who understands general relativity is no less awed by the universe than a shepherd who knows nothing of physics. The content of the awe changes, and the awe itself does not shrink. If anything, deeper knowledge brings deeper awe, because what is known shows more sharply and more vividly the outline of what is not known and cannot be. I take that to be a plain description of what happens to those who work at the edges of physics, and no consolation offered to them. The reason is

structural, and *An Immanent Metaphysics* calls it the *veil of secrecy*: the omniscient cannot contain or observe the transcendent. Knowing is itself an interaction, and so immanent, and from within the omniscient one cannot know the transcendent. The transcendent perceives the omniscient, and never the reverse. Every increase of the omniscient, every result added to physics, widens the surface across which the transcendent is met, and takes nothing at all from what remains on the other side.

So at the hard boundary, knowledge turns, and it does not end. Many things known reveal one thing unknowable, the meeting with that one thing unknowable returns inquiry to the unknown, and from the unknown new knowledge comes. The physicist and the shepherd stand before the same mystery. The physicist sees its outline more clearly. Seeing the outline of the mystery more clearly is all the difference between the physicist and the shepherd, and that difference is no small one.

Found or Made?

Physicists ask a question about their own laws. Was the uncertainty relation found in nature, or made by the minds that framed quantum mechanics? Draw a circle on the ground around where you stand, and call everything inside it you. Draw it tight around your body, and the laws you write down are found in the world outside the circle. Draw it wide enough to take in the whole world, and every law is your own making. The answer depends on where you draw the line.

The question, as posed, has an answer only once a choice is made, and the framework says which choice. Whether something is called created or discovered depends on where the boundary of the self is placed. Take the whole of being and divide it into self and not-self. Where the line falls decides what one will say. Take all of being as subjective, the position called idealism, and one will say one has created everything, the laws of physics included. Take all of being as objective, the position called realism, and one will say one has created nothing, and discovered whatever one has.

Neither extreme holds; Chapter 6 found the same of realism and idealism among the errors of extremes, where each pole, taken alone and absolutely, goes wrong. Whether a thing is created or discovered depends on where the boundary is drawn, and it is no feature of the thing itself.

For physics this has a definite consequence. A law of physics is a symmetry, a sameness of content across a difference of context, and a symmetry is omniscient: it is the same for every observer, and that is what makes it a law. So the law, as a law, lies beyond any boundary of self that a single physicist could draw, and in that respect it is discovered. But the framing of the law, the choice of what to compare with what, the choice of position and momentum rather than some other pair, the decision that this regularity and not that one is the one worth stating, is an act of the transcendent. Creation is cooperative across the whole community that made the framing, and in that respect the law is created. The law is discovered as a symmetry, and created as the comparison in which the symmetry was found. Each description holds in its own respect, and neither is complete, because the boundary of self that would complete one of them cannot be drawn. The answer leaves every physical claim as it stands; what it gives a physicist is the right way to speak of a law: found in one respect, and made in another.

What We Now Know

We now know the answer to the question this chapter began with. There is a limit to what physics can ever find out, and it is structural: it is set by what comparison is, and not by what has been found so far. Physics has two horizons, and the knowing self has a third, and each of the three is permanent as a kind.

Most of what this chapter has said belongs to the ground, the first kind of claim the Preface names; its readings of known physics, and its one exclusion, the singularity, are marked where they are made; and the Appendix lists every claim of the chapter with its kind. I mark where the unknowable begins, and I do not enter.

The three horizons matter for the question of this book, what cannot be and why. Chapter 12 found that a signal faster than light asks for a comparison beyond the outer edge of comparison. We now know what that edge is: one of three, each the place where the knowable gives way to the unknowable. What lies past a horizon lies outside the domain, however near or far it seems, and no instrument and no amount of effort brings it within the domain.

The unknowable gives knowledge its direction, and so the unknowable is no enemy of knowledge. The edges are horizons to be approached, and not walls to be torn down, and in the approaching, knowledge deepens. Without the horizon, inquiry would have no shape. The framework leaves the horizon in place. It says where the horizon is, and why it is there, and that it will always be there; and that is the shape of knowing, and no defeat.

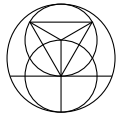
Summary

Physics has two edges, and the knowing self has a third. At the largest scale stands the speed of light, beyond which nothing arrives; at the smallest scale, the limit Planck's constant marks, below which nothing is distinguished jointly; at the scale of the self, the surface of unconsciousness, past which the self does not see. Each edge is a horizon: an edge of what can be known from here, which moves back as it is approached. Any one horizon, an instance, can be passed; the kind, the class, is permanent; and the absolute boundary of the knowable is set by the class. The unknow-

able is what no observer can know at any time by any method, because comparison does not reach it, and the three horizons are the three places where the knowable gives way to it, one for each modality, exactly three.

The outer horizon is marked by the speed of light, the C constant of the domain: a limit from the outside, on the side of context, where the theorem says to expect asymmetry, and where physics finds the light cone. The inner horizon is marked by Planck's constant, the H constant: a limit everywhere within, on the side of content, where the theorem says to expect discontinuity, and where physics finds the quantum. Between them lies everything that can be measured. Planck's constant is a number expressing a unit; it sets no least length and no least duration, and how fine a comparison can be made depends on the energy brought to it. The Planck length is a bound on findings, not a least step of space and time. Through the inner horizon the new enters every domain; from inside, its entrance looks like an outcome that nothing within the domain fixed; and the order of the everyday world is what many such events together become. The uncertainty relation is a case of the Incommensuration Theorem: physics supplies the size of the bound, and the theorem supplies the reason there is one. Where the two physical edges meet, in the black hole, the event horizon is one-way by the smooth leaning of every light cone, which the theorem permits, and the singularity, perfectly broken and perfectly one-way, cannot be in nature as it is described. In this chapter I say where the unknowable begins and why the unknowable begins there. The values of the two constants are numbers of physics, and what the two constants are, and why they have the sizes they have, is the next chapter's subject.

If the speed of light and Planck's constant mark the two edges of every measurement, what exactly are those two numbers, and why are they the size they are?



Chapter 19

The Two Constants That Bound Every Domain

A Floor and No Ceiling

Why is the speed of light the speed it is? And what is Planck's constant, the other number that every book about physics seems to mention, actually for?

Start with a dark room. Close the curtains, block the gap under the door, switch off the last small light on the charger, and the room grows darker, and darker still, until it cannot grow any darker at all. Darkness has a floor. Now open the curtains at noon, switch on every lamp, bring in a floodlight, and then a hundred more. The room grows brighter, and there is no point at which brightness has to stop. Brightness has no ceiling.

Every basic measure is built this way. An amount of light, an amount of sound, an amount of anything that can be measured at all, has a floor at zero and no top. A minus sign in a measurement, as when five meters to the left is written as minus five meters, says which way something lies; it never says that there is less than none of it.

The question of why the two constants have the numbers they have has a short answer. The speed of light is a little under three hundred thousand kilometers each second, and it is exactly two hundred ninety-nine million, seven hundred ninety-two thousand, four hundred and fifty-eight

meters each second, because since 1983 the meter has been defined by it. *Planck's constant* is what physicists call the quantum of action: the smallest package of process, measured as an energy multiplied by a time (or, which comes to the same, an amount of motion multiplied by a distance). By itself it sets no smallest length and no shortest time; how fine a comparison can be made depends on the energy brought to it. Since 2019 its number has been fixed in the same way as the speed of light, by the definition of the kilogram. So the numbers of the two constants say how large a meter and a kilogram are, and there is no fact about their values to explain. The numbers of the two constants are set by definition, and that is plain. What this world measures are the pure numbers, the ratios that come out the same in any units, and the two constants leave those untouched. One pure number, for example, is how many times heavier a proton is than an electron: about one thousand eight hundred and thirty-six, a measured figure, the same whatever units the two masses are weighed in.

The question of what the two constants are for is the subject of this chapter. The two constants mark the two edges of the range in which physics can compare anything at all: an outer edge, the farthest reach, and an inner edge, the finest grain. In this book I make claims about the two constants as edges. I quote the numbers of the two constants from physics, where those numbers are set by definition. I show why every range of comparison must have two such edges, and why the two edges are of two different kinds. In this chapter I go in three steps: first what a domain is, and the six things every measurement needs, with the grid that the six things make; then the two constants that bound every domain. The last step is what the two edges mean for the units, for zero, unit and extent, for mass, energy and light, for interaction, existence and creation, and for the kinds of equation that physics writes.

What a Domain Is

Chapter 12 used the word domain for the range in which comparisons can be made, comparison being the act of finding whether two things are the same or different, in some respect. The word now gets its full meaning. A *domain* is everything that a given kind of comparison can compare: all the contents that a given kind of interaction can bring into relation with one another, and all the contexts, the settings, in which that interaction can take place. Each kind of measurement brings its domain with it. Wherever there is measurement there is a domain, and the domain is whatever the measurement can reach. So a domain is a reach, and never a region of space, a collection of objects, or a set of laws.

What people call "this universe" is one domain: the domain of physical measurement, everything that physical interaction can compare. Whether it is the only such domain, or a special one, or one among many, the idea of a domain does not decide, and this chapter has no need to decide it. What matters is that the framework applies to every domain whatever: physical, mathematical, experiential (the domain of what is lived and felt), and imaginal (the domain of what is imagined). Each domain has a kind of comparison of its own, and so each domain has the structure that every comparison carries.

The structure that concerns us here is the structure of a domain's limits. Chapter 18 showed that what can be known has horizons, edges that no advance in knowledge moves, and that physics has two objective ones, an outer and an inner. The claim of this chapter is that every domain has two such edges, that the two are of different kinds, and that the distance between them is the single most important fact about the domain's structure. To see why the two differ in kind, we first need the six measures that every measurement uses.

Six Things Every Measurement Needs

Picture a bowling ball rolling down a lane toward the pins. There is the ball itself, heavy and solid, which keeps rolling once it is set going. There is the lane, the place where the ball is at each moment. There is the push of the bowler's arm, and the knock of ball on pin. There are the seconds the roll takes. There are the odds of a strike, given what is known of this bowler. And there are all the ways the ten pins could fall.

The things in the picture of the bowling ball, taken in order, are the *six measures*: mass, space, force, time, probability and possibility. Every idea of an act of measurement involves all six, and every measurement uses all six, even when it names only one of them. They are the six that physics handles: every physical quantity is a quantity of one of these six kinds.

Each needs a word of its own. *Mass*, in this sense, carries several meanings at once: inertia (the resistance of a thing to being set moving or stopped), pattern, form, shape and structure. Mass is what lasts, what resists change, what a measurement finds to be there. *Space* is the setting in which pattern is placed: here rather than there, this arrangement rather than that. *Force* is the coupling by which one thing engages another: how strongly change goes, and which way. *Time* is the medium through which change proceeds: before and after, the order in which a measurement unfolds. *Probability* is the weight given to each outcome on what is known: how much it is to be expected. *Possibility* is the range of outcomes, what could be the case, the alternatives among which any actual outcome counts as one.

Take away any one of the six, and the act of measuring falls apart. No pattern, and there is nothing to measure. No space, and there is nowhere for it to be. No force, and nothing engages anything. No time, and nothing happens. No probability, and nothing is weighed. No possibility, and there is nothing to weigh. That removal test is the first part of the argument that the measures are six. My book Comparison; The Necessary Basis completes it, by deriving the count as three roles, each taken twice: once as content and once as context.

The six measures are one of three sets of six that bear on physics, and it helps to keep the three apart. The first set is the six intrinsics of comparison, the six parts no comparison can do without (Chapter 1): sameness and difference, content and context, the subjective and the objective. The six intrinsics of comparison are the root, and every other structure in this book is a case of the six intrinsics of comparison. The second set is the six *intrinsics of interaction*: the six aspects into which, as my book An Immanent Metaphysics holds, any interaction between a self and the world finally resolves, exactly six, each built into interaction itself. The six intrinsics of interaction are the six intrinsics of comparison as the intrinsics of comparison appear when comparing is taken as an event, a happening between a self and a world, and not only as a relation among contents. The third set is the six measures.

The three sets are joined in this way: each later set is the earlier set described in one further way. Comparison, interaction and relation are one act, named three ways (the Root Tautology of Chapter 1), and interaction is that act described as an event. Measurement is interaction described in one further way: one side is taken as the measurer, the other as the measured, and what passes between them is taken to stand for what is measured. So the six measures are the six intrinsics of comparison read through two descriptions: comparison described as an interaction, and interaction described as a measurement. The sections below check each of the two readings on its own. With both readings checked, the six measures are not a new structure: they are the six intrinsics of comparison, seen from where physics stands. When I call the six measures the physical face of the six intrinsics, I mean that the six measures are the six intrinsics of comparison, seen from where physics stands.

The six measures can be matched with the six intrinsics of comparison in two ways, depending on what is in view. With comparison in view, as a relation, sameness goes with mass and difference with space; the objective goes with force and the subjective with time; content goes with probability and context with possibility. With measurement in view, the same six are matched by the role each plays: mass and space with content and con-

text; force and time with sameness and difference; probability and possibility with the subjective and the objective. The two matchings are equals, one a rotation of the other, and neither corrects the other. They answer two different questions, which is the rule Chapter 18 set out: which match a thing gets depends on what one is considering, as one woman is a mother to her son and a daughter to her own mother. The first matching pairs each measure with one intrinsic; the second pairs each column of two measures, a content with its context, with one pair of intrinsics. I look at measurement, so I work with the second matching. The rows of the grid in the next section are a third thing again: within every column, one measure is what is found and the other is the setting in which it is found.

The Grid of the Six Measures

The six measures fall into two groupings, and the two groupings cross each other, so the six sit in a grid of three columns and two rows.

The columns are the three modalities of Chapter 4. Mass in space is the column of the omniscient, structure, what holds the same for everyone; its character is the structural, the repeatable, existence and pattern. Force in time is the column of the immanent, process, what is happening here and now; its character is the processual, the interactive, change. Probability in possibility is the column of the transcendent, what makes things possible and new; its character is the evaluative (what weighs and values), the potential, choice and possibility. Each column is a content in a context, and I write them in that form on purpose: mass in space, force in time, probability in possibility.

The rows are content and context. An Immanent Metaphysics sorts the six measures into those that are the content of a measurement and those that are its context. Pattern (which takes in mass), force and probability are content: what is found. Space, time and possibility are context: where, when, and among what alternatives it is found. In a physical measurement the general idea of content is made up of the three particular ideas

pattern (mass among them), force and probability, and the general idea of context is made up of space, time and possibility. An Immanent Metaphysics calls each row a *triaxial*: three axes that together make up one kind.

	Omniscient (structure)	**Immanent (process)**	**Transcendent (possibility)**
Content	mass	force	probability
Context	space	time	possibility

Said in words, the grid has three columns, one for each modality, and two rows, content above and context below. The first column, the omniscient, holds mass as its content and space as its context. The second column, the immanent, holds force as its content and time as its context. The third column, the transcendent, holds probability as its content and possibility as its context. Read across the top row and you meet mass, force and probability. Read across the bottom row and you meet space, time and possibility. Every measurement fills all six boxes, even when it names only one of them.

The grid shows something that neither grouping shows alone. The columns say what kind of thing each pair is: structural, processual, or a matter of weighing. The rows say what part each measure plays inside its pair. And because content and context are two of the six intrinsics of comparison, the rows are not a sorting brought in from outside: they are that second pair of intrinsics, showing up inside the six measures. The other two pairs of intrinsics, sameness and difference, and the subjective and the objective, show up in the same way, as the next section shows.

The Six Parts of Interaction

An Immanent Metaphysics states the six intrinsics of interaction in a fixed order, and for each it gives the modality and whether it is a content or a context. Its own claim is that the six intrinsics of interaction match the six measures of any physical domain directly, one for one, because every measurement, as in the method of science, is necessarily an interaction. I take the six intrinsics of interaction in the order of An Immanent Metaphysics, and with each intrinsic I give the measure that the intrinsic matches.

The first intrinsic of interaction is the distinction of subject and object: of what is patterned and what is not, of form and feeling. It is an omniscient context. In every interaction there is always a difference, a break, between the subjective and the objective. It matches space. Space is the measure of that separation: the measurer is not the measured, and the measurement is made across a distance, even when the distance is only the thickness of a boundary. Subject and object are separate from each other and yet in balance: each is what the other is not, and neither is ranked above the other in the distinction itself. Space carries both features. Two points in space are distinct, and the relation between them is the same seen from either end.

The second intrinsic of interaction is the fixed character of the object's structure, what An Immanent Metaphysics calls the platonic eternity: the part of the object that does not change. It is an omniscient content. In every interaction there is always some degree of objectivity, a symmetry, with fixed characteristics. It matches mass, in its full sense of inertia, pattern, form, shape and structure. Mass is what stays the same through the interaction: the mass that resists being sped up, the pattern recognized from one meeting to the next, the structure that one measurement finds and a second measurement finds again.

The third intrinsic of interaction is that change cannot be undone. It is an immanent context. In every interaction there is always a one-way character on the subject's side, an asymmetry. It matches time. Time is the set-

ting of change, and its defining feature is that before is not after, and cannot be made so. An Immanent Metaphysics says that this one-way character is met within continuity: change is continuous, and within that continuity there is a direction that does not reverse.

The fourth is the direction in which change flows. It is an immanent content. In every interaction there is always a flow of change, with a direction and some strength. It matches force. Force has a direction and a size; it is the content of change: what flows, which way, and how strongly.

The fifth intrinsic of interaction is the transformation of information. It is a transcendent context. In every interaction detail is always being taken away or put in. It is taken away as the interaction moves from the world toward the self (as when a glance reduces a whole street to "busy"), and put in as it moves from the self toward the world (as when a plan becomes a built house). It matches possibility. Possibility is the range within which detail is put in or taken away: the range of what could be the case, narrowed by taking away to what has been perceived, widened by putting in to what has been specified. It is the context of the transcendent column because it is the context of the transformation itself.

The sixth is the specification of what the transformation puts in: detail that is not local and not determined, which An Immanent Metaphysics calls *gauge values*. 'Gauge values' is a term of *An Immanent Metaphysics*. In Chapter 20 I use the word 'gauge' in the physicist's sense, which is a different use. It is a transcendent content. In every interaction the details put in or taken away are always specified. It matches probability. Probability is the weighting of the alternatives so specified: what is said about an outcome that no law fixes in advance, the content that the transformation adds beyond what the symmetric and continuous structure fixes. An Immanent Metaphysics calls this content not local and not determined for a reason: probability is the measure through which the transcendent, the creative, the chosen, enters measurement.

The distinction of subject and object goes with space; the fixed character of the object, with mass; the one-way character of change, with time; the direction of its flow, with force; the transformation of information, with possibility; its specification, with probability. Omniscient context and content, immanent context and content, transcendent context and content: the grid of the last section comes back exactly, column for column and row for row.

An Immanent Metaphysics adds one more layer. The first two intrinsics, distinction and fixed character, are defined in terms of causality. The middle two, the one-way character and the direction, are defined in terms of change. The last two, transformation and specification, are defined in terms of choice. Causality is omniscient, change is immanent, choice is transcendent (the triple of the real, Chapter 7), so the three columns of the grid are the columns of causality, change and choice. Physics as it is ordinarily done works the first two columns, and it handles the third through probability alone, leaving possibility unspoken, as the range over which its probabilities are spread. That the third column is where choice enters is, I think, the single most consequential thing the grid has to say. Chapter 20 uses it, where the two questions physics leaves aside, who and why, turn out to be the two that ask for this column, and my book The Structure of Effective Choice develops it further.

Two Rows, Two Laws

Chapter 2 showed that every law is a law of symmetry or a law of continuity, or a compound of the two. The grid says which kind of law governs which row. Where mass, force and probability are the content of a perception, they are defined by symmetry law. Where space, time and possibility are its context, they are defined by continuity law. The content row is the row of symmetry; the context row is the row of continuity.

The definitions of symmetry and continuity lead one to expect that the content row is the row of symmetry and the context row is the row of continuity. *Symmetry* is sameness of content across difference of context: change the setting, and the content stays the same. A law of symmetry is therefore a law about content, about what stays the same as the setting varies. The conservation laws, the laws of amounts that stay the same through every change, are the physics of the content row. Mass is conserved together with energy, as one currency. Force is conserved as momentum, the amount of motion: a force is the rate at which it changes a body's momentum, and whatever amount of motion one push gives to one side, the other side loses, so the total stays the same. Probability is conserved as the rule that the probabilities of all the outcomes always add up to exactly one. *Continuity* is sameness of content within sameness of context: hold the setting fixed, and nothing jumps. A law of continuity is therefore a law about context, about the setting staying the setting, with no gap in it. The physics of the context row is the physics of the unbroken: space with no absolute gap, time with no absolute pause, possibility with nothing absolutely ruled out in advance.

Now bring in the Incommensuration Theorem of Chapter 3, which is a structural fact about the nature of knowing. When two of the four compounds (continuity, discontinuity, symmetry and asymmetry, which Chapter 3 also called the four characters) are applied absolutely and together, symmetry with continuity is forbidden, and so is asymmetry with discontinuity. Continuity with asymmetry, and symmetry with discontinuity, are the two pairs that survive. Within any finite domain, symmetry and continuity are both present to some degree; what they cannot do is both be made perfect at once.

An Immanent Metaphysics draws the consequence for the grid in three steps. First, whatever is known or defined as symmetric must also be discontinuous, since the theorem forbids anything to be held as both wholly symmetric and wholly continuous. Second, discontinuity has the same form as quantization, coming in whole units that cannot be split; so the idea of content goes with being quantized, and the idea of context

goes with not being quantized. Third, at the absolute microscopic scale of a domain, its very smallest, the mass, force and probability of a perception must be quantized: symmetric and discontinuous. At the absolute macroscopic scale of a domain, its very largest, the space, time and possibility of a perception must be continuous: not quantized, and asymmetric, running one way.

So the two rows of the grid are the two survivors of the theorem. The content row, at its limit, is symmetry with discontinuity. The context row, at its limit, is continuity with asymmetry. Neither row can be both at once. The grid as a whole is therefore two structures, and never a single structure that one description could make perfect. The two are held together only in the interaction, which fills all six boxes, and never in one description that finds all six perfect.

The words "at the absolute scale" set the scope of the claim. The claim is about limits, and not about every scale. At the human scale mass looks continuous and space looks divisible, and both are true to a finite degree. Push a domain to its smallest edge, and its content measures must show their discontinuity; push it to its largest edge, and its context measures must show their continuity. The very small belongs to the transcendent and the very large to the omniscient, and each is where the tool proper to the other is used. Chapter 17 called that the cross-modal reach of the two great theories, and the next paragraph shows it at each edge.

This fits the modalities' own characters. The omniscient, in its own character, is the quantized and the large, and anything actual comes in units; the transcendent, in its own character, is the unquantized and the small, and anything potential is continuous. At each limit, what physics finds is what the tool that physics brings to that limit can find. At the smallest limit physics brings symmetry, the omniscient tool, and finds the separate outcome, which is the actual; the amplitude between outcomes (the smooth wave by which quantum theory sets the odds), which is the potential, stays continuous. At the largest limit physics brings continuity, the transcendent tool, and finds a setting that is unbroken and directed.

More Than a Pattern Found by Looking

The grid is more than a pattern that was found because someone looked for it. A careful reader will still ask whether the grid is only that: six things here, six things there, lined up. I give three answers to the careful reader's question.

The first answer is the method. The method of the metaphysics makes two large claims. The first of the two large claims, *foundational triplication*, is that at the root of any domain there are three ideas, and the other ideas of the domain are built on those three ideas. The second of the two large claims, *type isomorphism*, is that the pattern of relations at the base of each domain matches the pattern of relations at the base of every other domain. The patterns match because every domain is made by comparison, and comparison has one structure. Measurement is a comparison. Interaction is a comparison. So the structure of measurement must match the structure of interaction, both must match the structure of comparison, and they must match modality by modality, because the modalities are what the structure is made of.

The second answer is that the match is two matches, each checked on its own. Comparison; The Necessary Basis checks comparison against measurement. I have checked interaction against measurement, using An Immanent Metaphysics's own list of the six intrinsics of interaction, a list written with its modalities and its rows already assigned, before any link to physics was drawn. The two checks fit together. The interaction list could have matched the six measures in some other order, and it matches them in the one order that the six intrinsics of comparison require. That order was at stake in the check: had the six intrinsics of interaction matched the six measures in any other order, the framework would be in trouble. They match in the required order, and that agreement is evidence of the kind the metaphysics is able to offer: agreement among structures stated independently of one another, and not experiment.

The third answer is a further sign. Every column of the grid is a content in a context, and every row is governed by one of the two kinds of law. Six quantities chosen at random would not do that. Pick six quantities from a physics textbook, and they will not sort themselves into three modal pairs, with the content of each pair governed by symmetry and the context by continuity. That these six do so is a consequence of the structure of comparison, as I read it.

What the Grid Says About the Two Theories

Chapter 17 showed, in full, that each of the two great theories reaches with the tool of one modality into the domain of the other. The grid shows the same thing briefly, as a matter of columns. What follows is a reading of known physics, and it starts with the mechanics of Newton, because the grid is already at work there.

Newton's second law says that force equals mass times acceleration. It relates three quantities: force, the immanent content; mass, the omniscient content, as inertia; and acceleration, the rate at which speed changes, speed being the rate at which position changes. The three have different units, different meanings in physics, and different ways of being measured: they are distinct. None is measured without the others. A force is measured by how much it speeds up a known mass, a mass by how much a known force speeds it up, and an acceleration by comparing positions at three or more moments. The law does not say that force equals mass, or that mass equals acceleration. It says that the three are related in a way that keeps them distinct and needs all three. The way Newton's second law relates force, mass and acceleration is the third axiom at work in mechanics.

The second law also shows the step of the second axiom, the step by which what happens (the immanent) lays down what is laid out (the omniscient). Forces act through time, in the immanent column; acceleration is the force on a body divided by its mass; speed is acceleration piled up,

and position is speed piled up; so forces acting through a stretch of time fix the position at its end, in the omniscient column. The piling up runs from the immanent to the omniscient, as the step of the second axiom does. The return from position to force is where Newton's mechanics leaves that order. In Newton's law of gravity, the arrangement of the masses fixes the forces on every body at once, at every distance, with no stretch of time between. That step runs from the omniscient back to the immanent with no interval, and Newton's world, which allows signals of any speed (Chapter 12), has no edge that would supply one. So the loop of Newton's mechanics, forces fixing positions through time and positions fixing forces at once, is a loop of two terms, the immanent and the omniscient, and the transcendent column does not enter it. Its instant return is the first of the limits that physics found after Newton.

Newton's mechanics has two later rewritings, by Lagrange and by Hamilton, and the three forms give the same predictions wherever all three apply. Forms that give the same predictions can be swapped for one another, and the third axiom requires the members of a triple not to be swappable, so the three forms of mechanics are not a triple of the modalities. The Appendix gives the technical form of the two rewritings, read in the six measures.

General relativity, the theory of gravity as the shape of space and time, works the omniscient column: mass in space, mass and energy in relation to the geometry they occupy. It reaches that column through the immanent column, force in time: the workings by which the bending of the arena and the motion of bodies answer to each other. It treats that middle column continuously. Quantum theory, the theory of matter and light whose findings come in countable steps, works the transcendent column: probability in possibility, the weighting of outcomes across a range of alternatives. It reaches that column through the same immanent column, the workings of fields and their interactions, and it treats that middle column in discrete steps.

So the two theories share the middle column, and they share it in two different ways: one applies continuity to it, the other symmetry. The overlap is the place where the two survivors of the theorem are both asked to be perfect at once, and the theorem says that the two survivors cannot both be perfect, so the overlap is never a meeting point. A unification would have to be whatever joins the omniscient column to the transcendent column without passing through the immanent column. The first axiom says that the omniscient and the transcendent are conjugate, with the immanent more basic than either: their connection runs through the immanent, never directly. That connection is the interaction itself, the one event of which all six boxes are parts, and it is not one more piece of mathematics.

Picture two neighbors, each taking over the strip of land between them from their own side. General relativity takes over the middle column from the left. It treats time as one more geometric direction, stretching the omniscient idea of space into the immanent; it treats force as curvature, a property of geometry and structure, pulling the immanent content into the omniscient frame of smooth geometry. Quantum theory takes over the middle column from the right. It treats force as the exchange of quantized go-betweens, pulling the immanent content into the transcendent frame of weighted possibility. The forces so treated are the forces of matter and light; gravity, as Chapter 17 showed, is the shape of the arena, and gravitation is not quantized. Quantum theory treats time as the dial along which the state changes, smoothly and by fixed rule, broken by the events of measurement. Quantum theory does not treat time as a geometric direction, and the dial is an idea of time different in kind from the geometric idea of time.

Each theory folds the middle column into its own modality, and neither theory handles the middle column in the middle column's own terms, as process. The axioms lead one to expect what the two theories do with the middle column. By the first axiom the immanent is more basic than the omniscient or the transcendent, so the immanent is derived from neither of the other two, and every formal theory that holds the immanent

takes it as given, without deriving it. General relativity takes its findings as events in its smooth space and time, and quantum theory joins its findings to its mathematics by a rule set down beside its smooth, fixed-rule evolution. The middle column belongs to neither side, and the disagreement between general relativity and quantum theory is how that shows up in physics. The grid does more than sort.

The Two Constants

With the grid in hand, the two constants can be seen for what they are. Planck's constant and the speed of light are basic measures of the physical domain because they set its smallest and its largest limits. They set the limits of what a knower who is part of the domain, standing at a place inside it, can know. Each divides what is known, and what could come to be known, from what is at last and without remainder unknowable. They divide it on different sides, and the difference is the thing to see.

The speed of light sets what is unknowable on the side of context. Space and time are context measures, the bottom row of the grid: they are the where and the when in which any content is found. The speed of light is a boundary between the times and places that can be known from here and those that, from here, cannot be known at all. It is the macroscopic limit, the limit of the very large, and in the alignment this book holds it is the omniscient boundary of the domain: the outer edge of the pattern of connection that can be shared and repeated.

Planck's constant sets what is unknowable on the side of content. Mass, force and probability, the still and the moving sides of pattern, are content measures, the top row of the grid: they are the what that is found in a setting. Planck's constant is a boundary between the happenings that can be known and those that cannot be known at all. It is the microscopic limit, the limit of the very small, and in the same alignment it is the transcendent boundary of the domain: the finest grain of the possible, below which no further difference is made.

When I mean the two as the edges of any domain at all, I call them, as Chapter 18 began to do, the C constant and the H constant, after the letters physicists write for the speed of light and Planck's constant, and when I mean the physical cases, I use the physical names. The difference between the general names and the physical names matters because my claim in this chapter is about every domain, and not only about the physical domain. A *domain constant* is either of the two constants that bound a domain, the C constant at its reach and the H constant at its grain.

The Incommensuration Theorem says what is found at each edge. At the large edge we ask the context measures to stretch on without a gap; that is asking for continuity absolutely, and what is found is asymmetry. The context measures are continuous there, and the light cone of Chapter 12 (at each event, the set of every event that can be joined to it by cause) opens one way. At the small edge we ask the content measures to be perfectly repeatable; that is asking for symmetry absolutely, and what is found is discontinuity: the content measures are quantized there, and the outcome cannot be refined. So the two constants are two limits of two kinds, and never two cases of one kind of limit. The C constant is where the continuous context meets its one-way character. The H constant is where the symmetric content meets its breaks. They are the physical faces of the two survivors of the theorem, and there are two of them because there are two survivors.

These constants belong to every domain. Every domain has the two in some form: a farthest reach, beyond which its kind of comparison does not extend, and a finest grain, below which its kind of comparison tells nothing apart. Each domain has its own values for them, particular to itself. A mathematical domain has a largest structure it can name and a smallest difference it can tell apart. A domain of experience has a horizon of attention and a threshold below which it notices no difference. The physical values are the values of the physical domain, and nothing in this framework says what they are. The framework says only that there must

be two, that they must be of the two kinds described, and that a domain lacking either would be a domain in which comparison had no edge, which the theorem rules out.

The Units and the Order of Nature

The way physics writes its units shows the same two constants again, this time in the order in which the units are defined. The units first used in physics were built outward from the omniscient, from structure. Length, mass and time were the first base units of mechanics, set up before the physics of heat and quantum theory brought the transcendent column into physics.

Since 2019 the International System of Units, the system of units used in science worldwide, has fixed every one of its units by the exact value of a *defining constant*, a constant whose value is fixed exactly, by definition. There are seven: the speed of light; Planck's constant; the elementary charge, the charge of one electron; Boltzmann's constant, which converts between energy and temperature; Avogadro's constant, the count of particles in a standard amount of substance; the luminous efficacy, which sets the unit of brightness as the eye sees it; and the frequency of one particular change of state in the cesium atom, which fixes the second.

The definitions make the units depend on one another like the branches of a family tree, with time at the root. Time stands alone. Length depends on time, through the speed of light. Mass depends on time and length, through Planck's constant. Charge depends on time, through the elementary charge, as a side branch. Temperature depends on time, length and mass, through Boltzmann's constant. The unit of brightness depends on time, length and mass, and the unit of amount of substance depends on no other unit; those two are conventions laid on top. The kilogram keeps its name among the base units, but its definition runs through Planck's constant, the meter and the second, so in this system mass is fixed by them and is not a base of its own.

The system derives the unit of force from the unit of mass, through Newton's second law, force equals mass times acceleration. In its table of units, then, force is made of pattern, space and time. In the grid, force is the immanent content: the directed flow of change, the push or pull that changes how something moves. Pattern, which gathers mass, inertia and structure, is the omniscient content: what is there, laid out. The first axiom states: "The immanent is more fundamental than the omniscient and/or the transcendent"; in plain words, process is more basic than structure or possibility. Read in the grid, the order in which the system builds force from mass runs against that axiom: it builds the immanent content out of the omniscient content.

Read through the chain of the defining constants instead, the system already holds mass as derived, since the kilogram is fixed through Planck's constant, the meter and the second. In that reading, the defining constant through which force enters the system of units is the elementary charge. Charge enters physics as what electric force acts on, and the ampere, the unit of current that the elementary charge fixes, is a rate of flow of charge: a measure of the force acting in time. The step is short: charge is what the electric force acts on, so a flow of charge, counted each second, measures the force at work through time. Mass then follows from the constants already brought in. So in the reordered reading, force is one of the boxes of the grid that a defining constant brings into the system of units, and mass is derived. Nothing changes in practice: Newton's second law holds in either reading, and every equation keeps its form. What changes is which quantity is written as basic and which as built.

The order in which the units were first built also hides the fact that two of the defining constants belong together. The revision of 2019 kept that old shape: it sent the new information that Planck's constant carries into the definition of the kilogram, and it treated Boltzmann's constant as a converter between units of energy and units of temperature. The transcendent measures themselves, probability and possibility, have no unit. What the order of discovery hides is that Planck's constant and Boltzmann's constant are the two constants of the transcendent column.

Planck's constant and Boltzmann's constant belong to the transcendent column because both work through counting. Planck's constant counts the states a system can be in (Chapter 18), and a count of states is a measure of possibility: it gives possibility its grain. Boltzmann's constant turns that count, or the spread of probability over those states, into energy per degree of temperature, which is what an entropy is: it gives the spread of probability its measure. So one constant serves possibility and the other probability, the two boxes of the transcendent column, as my book Measurement; The Steps from Comparison to Physics sets out. (The transcendent column is one of the grid's three columns. In Chapter 18 I placed Planck's constant also on the content row, as the bound of the small edge. A box of the grid lies in a row and a column at once.) The hiding of the transcendent column by the order of discovery is one case of a wider pattern. Description starts from the omniscient (the structural, the measurable, the knowable), and a description built from it tends to make the immanent and the transcendent look secondary to it.

Picture one room described with two different sets of directions. One visitor says "north, east and up"; another says "toward the window, toward the door, and up". The room is the same; only the directions chosen as the reference lines differ. Mathematicians call such a switch a *change of basis*, and the two ways of choosing units are related in just that way. Counting length, mass, time, current and temperature as the base kinds, as the international system does (leaving out the amount of substance and the brightness unit, the two conventions laid on top), gives five base kinds. The reading by the intrinsics takes instead the cesium frequency, the speed of light, Planck's constant (an action), the elementary charge (a charge) and Boltzmann's constant (an entropy, a measure of how many ways a state can be arranged), so that mass, current and temperature become derived. Both count five. Two different things are counted here, the base units chosen and the kinds of amount underneath them. Counted as independent kinds of amount, the same set has only the three of mechanics, once the constant of the electric force and Boltzmann's constant are each read as conversion factors (as physicists do in the units called Gaus-

sian, and whenever temperature is measured as energy). Then the units of current or charge and of temperature add no kind of amount beyond mass, length and time. All the pure-number ratios come out the same in both choices, so both carry the same information; only the reference lines differ. The claim of the Immanent Metaphysics is that the reading by the intrinsics is the one that lines up with the six intrinsics of comparison, and so it shows the structure of the modalities where the usual choice hides part of it.

The international system is a way of writing the relations among physical quantities that grew up over time, and its shape follows the order of discovery: mechanics first (length, mass, time), electricity and magnetism second (charge), the physics of heat third (temperature), quantum theory last (folded into mass through Planck's constant). The Immanent Metaphysics says: the order of discovery is not the order of nature.

Zero, Unit and Extent

Between the two constants lies a third scale, and the three together have a structure that every measurement with numbers already uses. Every such measurement needs three things: a zero, a unit and an extent. The *zero* is the point below which the measure has no definite meaning. The *unit* is the yardstick against which other amounts are compared. The *extent* is the reach beyond which the measure loses its structure, where the measure passes beyond what the domain can tell apart. The dark room of the opening is the principle at work, and light shows it most clearly: every basic measure has a zero and no maximum, and a minus sign gives a direction and never an amount. My book Unlocking Reality develops that principle; here it says something about the scales of a domain.

The zero matches the microscopic edge. It has the nature of the transcendent: below it, what could be outruns what is, and nothing definite can be pinned down. In physics this is the limit Planck's constant marks.

In perception it is the threshold below which no difference can be told: the faintest sound you can hear, the smallest difference you can see, the point at which two touches on the skin feel like one.

The unit matches the middle scale, which physicists call mesoscopic. It has the nature of the immanent: it is the ground of relating, the comparison itself, the act of measuring. In physics this is the human scale, the scale of the laboratory and the instrument, where a meter and a second are what they are. In perception it is the scale of ordinary life, where things have their natural sizes.

The extent matches the macroscopic edge. It has the nature of the omniscient: it marks the structural horizon beyond which the domain's description no longer applies. In physics this is the cosmological horizon, the surface from which nothing has yet reached us. It is the edge, as Chapter 12 described it, between the part of the universe whose light has reached us and the part whose light has not yet arrived. It is a different edge from the cosmological event horizon of Chapter 18, which bounds the events whose light never will arrive. In perception it is the limit of contact by cause: the farthest thing that can affect you.

So each domain has three scales, and the three are the three modalities: the microscopic is transcendent, the middle immanent, the macroscopic omniscient. The matching of the three scales to the three modalities is the alignment that I use throughout this book, now seen from the side of measurement. Quantum theory is applied in the transcendent domain, general relativity in the omniscient domain, and the immanent scale between the two domains is where the measuring is actually done.

The three scales are themselves a triple, in the sense of the third axiom: distinct, inseparable, and not interchangeable. They are distinct: the smallest scale gives separate, countable, particle-like things; the largest gives continuous, spread-out structure on the scale of the cosmos; the middle gives the familiar workings of forces and motions. The three scales are three things, and not three names for one thing. They are inseparable: no scale exists without the others. "Very small" means something only

against a "very large", and the middle scale is where that contrast is drawn, the scale from which both are defined as limits. Without all three, the idea of scale itself is empty. They are not interchangeable: quantum theory cannot do the work of general relativity, the mechanics of Newton cannot replace quantum theory, and each scale has a part to play that the others cannot.

What comparison finds takes a different form at each scale, because each limit has a different character. At the middle scale comparison is performed. No constant of physics bounds it; its limit is the act of comparing itself, which no comparison takes as what it compares. What shows at this scale is the frame in which mechanics is written: mass in space, force in time, and probability in possibility, the three columns of the grid, with the three axioms relating them.

Read from small to large, what is found grows coarser in grain. At the finest grain, the smallest limit, there are the six intrinsics of comparison, found again as the kinds of particle. My book The Story of the Particles calls that finding *re-projection*, the six parts of every comparison showing up again as the kinds of thing that physics finds at the bottom. At the middle grain there are the three columns of the six measures, and the axioms that relate them. At the coarsest grain, at the C constant, there is the threefold idea into which the idea "universe" resolves: existence, interaction and creation. Two things here are my identifications, made in this book: the use of the word projection, what comparison finds at a scale, for the middle and large scales as well; and the ordering from finest to coarsest. At the middle and large scales the word projection says that what is found is shaped by the character of the scale at which it is found, and it leaves the kinds found there as they are, without calling them the intrinsics. The ordering rests on what the grid says of the two limits: content must be quantized at the smallest scale, so the smallest limit is where content is divided into separate kinds; context must be continuous at the largest scale, so the largest limit is where only the broadest feature of structure appears.

An Immanent Metaphysics also names the three ideas of measurement as position, scale and direction: position, where a thing stands in an ordering, omniscient; scale, the change from one level to another, transcendent; direction, the flow from toward to away, immanent. These matchings sit one place over from those of zero, unit and extent, because what is in view has changed, by the rule Chapter 18 set out. The shift has a name in my books, an *Axiom II rotation*, said aloud as "axiom two rotation". When the thing being considered changes, here from the structure being measured to the act of measuring, each matching to a modality moves one step along the cycle of the second axiom. The cycle of the second axiom is the cycle in which a class of one modality comes before an instance of the next modality, omniscient to transcendent to immanent and back. The rotation is an idea of its own, used beside the second axiom, which it leaves as it is: a change of the point of view, with the teaching unchanged. It comes back below, in the three kinds of equation.

The Width Between the Edges

My central claim is that the degree of objectivity that knowledge can have inside a domain, for a comparer inside it, rises in step with the span between the domain's two edges, counted in powers of ten. The count is how many times one must multiply the smaller by ten to reach the larger. The wider the span, the richer the domain, the more structure it allows, and the more can pass between self and world inside it.

Chapter 4 showed that objectivity is a matter of degree, the agreement of many findings across comparers, and my book Accessing the Real argues it at length. Nothing is purely objective. Objectivity belongs to clusters of related perceptions, never to any single one, and its degree is three things multiplied together: how strong the contact between self and world is; how consistent that contact is across settings, which is a matter of the large; and how much real newness it carries, which is a matter of the small: newness that no earlier state of the comparer could have pre-

dicted. The second factor is the part that is symmetry: sameness of content across difference of context. The third factor is the part that most accounts leave out: randomness, without which a perception brings no news from outside the self.

Now set those factors against the two constants. Consistency is found across settings, and the large edge caps how many settings there are: the farther the domain reaches, the more settings across which a pattern can be found the same. Newness enters at the finest grain, and the small edge sets that grain: the finer the grain, the more room for what is new to enter from below anything already specified. So the span between the two edges is the range of scales at which perceptions can be both consistent and informative. A domain with a wide span allows many levels of organization, many degrees of objectivity, many scales of interaction, each level able to serve as content for the level above it and as context for the level below it. A domain with a narrow span allows little structure and little objectivity: there is almost no room between the point where consistency fails and the point where information disappears.

By a domain's *degree of structure* I mean the count of how many times content can sit inside context between the grain and the reach. It is also why I describe the span as the *bandwidth* of a line of communication between self and world, bandwidth being how much a line can carry: the line carries as much as the number of levels that can be told apart between its floor and its ceiling.

For the physical domain the span is very large. The Planck length (Chapter 18), written out in meters, is a decimal point, then thirty-four zeros, then the digits one and six. The radius of the observable universe, in meters, is about forty-four followed by twenty-five zeros. The two are about sixty-one powers of ten apart. The figures for the Planck length, for the radius of the observable universe and for the sixty-one powers of ten between the two lengths are figures of physics, computed from measurements, and I quote the figures here from physics. The span is a ratio of two lengths, and its lower end is the scale at which the description of context and the description of content meet: a bound on findings, as Chapter

18 showed, and not a least length. What the framework says of the span is only this: a span of that size is why physics can reach objectivity of such high confidence.

Picture the sciences as the rungs of a ladder that runs up that span: the nuclear, the atomic, the molecular, the biological, the planetary, the stellar and the galactic. There are that many levels at which a pattern can be found consistent and at which newness can still enter, and each rung is a layer of structure made possible by the room between the two constants. A domain only a few powers of ten wide would have no such ladder, and no such confidence. The rise of objectivity with the span is a correspondence of structure: the degree rises with the span, for the reason just given.

Mass, Energy and Light

The two constants have a companion triple among the basic ideas of physics, and it follows the same pattern of the modalities.

The idea of a state, a structure, an arrangement, and in the end the idea of matter and of mass itself, has the nature of the omniscient. Mass is pattern: the fixed, structural, repeatable side of a domain, what stays the same and can be counted on to stay the same.

In the reordered reading of the units, Planck's constant does not bring mass in as a base quantity. Mass is derived, fixed by the constants already brought in together with a length: the grain of possibility (Planck's constant), the bridge between space and time (the speed of light), and a length in space. Physicists define a length for every body, its *Compton wavelength*, as Planck's constant divided by the product of the body's mass and the speed of light. Turn that around, and the body's mass is Planck's constant divided by the product of the speed of light and that length. Mass is the ratio of the grain of possibility to the product of the speed of light and a length in space. In that ratio, mass is derived from what comes

before mass in the chain: possibility, through Planck's constant; space, through the length; and the bridge between space and time, through the speed of light.

The same definition can be written as one product: the body's mass, times the speed of light, times its Compton wavelength, equals Planck's constant. On one side stand mass (pattern, the omniscient content), the speed of light (the bridge between space and time), and a length in space (space, the omniscient context). On the other side stands Planck's constant, the grain of possibility in the transcendent column. The relation shows how a mass is fixed once Planck's constant, the speed of light and a length are given: the omniscient specification of a particle, its mass and a length in space, joined through the speed of light, equals the grain of possibility. The relation is a definition. The relation shows that, in this reading, mass is fixed by the grain of possibility, the bridge between space and time, and a length. So mass is not a base quantity.

The idea of a relation between states, a comparison of possibilities by their probability, and in the end the idea of energy itself, has the nature of the transcendent. Energy is potential. To speak of an amount of energy is to speak of a difference between possible states, of what could happen rather than what is.

The idea of measurement, of interaction, of a change of state, of force, power and strength, and in the end the idea of observation itself, which is light, has the nature of the immanent. Picture the simplest comparison in physics: two atoms, and the light that passes between them. In that picture light is interaction: the act in which pattern and potential are brought into relation. In the physical domain it is what anything is seen by at all, and I use the word light in that widest sense, for whatever carries any measurement, and not for the visible colors alone.

The three axioms then apply, one by one. By the first axiom, process is more basic than structure or possibility, and structure and possibility are conjugate, each perfected only at the cost of the other, neither prior. So interaction, light, is more basic than both matter and energy, and matter and

energy are the omniscient and the transcendent faces of light, related through it. That matter and energy are the omniscient and the transcendent faces of light is my reading of the equivalence of mass and energy. The equivalence of mass and energy is the finding of physics that a body's mass and its energy always stand in one fixed proportion. They are one currency, like one money in two denominations, the two faces of a single interaction, held in proportion and never traded one against the other. The constant relating them is the constant that bounds interaction, the C constant. My account of mass, energy and the C constant is a reading of a known result, the equivalence of mass and energy. In my reading I say why the relation of mass to energy has the form it has, and the relation itself comes from physics.

By the third axiom, the three are distinct, inseparable and not interchangeable. There can be no matter wholly without energy and interaction. There can be no interaction wholly without structure and potential. There can be no potential wholly without substance and access. Each of the three needs the other two, and none reduces to another.

By the second axiom, a class of one modality comes before an instance of the next, around the cycle from omniscient to transcendent to immanent and back. Applied here: before there is any small degree of energy, of potential, there must be large degrees of matter, since a class of the omniscient comes before an instance of the transcendent. Before there is any small degree of light, of interaction, there must be large degrees of energy, since a class of the transcendent comes before an instance of the immanent. Before there is any small degree of matter, of the actual, there must be large degrees of light, since a class of the immanent comes before an instance of the omniscient. The three come before one another in a circle, and none of them is first.

One more word on mass. In this framework mass and pattern are taken as the same in form, by way of the principle Ernst Mach set out in The Science of Mechanics (in English, 1893). By that principle, the inertia of a thing, its very thing-ness, is set by how it changes in relation to all other things, and the whole structure of those relations, taken in terms of

their changes, is the being of the thing as pattern. Here the interactions among interactions are those relations of change, taken together as they act on one another. From this it follows, as a correspondence of structure and not as a calculation of amounts, that mass varies inversely with the potential of the interactions among interactions. The more confined and isolated an interaction is, the more liveliness and potential it develops; and the smaller and more confined it is, the more energy it takes to perceive it. That last clause is the content of the uncertainty relation of Werner Heisenberg (1927) as Chapter 18 read it, and it is where the triple of mass, energy and light meets the H constant. The finest grain of the domain is the scale at which the energy needed to perceive a thing becomes the energy that destroys what is perceived.

Interaction, Existence and Creation

The principle that every basic measure has a zero and no maximum reaches past measurement, into what there is, and three of its consequences make the two constants constants of the kind they are. Call a degree *zero-positive* when it has a floor at zero and nothing below it, and no ceiling above.

Interaction is zero-positive. One cannot fail to interact. There is no anti-interaction, no negative degree of contact between self and world. The degree of interaction can be small, can come close to its floor, but there is nothing beneath the floor. Because the degree of interaction has nothing beneath its floor, the smallest limit is a limit. Below the H constant there is no interaction that can be told apart, rather than less interaction. The difference between less interaction and none is the difference between a scale inside the domain and the edge of the domain.

Existence is zero-positive. There are degrees of existence, which in this framework means the lasting of a pattern across settings, but there is no negative existence, no thing that lasts less than not at all. Because there is no negative existence, the largest limit is a limit. Beyond the C constant

there is a region about which no one can even ask, from here, whether anything in that region exists, because no pattern from that region lasts into here. The region beyond the C constant is not a region of lesser existence.

Creation is zero-positive. The power of creation, the coming in of what is new, can be greater or less, but it cannot be negative. There is no anti-creation, no removal of possibility that goes below the possible. That creation is zero-positive is what makes newness, the third factor of objectivity, a factor at all. The newness that enters at the finest grain has a floor of zero and no ceiling, and so the domain can always be surprised.

The three consequences are once again a triple: interaction immanent, existence omniscient, creation transcendent. Because the three are inseparable, wherever interaction and existence are, creation is too. The order the axioms set among the modalities is an order of what presupposes what, and not of clock time, so nothing in the axioms places creation at an earlier time than interaction, and creation is not confined to the past. Creation is a side of every event, and creation is not itself an event. That conclusion parts from every account that puts creation at a first moment, and it agrees with the traditions that hold creation to be continuous.

The consequence for physics is direct. Physics as it is ordinarily done treats creation, existence and interaction as separate: creation comes first and then stops, existence follows as a standing condition, and interactions follow with nothing further new. Since creation is ongoing and not a thing of the past, the universe is not a system whose whole history is fixed by how it began. Something new enters in every interaction that makes a finding not contained in what came before, and real newness belongs to the structure of every moment.

The degree of interaction is always above zero. It can never be zero in an absolute sense, because zero is the edge of the domain and not a state inside it. So there is always at least this: the plain fact of contact between

subject and world. That, in the framework's terms, is why there is something rather than nothing; Chapter 20 shows why physics itself cannot ask that question.

A test confirms that the triple leaves nothing out. Is there anything in the idea "universe" that existence, interaction and creation do not hold? Try every candidate for a fourth part. Space is a side of existence: the structural spread of what is. Time is a side of interaction: the ordering of events in process. Possibility is a side of creation: the range of what could come to be. Law is a side of existence: regularity laid out as pattern, since the omniscient is the structural and the repeatable. Matter is a side of existence: the structural content of what is. Force is a side of interaction: forces unfolding across time. Energy is neither a fourth part nor a side of one of the three alone: it is a relation that can be written in the six measures, across all three columns. The energy of an arrangement is a relation of pattern in space across possibility, and the energy of motion is written through momentum, force and time. Energy has its place in another triple, the triple of mass, energy and light, where energy is the transcendent face of light. A thing can have a place in each of two triples, because the two triples are about different things, as one woman can be a mother in one family relation and a daughter in another. Every candidate comes down to one of the three, or to a relation among them. The test runs through the candidates named here. My book Incommensuration; The Theorem and the Three Axioms makes and argues the full claim: if one knew everything about existence, interaction and creation, nothing in the universe would be left out.

The usual picture of the universe is a box, a container in which things happen. That picture is the omniscient view of the triple, laid out over space and time: it treats existence as the main part, and folds interaction and creation into the box as events and processes inside it. The reading of this book reverses it. The universe is the triple of existence, interaction and creation, and not a box with things happening inside it. The arena of space and time is the omniscient face of the triple: the face that can be mapped, measured and described as structure. It is only one face. The im-

manent face is interaction, the ongoing ground of relating that everything stands on. Every measurement is an interaction. Every comparison is an interaction. That the universe is made of relations is a feature of what it is, and not a feature of what it contains. The transcendent face is creation: not an event in the past but the ever-present potential for what is new. The transcendent will not be held in categories of space or structure. It cannot be located, mapped or fully described. It can be recognized only from the inside, by taking part.

Three Kinds of Equation

Start with three kinds of counting. Two friends pool their coins: three coins and four coins make seven, and each friend's share can still be picked out of the pile. A café offers three soups and four breads, and every soup can go with every bread: that makes twelve meals, three times four, the full table of pairings, and no meal is one side's share alone. Now hand each of three guests one of four colored hats, and count every possible rule of handing them out: four times four times four, sixty-four rules, every way the handing out could go. These three kinds of counting are adding, multiplying and raising to a power, and each answers to one modality.

Take two selves in one shared world, and ask how the total information each holds can be related. The framework names three kinds of equation for it, one for each modality, and my book Category Theory and the Immanent Metaphysics proves which arithmetic operation goes with which modality, the matching the next sentences give. Adding joins two things into one; it is interaction, and in a purely immanent setting the equation of information is a pure sum: what the one holds and what the other holds add up to what is between them, and each one's share can be told apart. Multiplying forms the complete table of all pairs, a pattern given all at once, and in a purely omniscient setting the equation is a pure product: the shares cannot be told apart, and the whole is the full table of

their pairings. Raising to a power forms the range of every possible rule, every way things could go, and in a purely transcendent setting the equation takes the form of a power, with its reverse, the logarithm, which asks what power gives a number. Many additions taken together make a multiplication (the twelve meals are four breads added three times over), and many multiplications taken together make a power (the sixty-four rules are three fours multiplied together), in the order of the second axiom.

At the edges of a domain, the kinds appear in this way. At the large edge, relations are written as sums, in the form "this amount plus that amount equals zero", with a basis in symmetry. A sum equal to zero is the form of a conservation law, where what comes in and what goes out add up to nothing. What goes out is counted with a minus sign, which gives its direction, as the opening of this chapter showed. And Chapter 2 showed that conservation is what symmetry requires of a domain. At the small edge, relations are written as products, in the form "this amount times that amount equals one", with a basis in continuity. A product equal to one is the form of the uncertainty relation, where two spreads multiply to a fixed amount, and neither spread can be made smaller without the other spread growing. The "one" of the form is that fixed amount, taken as the unit: in the uncertainty relation of Chapter 18 the fixed amount is half of the reduced Planck constant, and with each product counted in that amount, the floor of the product is one. In between, relations take the form of a power law or a logarithm, and that is the form that builds shapes repeating their structure at every scale: the fractal forms, like coastlines and branching trees, that the middle-sized world is full of.

A careful reader may ask here whether the pairing of edge and tool has turned over. In Chapter 18, and in the section on the two rows, I said that physics asks for continuity at the large edge. In both places I also said that physics asks for symmetry at the small edge. Now the large edge has a basis in symmetry, and the small edge has a basis in continuity. Both teachings stand, because they name two different things at each edge. What is asked for at an edge is the tool that physics brings to it, and each

edge is where the tool proper to the other modality is used: the large edge is omniscient, and physics brings it continuity, the transcendent tool; the small edge is transcendent, and physics brings it symmetry, the omniscient tool. The basis of the relations written at an edge is the tool proper to the edge's own modality: symmetry, the omniscient tool, at the omniscient large edge; continuity, the transcendent tool, at the transcendent small edge. So at each edge physics brings the tool of the other modality, and the relations it writes there rest on the tool of the edge itself. The cycle of the second axiom shows the same fit from the side of the equations. At the large edge the equation is a sum, which is immanent, and its basis, symmetry, is omniscient: one step along the cycle, since a class of the immanent comes before an instance of the omniscient. At the small edge the equation is a product, which is omniscient, and its basis, continuity, is transcendent: again one step along the cycle, since a class of the omniscient comes before an instance of the transcendent.

One rule governs all three kinds, and it is the Axiom II rotation again. The sum, which is immanent, appears at the large edge, which is omniscient. The product, which is omniscient, appears at the small edge, which is transcendent. The power and the logarithm, which are transcendent, appear at the scale between, which is immanent. Each kind of equation stands one step along the cycle of the second axiom from the scale where it appears: a class of the immanent comes before an instance of the omniscient, a class of the omniscient before an instance of the transcendent, and a class of the transcendent before an instance of the immanent. That conservation laws take the form of a sum and the uncertainty relation the form of a product are facts of physics; that the arithmetic operations have the characters they have is proved; and the one step along the cycle is what joins the two. Joined that way, the sum at the reach and the product at the grain are a reading of known physics: the step gives the structural reason the two results have the forms they have.

What the Edges Mean for This Book

We now know that every domain has a floor and a reach, and that the two constants mark them for physics. Chapter 12 showed that a signal faster than light asks for a comparison beyond the outer edge of comparison, and so names nothing. In this chapter I show what that outer edge is: the C constant of the physical domain, the edge of the context row, where the unbroken setting meets the one-way character of the setting. I also show that the outer edge never comes alone. Every domain has an inner edge too, the H constant, where repeatable content meets its breaks, and every domain has the two because the theorem has two survivors. What lies past either edge is outside the domain, and being outside the domain is a different thing from being far away.

Most of what this chapter has said is a claim of the structure of comparison itself, the first kind of claim the Preface names; its readings of known physics are marked where they are made; and the Appendix lists every claim of the chapter with its kind. The grid is a structure of kinds, and it leaves every amount to physics. I add no new exclusion; I name the two edges of the domain in which the book's exclusions are made.

Summary

Every measurement uses six measures, mass, space, force, time, probability and possibility, and they sit in a grid of three columns and two rows. The columns are the three modalities, each a content in a context: mass in space (the omniscient), force in time (the immanent), probability in possibility (the transcendent). The rows are content (mass, force, probability), governed by symmetry, and context (space, time, possibility), governed by continuity. The six measures are the six intrinsics of comparison read through two descriptions, as interaction and as measurement, and

the six intrinsics of interaction match them one for one, column for column and row for row. So the steps from comparison, through interaction, to the measurements of physics are complete.

The Incommensuration Theorem says what the two rows must be at their limits. Content must be quantized at the smallest edge of any domain. Context must be unbroken, and directed, at its largest edge. The two rows are the two survivors of the theorem, and no single theory finds both perfected. Only the interaction holds them together, and it does so by being the event in which all six occur, not by describing them.

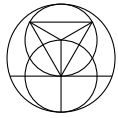
Every domain is therefore bounded by two constants. One is the large limit of context, of which the speed of light is the physical case: the omniscient edge beyond which no pattern arrives. The other is the small limit of content, of which Planck's constant is the physical case: the transcendent edge below which no difference is made. Between them lies the middle scale of the measuring self, and the three together are the zero, the unit and the extent of every measurement the domain holds. The span between the two constants is the domain's degree of structure, the bandwidth of the line between self and world, and the measure of the objectivity the domain allows. The physical domain has a very wide span, and that is why physics can be as sure as it is.

Mass, energy and light form a triple, omniscient, transcendent and immanent, in which interaction is more basic than substance or potential. Interaction, existence and creation each have a zero and no maximum, which is why the edges of a domain are edges, and why there is always something rather than nothing. The relations at the two edges take two forms, a sum at the reach and a product at the grain, each one step along the cycle of the second axiom from the scale where it appears.

The structure of physics, its equations, its constants and its basic ideas, is one case of the structure that every domain must have. The two constants belong to every domain. Their numbers express the units in which this domain is measured: the speed of light has been fixed by the definition of the meter since 1983, and Planck's constant by the definition of the

kilogram since 2019, and there is no fact about their values to explain. What this domain measures are the pure-number ratios they leave untouched.

If physics has edges, what lies beyond what any theory of physics can reach?



Chapter 20

What Physics Cannot Reach

Will Physics One Day Explain Everything?

Will physics one day explain everything? It is easy to think so. Drop an apple from a branch, and physics will tell you how long it takes to reach the ground, how fast it is moving when it lands, and why it fell at all. Ask when the next total eclipse of the Sun will cross your town, and astronomers will give you the year, the day and the minute, decades ahead. Each answer can be checked, and each kind has been checked, again and again. With answers like these behind it, physics seems to need only time before it answers every question there is.

Now listen to a child who has just learned the word "why". Why did the apple fall? Because heavy bodies attract each other. Why do heavy bodies attract each other? Because mass bends space and time. Why does mass bend space and time? Because the equations of gravity say so. Why do the equations of gravity hold? Here something changes. Each earlier answer came as a step up from the one before it. This one does not come in the same way, and I show why.

The last two chapters found the edges of physics. Chapter 18 showed that physics has three horizons: an edge at the large end, marked by the speed of light; an edge at the small end, marked by Planck's constant; and an edge on the side of the one who looks, the place past which a self cannot see itself. Chapter 19 showed that the two constants mark the reach

and the floor of the *domain of comparison*, the range within which comparisons can be made at all. In this chapter I ask what lies beyond the reach of every theory of physics, and why. The answer comes in two parts. The first part is about questions: there are six kinds of question, and one of them, asked in its deepest form, is a question no measurement can answer. The second part is about theories: every theory has three parts, and the three parts can never be made one, so no theory ever holds the whole of its world.

Six Questions

Every language that lets people ask things has a way to ask about the one who acted, the thing itself, the method, the reason, the place and the time. The words differ from language to language; the distinctions underneath the words are the same in every one of them. There is always a way to ask who acted, what is there, how it was done, why it happened, where it is, and when it occurred.

That sameness is no accident of grammar. The six questions are the six measures turned into requests. Chapter 1 named the *six measures*, the six things a physical measurement measures: mass, space, force, time, probability and possibility. The match between the six questions and the six measures runs one to one. A what question asks for mass, in the wide sense that takes in inertia, pattern, form and structure: what is there, what lasts, what a measurement finds. A where question asks for space: the position, the arrangement, the setting of the pattern. A how question asks for force: the mechanism, the coupling by which one thing acts on another, the way an outcome comes about. A when question asks for time: the moment, the order of events, the place in the flow. A who question and a why question stand together, in the column of probability and possibility. A who question asks for the agent: the one whose choice it was, that choice

weighed among the other things the agent could have done. A why question asks for possibility: the range of alternatives within which this outcome is one, and the principle from which it follows.

The six fall into three pairs, and each pair is a content in its context, as Chapter 19 set out the six measures in its grid: mass in space, force in time, probability in possibility. Each pair has the character of one of the three modalities that Chapter 4 named: the *omniscient*, the mode of structure, what holds the same for everyone; the *immanent*, the mode of process, what is happening here and now; and the *transcendent*, the mode of possibility, what goes beyond anything already found.

What and where are mass in space, and they have the nature of the omniscient: the structural, the repeatable, the shareable, what exists and where it is. An answer to a what question names a pattern; an answer to a where question locates it. Together they say what the world contains and how it is arranged. How and when are force in time, and they have the nature of the immanent: the process, the interaction, the here and now. An answer to a how question gives the mechanism by which the outcome came about; an answer to a when question places it in order. Together they say what is happening, and by what workings. Who and why are probability in possibility, and they have the nature of the transcendent: what weighs and values, what could be, what brings something new about. An answer to a who question names the one whose choice is at issue. An answer to a why question, in its full sense, gives a principle from which an actual structure follows. Both reach past pattern, toward the possibility within which a pattern is one alternative.

The six can also be grouped by what each question does, instead of by what it is about. Grouped that way, who and what pair as the subjective and the objective, the one who acts and the thing that is there; where and when pair as space and time; and why and how pair as principle and instruction. In that grouping, how and why run in opposite directions across one span: how runs down from a principle to a single case, and why runs up from a single case to a principle. The two groupings are two views of one structure, the first by content and the second by use, and

they agree on the point this chapter turns on: a what question is answered by a pattern, and the full sense of why reaches past pattern toward possibility.

One more matching does the work here. In *An Immanent Metaphysics* I aligned laws with the omniscient, axioms with the immanent, and principles with the transcendent; and I aligned knowing with the omniscient, and understanding, taken as what guides action, with the transcendent. So a what answer is shaped like a law: it says what is the case, in a form that can be checked, repeated and shared. A why answer, in its full sense, is shaped like a principle: it says what makes this the kind of thing that could be the case at all. Knowing what is there and understanding why it is there are answers of different modal character, and they differ in kind, not in amount.

What Physics Answers

Physics answers what questions. What is the structure of the atom? What colors of light does hydrogen give off? What is the state of this system at this moment? What questions are requests for pattern, and physics returns pattern: a structure, a value, a description of the state. The answers have the character of the omniscient. They can be written down, passed on, and confirmed by anyone who has the instruments.

Physics answers how questions. How does an electron move from one energy level of an atom to another? How does a gas settle into a steady condition? How questions are requests for the path by which an outcome comes about, and physics returns a mechanism, an equation of motion.

Physics answers when and where questions, sometimes with great precision. When will the eclipse occur? Where is the particle? At the smallest scale physics cannot answer these exactly, and there it says so, in the form of a spread of likelihoods over the possible places and times. That spread is itself an answer to a what question: it says what the limits of where and when are.

Physics seldom answers who questions, and when it does, the who is usually a which: which body exerted the force, which source sent out the light. The agent, in the full sense of the word, is not among the things physics measures. The absence of the agent from what physics measures is no accident. Who and why stand together in the column of the transcendent, and the two questions that physics leaves aside are the two that ask for possibility, and for the choosing within it.

And physics answers a certain kind of why question. Why does the apple fall? Because heavy bodies attract each other. Why is the sky blue? Because sunlight of shorter wavelength, the blue end of the rainbow, is scattered by the air more strongly than the red. Why does a pendulum keep time? Because, for small swings, the pull that brings the bob back grows in step with how far it has swung. Each of these is a real why question, each receives a real answer, and I do not want to say otherwise.

But look at the form of each answer. Each is a more general what. The apple falls because there is a law of attraction; the law is a pattern, a what, and the fall of the apple is one case of the law. In every case the why is answered by placing the case under a law, and a law is a what at a higher level of generality: a pattern that holds, and so forbids whatever would depart from it. I call this the *shallow why*. The name is not meant to belittle it: nearly all the explaining that science does is done at this level. The name marks its structure. The shallow why asks for the pattern of which this is a case, and physics answers it by supplying the pattern. The answer has the character of the omniscient. It is a law.

The Deep Why

The shallow why can be asked again, as the child asked it. Why do heavy bodies attract each other? Because mass and energy bend space and time. Why do mass and energy bend space and time? Because the equations of gravity say so. Why do those equations hold? At that question the answers stop, or, more exactly, the answers change kind.

When the questions reach the laws themselves, three replies are usual, and each of them grants that the answers change kind there. The first reply is that the laws simply hold, and there is no further why. That reply says nothing false, and it declines the question instead of answering it. The second reply is that the laws could not have been otherwise, that some need for consistency picks them out. The second reply is a claim about the range of possible laws, which is to say a claim about possibility, and the method that makes the second reply is no longer measurement. The third reply is that among all the worlds that might have been, only some have anyone in them to ask. The third reply too is a claim about the range of worlds that might have been, and again the method that makes the claim is not measurement. Each reply, in its own way, steps out of the physical domain, into the question of what the domain could have been.

The clearest form of the deep why is the question that Gottfried Wilhelm Leibniz put in 1714: why is there something rather than nothing? Leibniz's question is a why question in the full sense. It does not ask for the pattern of which this world is a case, because it asks about the existence of any pattern at all. It asks for the possibility within which a world is one alternative, and no world is another. Among the six measures, possibility is a measure of context, the widest of the three: the whole range of outcomes, against which any actual outcome is one. But the possibility this question asks for is more than the range of possibilities inside one world. It is the possibility that goes beyond every particular world: the possibility of there being a world at all. I call the why that asks for it the *deep why*: the why that asks for the possibility from which a whole world is chosen.

Physics is the method that belongs to a single domain. A measurement is an interaction, and an interaction happens between things that share a domain of comparison. Every what that physics returns is a what within the world it measures. Every how is a path through that world. Every when and where is a position in it. And every shallow why is a law of it. The deep why asks how this world stands to the possibility of worlds, and that relation is inside no world. It cannot be measured, because from with-

in there is nothing to compare it against. In *An Immanent Metaphysics* I put this as a difference of subject: metaphysics is the study of the relations between domains, and physics is the study of the interactions within one.

So the deep why is a question of a different kind from the questions physics answers. It is no question that physics has failed to answer and might yet answer with better instruments. The method of physics does not reach it, for the same reason that a ruler does not measure temperature. Lay a ruler against a cup of hot tea and you learn the height of the cup, exactly, and nothing about its warmth. Nothing is missing from the ruler, and in the same way nothing is missing from the method of physics.

Two other claims are sometimes made in place of this one, and where the deep why sits among the six kinds of question settles both: it is a why question, and it asks for possibility. The first of the two other claims is that the deep why is meaningless, because no observation could bear on the deep why. That claim mistakes the limit of what physics can reach for the limit of what has meaning. The question is well formed: it asks for possibility, and possibility is one of the six measures, which every act of measurement already takes for granted. The second claim is that physics has already answered it, by showing how a world can arise from empty space, or from a law such as gravity acting on nothing. Stephen Hawking and Leonard Mlodinow argued so in 2010, and Lawrence Krauss in 2012, and there have been others. Such accounts answer a how question and a what question: how a world of this kind grows from a starting state of that kind, and what the laws are under which it does. They take the starting state and the laws as given. The deep why asks about exactly what they take as given, and it is untouched.

I add one thing to the old question: I say which of the six kinds of question the old question is. It is a why question; it asks for possibility; and possibility of the kind it asks for is the one measure that no measurement within a world can return.

The What Questions of Metaphysics

Metaphysics asks what questions too, and most of its work is done there. What is the nature of existence? What is the nature of creation? How does the self stand to the world? What is the structure of comparison? The four questions about existence, creation, self and world, and comparison are what questions, which ask for pattern, for the structure that lasts, for what is there. The four questions are what questions about the structure of possibility itself, as distinct from what questions about the contents of any one domain. The Incommensuration Theorem is a what answer: it says what the structure of comparison is, and which pairings that structure forbids. The three axioms are what answers, about how the three modalities stand to one another. The six intrinsics are a what answer, about what any comparison involves. So it would be a mistake to conclude that metaphysics asks why and physics asks what, as if the two divided the six questions between them.

The difference between physics and metaphysics lies in what the question is asked of, and the question word does not mark it. Physics asks what, where, how and when of a world, from within it, by measurement. Metaphysics asks what, how and why of the relation between worlds and their possibility, by a method other than measurement, because measurement already takes for granted what metaphysics is asking about. The deep why is the question at which the difference cannot be missed, because it is the one question that cannot be rephrased as a what within the world.

That difference tells you how to read this book. Nearly every claim in it is a what claim. The chapters say what the structure of comparison is, what physics is a case of, and what, given that structure, physics will not find. They do not say why there is comparison. That there are six questions, each asking for one of the six measures, is a claim of the metaphysics itself. That physics answers what, where, how and when, and the shallow why, and not the deep why, is a reading of what physics does, with the match of questions to measures as its structural reason. In this chapter

I do not say that metaphysics answers what physics cannot. I locate the deep why, and I say what kind of answer the deep why would take. That answer would be shaped like a principle, as I said in the section on the six questions: an answer that says what makes a world the kind of thing that could be at all.

Three Parts of Every Theory

Now turn from questions to theories. Picture a walker in a strange town with a paper map. There is the town itself, with its streets, its noise and its weather. There is the map, neat and complete in its own way. And there is the walker's act of turning the map until its streets line up with the streets in view. To make a theory of something is to set the thing you want to understand in an ordered relation to a framework in which the understanding can be expressed. Every theory has exactly three parts, and I set them out in my book *Unlocking Reality*.

The first part is what the theory is about. I call it the *object domain*, and in physics it is the *world*, the physical world itself. It is always richer than any description of it. It holds both what the theory captures and what it misses, and the line between the two cannot be drawn from inside the theory. It has the nature of the transcendent: it comes before the theory, it goes beyond the theory, and it is described in terms of interactions, the events the theory is made to represent.

The second part is the *model*, which I also call the model domain: the symbols, equations, categories and formal relations in which the object is described. The model is always more structured than its object, more definite, easier to handle, and that is both its strength and its limit. It has the nature of the omniscient: a finished structure, which can be inspected from outside, checked for consistency, and passed on without loss, and it is described in terms of relations.

The third part is the set of links between the first two, which I call the *correspondence domain*: this piece of the model stands for that feature of the world; this relation in the model mirrors that relation in the world. The links are a domain in their own right, with their own structure and their own constraints. The links have the nature of the immanent: the links are not given in advance, but are made by acts of comparison and described in terms of comparisons. A *theory* is a comparison, the comparison of a world with a model, and the correspondence is where the comparison is made. By the correspondence I mean the practice of comparing in the world, the acts in which a model is held up against what it is about, and that practice is immanent. In the picture of the walker, the town is the world, the map is the model, and the turning of the map to fit the streets is the links.

Interaction, relation and comparison are the three names that the Root Tautology of Chapter 1 gives to one act, and here each part of a theory is described under one of the three names: the world in terms of interactions, the events that happen in it; the model in terms of relations, laid out in its structure; the links in terms of comparisons, the acts that make them. The act is one; which name fits depends on which part of the theory is in view, by the rule Chapter 18 set out.

Two consequences follow from the axioms. The first axiom says that the immanent is more basic than the omniscient and the transcendent, and that the omniscient and the transcendent are *conjugate*, a pair each of which can be perfected only at the cost of the other, with neither one prior. Applied to the three parts of a theory, the first consequence is that the links are more basic than either the world or the model. In the order of logic, the comparison comes before the split between what is compared and what it is compared with. A model stripped of its links to the world is a game played with symbols, and a world with no links to a model is experience that no theory has touched; neither the stripped model nor the unlinked world is a theory. The second consequence is that world and model are conjugate. World and model differ, and more than that, they are con-

jugate: each is defined by how it stands to the other across the links, and each goes beyond the other in a direction the other cannot follow. The next two sections take the two directions in turn.

The World Exceeds the Model

A scientific theory is one case of a theory with the three parts: a world, a model and the links between them. The object domain is the physical world. The model is the mathematics, the formalism. The links are the rules that connect the mathematics to measurement: this quantity in the equations corresponds to that reading on that dial; this number corresponds to that position on that instrument; this term in the equation is the thing the detector counts.

The first direction is the familiar one. The physical world, which is transcendent, always goes beyond the mathematics, which is omniscient. No model completely describes what it is about. I want to say exactly why this is a lasting condition, which no better model will remove. The transcendent is the mode of what goes beyond specification; the omniscient is the mode of what is specified. A complete model would be a specification of what goes beyond specification: no difficult achievement, but an incoherent one, a phrase with nothing behind it. Every improvement in a physical theory brings more of the world within the model, and every such improvement leaves a world that still goes beyond the model, because going beyond is what the world is, in its mode, and the excess is no amount that improvement shrinks toward nothing.

That the world always goes beyond the model is the structural reason for a fact that every physicist knows and that philosophers of science have struggled to state. Theories are replaced. A replacement is usually no disproof of the old theory in the range where its links to the world hold. Newton's mechanics is as good as it ever was in that range: for the falling apple, the thrown ball and the orbit of the Moon. A replacement is the discovery that the world went beyond the model in a place the model

could not have shown, because a model, being structure laid out, can show only what it contains. The excess was always found from the side of the links, by a measurement that failed to match, and never from the side of the model alone.

The Model Exceeds the World

The second direction is less often remarked. The model always contains structure that has no counterpart in the world. This follows from the same conjugacy, read the other way. The omniscient is the mode of structure, of pattern, of relation. A model is structure through and through, and a structure carries relations among its own parts whether or not each of those relations stands for anything. The world is no more the kind of thing that has a relation for every relation in the model than the model is the kind of thing that has a piece for every excess in the world.

To picture structure in a model that has no counterpart in the world, measure the hills around a valley twice: once as heights above sea level, and once as heights above the valley floor. Every number on the second list differs from its partner on the first. Yet every slope is the same on both lists, and a ball set down on any hillside rolls the same way, whichever list you read. The choice of where to put zero is real in the numbers and nothing in the hills.

Physics has a name for structure of this kind, structure that refers to nothing. It calls it *gauge*. The word gauge now meets the reader in a second use, and a third is coming, so the three are set apart here. Chapter 19 met the gauge values of An Immanent Metaphysics: the detail, not local and not determined, that the transformation of information puts in, matched with probability. Here gauge, in the physicist's sense, is structure in the model that refers to nothing in the world. And the gauge group, below, is the family of symmetries that the forces have, which does refer to something. I use the word in the second and third uses only. In the mathematics of electricity and magnetism, the fields that are measured are

computed from quantities called potentials. The potentials can be shifted, by an amount that may vary freely from place to place and from moment to moment, in a way that leaves every measured field exactly as it was. The shift is real in the model and nothing in the world. In quantum theory, the wave that describes a system carries at every point a size and a phase, a setting like the angle of a clock hand. Turn every hand by the same amount, everywhere at once, and no probability of any outcome changes. The overall phase is a relation among the parts of the model, and the overall phase corresponds to no feature of the world.

In general relativity the case is sharper still. The *coordinates* of the theory, the grid of numbers used to label places and times, carry no physical identity of their own. The same physical situation can be described in any grid, and features that show in one grid vanish in another. Think of the date line on a map of the world: cross it, and the calendar jumps by a whole day, yet the sea at the date line is like the sea anywhere else. The classic case in physics is a black hole. In the grid in which a black hole was first written down, the numbers break down at a certain distance from the center, the distance called the Schwarzschild radius, where the horizon lies. In another grid, that breakdown vanishes. Physics calls such features *coordinate artifacts*, and the practice of the discipline is to find them, and to count as physical only what stays the same when they are removed.

That practice is correct, and I do not propose to change it. What the framework adds is the reason the artifacts are bound to appear, so that their appearance in each new theory is expected, and no embarrassment. The third axiom says that the three modalities are distinct, inseparable and not interchangeable. Model and world belong to different modalities. They cannot be exchanged for each other, and so no matching between them can be exactly one for one. An exact one-for-one matching would let the model stand in for the world with nothing left over, and the world stand in for the model with nothing in excess, and that is exactly the exchange the axiom forbids. Gauge freedom is the model's leftover, the surplus structure that the axiom guarantees, just as the unspecified excess of the world is the world's surplus, guaranteed by the same axiom on the

world's side of the matching. In this sense the freedom of gauge (the free choice of potential and of overall phase) and the freedom of coordinates are structure that refers to nothing. I read the freedom of gauge and the freedom of coordinates, a known feature of physical theories, as structure that refers to nothing. The reason I give for that reading is the third axiom.

The *gauge group* itself, the family of symmetries that the forces of nature have, is no surplus. It refers to something: the exact charges that are found in nature. My book *Non-Commutative Geometry and the Immanent Metaphysics* derives the gauge group from the structure of comparison. In 1918 Emmy Noether stated, in physics, the difference between a freedom of gauge, which refers to nothing, and a symmetry that refers to something. In one paper she proved two theorems. Her first, the well-known one, pairs each symmetry of a certain kind with an amount that is conserved. Her second shows that a symmetry that can vary freely from place to place, as a freedom of gauge does, gives no conserved amount of its own, only an identity among the equations, a relation the equations satisfy automatically, whatever their solutions. And a theory in which every piece of the model matched a feature of the world would be a theory in which the model had stopped being a model.

The old saying that the map is not the territory is true, and it tells us little. The Incommensuration Theorem makes it exact. Recall the theorem: symmetry and continuity cannot both be applied, at once and absolutely, to one comparison, and neither can asymmetry and discontinuity; it is a structural fact about the nature of knowing. A model is symmetric with what it describes: it is the same content, the same structure, shown in a different setting, the setting of the mathematics instead of the world. The sameness of a model with what the model describes is what makes the model a picture of anything at all. But what is perfectly symmetric cannot also be perfectly continuous. So the model cannot be, at one and the same time, a perfect copy of the world and an unbroken continuation of the world, with no gap where the model ends and the world begins. There must be breaks between the description and what it describes. Gauge freedom is one such break, on the side of the model; the world's unmodeled

excess is another, on the side of the world. The links are the set of places where the two are joined, and the joining is finite, which is to say that it is in some places and never everywhere.

Physics as Instance, Metaphysics as Basis

The usual view takes physics as the foundational science and metaphysics as something extra built on top of it, to be done after the physics is finished, by those with a taste for it. The framework turns this around, and the reversal follows from the three parts of a theory.

Start with a definition. Metaphysics is often defined as what is not physics. A definition of that kind tells us nothing about what the thing is. The word "not" has no content of its own. To say "metaphysics is what physics is not" is like saying "this is whatever that is not": it tells us about that, and nothing about this. Because a negation says nothing of the thing defined except what it is not, a definition must say, positively, what the thing is.

Positively defined, physics is the study of the interactions within a given domain. Metaphysics is the study of the relations between domains. Put that way, which of the two is the basis of the other is already plain. Interactions within a domain take for granted that there is a domain; that it has edges; that what is inside it can be compared with what is inside it; and that each comparison has a subject and an object, a content and a context, a sameness and a difference. None of these is a result of physics. Each is a condition of there being physics at all.

Put it in terms of the three parts. Physics uses one particular object domain, the physical world; one particular model, the mathematics; and one particular kind of link, the method of measurement. Each of these is a case of a general idea (object, model, link) that the metaphysics describes and physics does not. Physics cannot say, from within physics, what a model is, what it is for a model to match a world, what a measurement is, such that it links the two, or why the mathematics should carry structure

that the world does not. Physics uses all of these. It does not ground them. Grounding them is metaphysics, and the grounding is the basis on which physics stands, and no extension of physics.

By the reversal of physics and metaphysics I mean that metaphysics describes the structure of finding, and physics finds what is there. I do not mean that metaphysics tells physics what physics will find. In *Unlocking Reality* I was explicit that the axioms alone cannot tell you what particles exist, what laws hold, or what will happen tomorrow. Taken alone, they say only that any world must involve interaction; that interaction must involve symmetry and continuity, though never both perfected; and that the subjective and the objective are both present in every comparison. The form of the physical world, its algebra (the rules by which its quantities combine), its symmetries and its charges, is reached from comparison in other volumes of this series, through identifications that those volumes name as such where they make them; for example, my book Non-Commutative Geometry and the Immanent Metaphysics derives the gauge group, the family of symmetries that the forces have, from the structure of comparison. The values of its numbers, and each single event, are matters for the sciences concerned, or for no description at all. The relation is that of a kind to its cases, and physics is one of the cases. To say that physics is an instance of metaphysics is to say that physics belongs to the kind of thing that metaphysics is about, and that physics cannot step outside that kind to examine that kind. Calling physics an instance of metaphysics does not mean that every value of physics can be deduced from metaphysics.

I call the layer beneath every physical theory *infra-physics*. It is the foundation of physics, and no replacement for it: the layer of structural necessity on which the facts of physics, which could have been otherwise, rest. Infra-physics asks what must be true of any possible universe, what structure is needed for there to be physical law at all. It names the conditions that any physical theory must meet, whatever the values of its constants and whatever the particular forms of its laws. The axioms, and the structural principles this book sets out, are that layer. They are the scaf-

folding on which physics is built: the conditions needed for anything to be measured, for causes to come in order, and for conservation and symmetry.

The whole of this book is an exercise in the reversal. Each chapter takes a result of physics and reads it as a case of a structure the metaphysics describes, with the claim that the structure is the reason the result holds. The reading runs from the basis to the case. It never runs the other way, from the case to the basis, because a case cannot see its own basis, any more than an interaction within a domain can see how that domain stands to another.

The Right Test for Each Kind of Claim

How would anyone know whether the claims of this framework are right? Each kind of claim has its own right test, and the six questions say which test belongs to which claim. Three of the questions are needed here. A why question asks for the principle from which an outcome follows; its answer is an *explanation*. A how question asks for the path by which an outcome comes about; its answer is a *prediction*, an instruction that, followed, yields the outcome. (Physics answers how questions, as the section on what physics answers showed, with a mechanism, an equation of motion; a mechanism, followed forward from a starting state, is just such an instruction, and so the same answer is a prediction.) A what question asks what is there, what is meant, what is the case; its answer is a *description*. Physics, in its explaining, answers why, in the shallow sense, by placing a case under a law. Technology, and physics put to use for prediction, answer how. Metaphysics, in its work of describing and defining, answers what: what comparison is, what a theory is, what a domain is, what it is for something to exist or to be real.

A claim is *falsifiable* when it forbids something: when it says that a specified observation will not occur, so that the observation, if it occurred, would show the claim false. The philosopher Karl Popper set out this test,

and his book on it appeared in English in 1959. Forbidding is exactly the structure of an explanation and of a prediction. A law that explains why the apple falls forbids the apple to rise. An instruction that predicts how the eclipse will happen forbids it to happen at another hour. The test of falsifiability is a test of why and of how, and it is a good one.

A description says what is there, and it forbids nothing. When the framework says that comparison has six intrinsics, or that a theory has three parts, a world, a model and the links between them, it says what is meant, what is there when one looks at the structure of comparing, and it names no observation that will fail to occur. A description is found adequate or inadequate. *An Immanent Metaphysics* gives the test of adequacy directly: a description is judged by a sense of its significance, its completeness and its relevance, and by no fixed procedure. The three questions of that test are whether a description describes what is there, whether the description leaves out anything that matters, and whether the description bears on what one wanted to understand. These questions are no weaker than the test of falsifiability. They are questions of a different kind, the questions one asks of a definition, a map, a scheme of classification, or a metaphysics.

So the metaphysics, as description and definition, is found in every act of comparison, and logic and proof are themselves comparisons that use it. It is not proved in the way a theorem is proved inside a formal system of rules. Nobody disproves the concept of comparison; one uses it, or fails to use it. If a description of comparison is muddled, or incomplete, or beside the point, one says so, and that is the criticism proper to it. To demand that a description forbid an observation is to ask a what to behave like a why. That demand is a *category error*, the mistake of applying a concept where the concept does not belong.

So a reader who asks of any metaphysics whether it is falsifiable finds the answer by asking which kind of question each of its claims answers. The test of falsifiability belongs to claims of why and how, so it is the right test for physics. Metaphysics is open to the criticism of its own kind, which asks the three questions of adequacy.

The same point can be put by levels, and the two levels use the word falsifiable in two senses. At the first level, the usual sense, a single observation within a domain can contradict a law of that domain. The axioms are no law of any one domain, so at that level they are judged by the tests proper to a description, and no single observation contradicts them. At the second level, the level of consequences, the structural results drawn from the axioms reach into physics, and there experiment bears on them; at that level metaphysics done this way is falsifiable.

The framework itself is description. The readings and exclusions of this book step from description toward explanation: they say that a result holds for a structural reason, and in the strongest cases they say that something will not be found. Four of them bear on this chapter. The first exclusion is no superdeterministic completion: no account of quantum theory in which the experimenters' choices of setting and the particles' answers were fixed together, from the start, by one shared past. In Chapter 11 I showed that the superdeterministic escape is not there. The second exclusion is no signal across the boundary: nothing passes between two events outside each other's light cones so that a record at one event depends on an act at the other event. Chapter 12 showed why there is no such signal. The third exclusion is no faithful reversal of time: no reversed process passes back through exactly the states that the original process passed through and arrives at exactly the condition that the original process started from. Chapter 13 showed why there is no faithful reversal. The fourth is no single theory that finds both survivors perfected: no one description makes perfect both the smooth survivor that goes somewhere (continuity with asymmetry) and the countable survivor that is the same everywhere (symmetry with discontinuity); Chapter 17 showed why. Those four are consequences of the structure of comparison that reach into physics, and there measurement speaks to them, as it speaks to any claim shaped like a why. The test of adequacy belongs to the description of comparison, and the test of measurement belongs to these consequences.

The Theory Turned on Itself

The framework is itself a theory. Its object domain is the relation between the self and the world. Its model is the axioms, the modalities, the intrinsics and the theorem. Its links are the identifications by which that structure is tied to particular fields of experience, and the readings in this book are examples of them. Being a theory, it falls under what it says about theories.

The first thing it says about theories is that the model is omniscient, the object transcendent, and the object goes beyond the model. So the framework, as a structural description, is exceeded by what it describes. The first axiom says the same in another form: the immanent (practice, process, interaction) is more basic than the omniscient (theory, structure, pattern). The theory of comparison is less basic than the comparing it describes.

The theorem makes the excess exact. The framework, as a description, is symmetric with the reality it describes: it shows the same structure, in the setting of a text. So it cannot be perfectly continuous with that reality. There are gaps, edges and breaks between the description and what it describes, and they are features of every description whatever, this one included.

The framework refers to itself here, and that self-reference is no circle. A *circular* argument uses its conclusion as one of its premises, and so establishes nothing. A *self-referential* description counts its own activity among the things it describes. The framework describes comparison, and the description is itself a comparison. It describes the structure of theory, and it is itself a theory with that structure. It describes the making of links, and its own links are cases of that making. The framework applies to its own application. Ordinary theories describe things that do not include themselves: a theory of gases is not a gas. The framework is unusual, in that what it is about includes the act of making theories, and so includes the framework. Think of a grammar of English written in English: the book is itself a piece of English, and its own sentences follow the rules

it states. I call that *self-descriptive closure*, and it is what makes this part of the chapter possible: the framework can say what its own limits are, because its own limits are cases of what it describes.

This self-reference is the opposite of the one Kurt Gödel used. In 1931 Gödel showed how to build, for any formal system of rules rich enough for arithmetic, a sentence that turns the system against itself and finds what the system cannot prove. Gödel's incompleteness is a case of the Incommensuration Theorem, in the domain of formal systems. In such a system every proof compares premise with conclusion, so a system rich enough to describe its own proofs compares itself, and there consistency and completeness cannot both hold, as my book Incommensuration; The Theorem and the Three Axioms sets out. The framework's self-reference confirms the framework, where Gödel's turns a system against itself: describing comparison is itself a comparison, and in being made, the description shows the structure it describes. The description closes on itself at three levels: first, the intrinsics, the parts of every comparison; second, the theorem, which is about comparison; and third, the recognition that the theorem is itself an act of comparison. A fourth level adds nothing new: recognizing that recognition would be one more act of comparison, the kind the third level already names. Chapter 4 counted a different closing, the axioms described by means of the axioms, which closes at its second step; the three levels here are the levels of the description of comparison. What no description contains is the act of comparing itself, which is the ground of the description, and no further piece of it.

Three Things the Framework Cannot Do

The framework names three things it cannot do, and they follow from the three modalities.

The first limit is the particular, which the framework does not describe. The form of the physical world, its algebra, its symmetries and its charges, is reached from comparison in other volumes of this series, as the section

on physics as instance said. The values of this universe, the masses of the particles, the way the particles mix, and the strengths of the forces, go beyond the framework in the way a world goes beyond its model; and the single event, which outcome happens and which tomorrow comes, no description reaches at all.

The second limit is the edge of the unknowable, which the framework marks and cannot enter. Chapter 18 showed that the *unknowable* is what lies beyond the reach of comparison itself, as distinct from the not yet known, and that it has a shape: three horizons, two physical and one on the side of the self. The framework can say: here is where knowledge ends and the unknowable begins. It cannot say what lies beyond, because to say it would be to know it, and to know it would be to bring it within the edge of the known, and then it would not have been the unknowable after all. The physical location of the edge is the pair of extremes that the last two chapters found: at the small end, the scale marked by Planck's constant; at the large end, the scale marked by the speed of light; the two constants that bound the domain of comparison, with everything that can be measured lying between them. At each extreme the theorem makes complete specification impossible, since to specify anything takes a comparison, and at the edge the conditions for comparison give out. The framework identifies these limits. It does not pass through them.

The black hole comes back here because it is the one place where this second limit can be seen in a single object. It is the one place in physics where the two extremes meet, where the very large closes on the very small, and Chapter 18 treated it. The point at its center, the *singularity*, as general relativity describes it, joins perfect asymmetry with perfect discontinuity, which the theorem forbids; the horizon is one-way by the smooth tipping of every light cone inward, which the theorem permits. The framework can say that much and no more; what replaces the singularity is a matter for physics.

The third limit is the living process of discovering, working through and using the framework, which the framework cannot capture. The third limit is the subtlest of the three limits. The framework is symmetric, struc-

tural: a description. The discovering of the framework, the working through of it, and the changes of view in someone who comes to see comparison this way are continuous: a process, lived. The theorem says that the description cannot be perfectly continuous with the process. So the lived experience of the framework cannot be passed on by the framework. It can be pointed toward, shown by example and suggested by story, and, by necessity, it cannot be handed over as structure. The description and the doing are a dual pair. They cannot be made one, for the same reason that the two great theories of physics cannot, the reason Chapter 17 gave: making them one would perfect symmetry and continuity together, in one description.

The Blind Spot Is a Place to See From

A blind spot is usually thought of as an absence. The framework's blind spot, the living process of discovering, working through and using the framework, is a place to see from, and some things can be seen from it. From the side of structure, the process is what the theory describes and cannot be. From the side of process, from the lived experience of anyone who uses the framework, the theory is what the process uses and is never used up by. Neither side is complete without the other. That neither structure nor process is complete without the other is the conjugacy of the first axiom again, applied now to the framework and to the practice of the framework. And the framework, unlike most theories, knows this about itself, because it predicts its own incompleteness and locates it exactly: in the gap between the symmetric and the continuous, between structure and process, between the theory of comparison and the act of comparing.

The incompleteness that the framework predicts in itself has a consequence for how the framework must be presented, and a consequence for physics. For presentation: beside the structural arguments there must be something else, story, example and personal encounter, because process can be passed on only by process. A book of arguments cannot sup-

ply that; it can say that what it cannot supply is required. For physics: the same structure holds of every theory, and of physics in particular. The mathematics of physics cannot contain the doing of physics. The judgment of which measurement to make, the recognition that a result does not match, the choice to chase one oddity instead of another, the moment at which a physicist sees that the model has been outrun: none of these is in the model, and none could be, since they are immanent and the model is omniscient. Physics has always known this in practice, and has sometimes been embarrassed by it in principle. The framework says there is nothing to be embarrassed by. The blind spot of a theory is where its practice lives.

The blind spot of the framework is also where the framework's practice lives, and that finding is the deepest result I reach about the framework itself. The theory that finds the limits of all theories finds its own limits as a case of what it finds, and so its own limits bear the theory out. A framework that claimed to have no blind spot would be claiming to be both perfectly symmetric with what it describes and perfectly continuous with it, and that is the one thing the framework says cannot be. That the framework closes on itself as a description, and that the act of comparing stays outside every description, are together the framework's evidence for itself.

The Standard Model as Inherited Code

What does all this say about the most successful theory physics has? Start with a picture from software. Every experienced software engineer has met a *legacy system*: a large program written and rewritten over decades, by many hands, with clashing habits, incomplete notes, and no single designer who understands the whole of it. The system works: it gives correct answers for the inputs it is known to handle. But no one can explain why it works, and every attempt to change it risks breaking something that people rely on and do not understand.

The *Standard Model* of particle physics, the theory of the particles of matter and of the forces other than gravity that act between them, is such a system. It has been built over more than fifty years, by thousands of contributors, using several kinds of mathematics, with numbers that are measured instead of worked out, and with tensions inside it that everyone acknowledges and no one has resolved. Three of those tensions have names. The hierarchy problem asks why the mass of the Higgs particle (the particle of the Higgs field, which gives the other particles their masses) is so very much smaller than the scale at which gravity becomes strong. The strong CP problem asks why the strong force treats a process and its mirror image, with matter and antimatter exchanged, so exactly alike (antimatter is matter whose particles carry the opposite charges). The cosmological constant problem asks why the energy of empty space is so very much smaller than estimates of it made with the theory. And the Standard Model gives correct answers. Its predictions have been confirmed to extraordinary precision: for the strength of the electron's magnetism, to about one part in a trillion, in the measurement that Fan, Myers, Sukra and Gabrielse reported in 2023. That figure is measured: the theory's computed value and the laboratory's value agree to within it. By any practical measure it is the most successful scientific theory ever built. And the Standard Model is a mess.

Faced with a legacy system, the first job is to sort it. Not everything in the store of files is the same kind of thing. The store of files holds code that runs, settings, the record of test results, and commentary. Treating these as interchangeable is the first and the commonest mistake.

The first kind is the *action machinery*: the code that runs. In the Standard Model it is the master equation from which the rest is computed (physicists call it the Lagrangian), the method of adding up every possible history of a process (the path integral), the rules for computing each possible process step by step (the Feynman rules), and the equations that say how the theory's numbers change with the energy at which one looks (the renormalization group equations). This machinery turns inputs (starting states, motions, the labels that say which particle is which) into

outputs (how often collisions of each kind happen, how fast particles decay, the colors of light they give off). It is the engine that turns inputs into outputs.

The second kind is the *configuration*: the settings you choose before you press "run". In the Standard Model they include the choice of symmetry for its three forces (the strong force, the weak force and electromagnetism); the particles of matter, the quarks and the leptons (the electron and its kin, which feel no strong force), which come in three families, three copies of one set of particles, alike except in mass; the arrangement of the Higgs field, which gives the particles their masses; and the free numbers, which are measured: the strengths of the forces, the masses, the angles that say how particles of one family mix with those of another, and the phases that make matter and antimatter behave slightly differently. There are nineteen such numbers if the neutrinos (particles of matter with no charge and almost no mass) have no mass at all, and twenty-six once each neutrino has a mass, which adds three masses, three mixing angles and one phase. Those counts are counts of the inputs the theory must be given; each number's value is measured.

The third kind is the *outputs*: the computed results, the predictions the Standard Model makes when it is set up and run. How often each kind of collision happens; how often a particle decays one way rather than another; the masses of particles; the mixing angles and the differences between matter and antimatter. The outputs are what you check against experiment.

The fourth kind is the *documentation*: the explanation wrapped around the code. In the Standard Model it is the symmetry principles, the arguments about what is natural, the hierarchy problem, the plans for development (grand unification of the forces, supersymmetry, string theory), and the guiding ideas (the symmetry of gauge, the breaking of a symmetry by the state the world settles into, and the Higgs mechanism, the way the Higgs field gives the particles their masses). The documentation tells you what the system is supposed to be doing. In a legacy system, the documentation is often out of date, incomplete, or wrong.

The four kinds must be treated separately, and in the right order. The commonest mistake with the Standard Model is to begin with the outputs and work backward, or to take the documentation at its word. Both go wrong. The outputs mean something only once the machinery and the settings are understood, and the documentation may or may not describe what the code actually does. Take the four kinds in the right order: read the documentation first, for a general sense of what the system is meant to do; then set it aside, and trust nothing in it until the code confirms it. Second, find the settings: which parts are loaded, which numbers are set. Third, trace the machinery: read the code, and see how the settings flow through it to make the outputs. Then take the machinery and the settings together, as one fixed rule that turns inputs into outputs. Fourth, and only fourth, examine the outputs. Now any gap between the predicted outputs and the observed ones can be traced: is the error in the machinery, in the settings, or in the observation? The documentation is consulted for orientation and trusted for nothing.

The order for a legacy system, the documentation first, then the settings, then the machinery, and the outputs last, is exactly the order in which I approach the Standard Model. Its documentation (the arguments about what is natural, the hierarchy arguments, the stories of grand unification) is consulted for orientation and trusted for nothing. Its settings (the symmetry groups, the three families, the masses) are identified and analyzed for their structure, in the books of this series on the kinds of particle. Its machinery is examined to see whether it obeys the framework's structural constraints. And its outputs become readable only after the machinery and the settings are both understood.

In a well-designed computer, the line between program and data is sharp. The program is fixed, and the data varies. The program works on the data, and the data never rewrites the program. Engineers call a design that keeps the two in separate memories the *Harvard architecture*. In the Standard Model, the line leaks. Many numbers that look like settings are in fact results computed partway through: outputs of one stage of the calculation, fed in as inputs to the next. The strength of each force "runs"

with energy: its value at one energy is computed from its value at another, by the renormalization group equations. The masses of the particles that are built of smaller ones, the proton and the neutron, are computed from the workings of the strong force that binds their quarks together. The shape of the Higgs field's energy is computed from corrections that involve all the other numbers of the theory.

So the Standard Model is no Harvard architecture. It is closer to a program that rewrites itself, or to a system in which even the rules of the programming language can change. Where "what the program is" ends and "what the program works on" begins shifts with the energy, the method of approximation and the level of description. That leak is why the Standard Model is hard to rebuild: you cannot separate the code from the data, because the code makes some of the data it then uses. Any attempt to fix a definite list of the "true inputs" needs a choice of energy, of method of approximation and of level of description, and each of those choices is a convention, and no fact.

The framework stands to the Standard Model as a computer's hardware stands to its software. The hardware does not decide which program runs: many different programs can run on the same machine. But every program must obey the machine's limits. No program can break the machine's set of basic instructions, its layout of memory, or its rules of timing. The framework specifies the hardware. The three axioms are the set of basic instructions, the operations the machine can perform. The six intrinsics are its registers, the few slots the machine works on directly in every step. The Incommensuration Theorem is its rule of timing, the rule for what can be done together. The structure of the domain is its layout of memory: the small-end limit at Planck's constant, the large-end limit at the speed of light, and the middle limit, at the act of comparison itself, which no constant of physics marks. The Standard Model is the software. It runs on the framework's hardware, and it must obey the axioms, the intrinsics, the theorem and the structure of the domain. Whatever those constraints leave open in a particular theory is the software's business, not the hardware's. How much they leave open is settled item by item in the

books on the kinds of particle: the gauge group, for example, is derived from the constraints, and so it is a question for the hardware, and no free choice of the software.

The picture clears up a persistent confusion. The framework is no rival theory to the Standard Model. It proposes no different forces and no different equations. The particles it states, in the books of this series on the kinds of particle, are the measured ones, with each neutrino having a partner that spins the other way relative to its motion, and a mass of the ordinary kind, which physicists call a Dirac mass; that is the case of the twenty-six numbers counted above. The framework names the structural constraints that every correct theory of particles and forces must meet, and the books on the kinds of particle argue that the Standard Model meets them. The framework explains why the Standard Model has the structure it has. The Standard Model says what its particular content is.

The Map, the Compass and the Position

A traveler crossing unknown country uses three tools. A map gives the layout of the country. A compass gives direction, measured against a reference outside the traveler. And a fix of position gives the traveler's own location in the country, found by an actual sighting of a landmark. The three form a triple. The map has the nature of the omniscient: a complete, outside, structural picture, laid out and held in view as a whole. The compass has the nature of the transcendent: it is the standard against which every bearing is read, and it points to something beyond the traveler's immediate setting. The position has the nature of the immanent: the actual here and now of the traveler meeting the country.

In physics the Standard Model, as a theory, gives the map, since in every pairing of theory with practice, the theory is the omniscient side. The measurements give the position: each finding is an act, here and now. The framework gives the compass: the frame that relates the theory to its findings, which is the transcendent part when the practice is what is in

view. This matching differs from two earlier ones, by the rule of Chapter 18 that what is in view decides the match. When a single theory was in view, its links, the acts of comparing, were immanent; when the framework was taken as a theory of its own, its model was omniscient. Here the practice of physics as a whole is in view, theory and measurement together, and the framework is the standard against which both are read. A map with no position is an abstraction that cannot be applied. A position with no map is an experience that cannot be understood. The compass is what lets the map and the position be brought into line.

Computing gives a second picture of where the framework stands. A process of computation can be seen at three levels. Take sorting, the putting of things in order, as the example. The first level is the running program. The second level is the rewriting of one program into another program that computes the same thing. The third level is the design of the language in which any such program is written, and the deepest power of the third level is that the third level defines what 'sorted' means. At the first level the Standard Model is a procedure for sorting: it takes unsorted inputs, the starting states, and gives sorted outputs, the possible end results, each ranked by its weight of likelihood. Its several methods of calculation (the Feynman diagrams, the drawings by which the Feynman rules are applied; the path integral; calculation on a grid of points) are several procedures for sorting. They give the same order by different steps.

At the second level the question is whether the different procedures compute the same thing. Does adding up every possible history give the same quantum theory as the rules applied directly to the equations (the method called canonical quantization)? Does the strong force computed on a grid of points by computer agree with the strong force computed as a series of small corrections, where both apply? Work at this level turns one sorting procedure into another while keeping the order the same.

At the third level the question is: what must every sorting procedure obey? What does it mean for the outputs to be "correctly sorted"? What structure must every valid computation keep? The framework answers these questions of the third level. The three axioms define "correctly sor-

ted" for every representation of the physical world. The six intrinsics define the categories that must be kept. The Incommensuration Theorem holds of every comparison, and so it sets the trade between symmetry and continuity that the findings of every theory must respect: the two cannot both be perfected in what a single theory finds. These requirements are properties of no particular procedure. They are properties of being sorted at all.

Why Mathematics Works in Physics, and Why It Is Never Complete

In 1960 the physicist Eugene Wigner asked why mathematics, worked out by mathematicians for their own reasons, describes the physical world with such precision. The number pi, defined from the circle, turns up in the way errors of measurement spread, and in the equation of quantum mechanics. The complex numbers, invented to solve a problem in algebra, are needed for the amplitudes of quantum theory, the quantities from which its probabilities are computed. And the prediction of quantum electrodynamics, the quantum theory of electrons and light, agrees with the measured magnetism of the electron to about one part in a trillion, the measured figure quoted earlier in this chapter.

The fit is no accident. A physical theory is a comparison of its mathematics with its findings. By the Root Tautology, the mathematics and the physics are two sides of one act of comparison, built from the same six intrinsics and the same four compounds of sameness and difference. The reason for the fit is the principle of cross-modal reach, which Chapter 17 used for the two great theories. Its first part says that an exact theory's mathematics and its findings are of opposite characters: smooth mathematics for countable findings, symmetric mathematics for a smooth setting. Its second part says that the mathematics reaches its findings densely, a mathematician's word meaning as close to each finding as one likes, without ever covering them all: a continuous mathematics gives values to

separate outcomes to any precision required, and a symmetric mathematics carries a smooth content along any path to any precision required. That dense reach is why mathematics works so well in physics. And mathematics has structures of both characters: groups and quantities that stay the same under change, on the one side; smooth spaces and continuous functions, on the other. So for findings of either character it has a mathematics of the opposite character, as the first part of the principle requires.

Pi and the complex numbers show the shared structure in particular cases. Pi stands for the symmetry of turning, and it turns up wherever that symmetry does: in a spread of errors over two directions that can be exchanged for each other, as much as in a circle. The amplitudes of quantum theory have a size and a phase, and the phase comes back to where it started after a full turn; the complex numbers are the smallest system of numbers that carries that return. In each case the mathematics was made to capture an idea of structure, and the same idea is at work in the physics, because both are sides of the same comparison.

The second part of the principle also says that the reach is never complete, so how well mathematics works in physics has a limit. The structure that makes mathematics work in physics is the same structure that sets the limit. A mathematics that contained its findings exactly would perfect both characters in one, which the theorem forbids, and so every theory meets a *residual*, a remainder that no lawful description fixes. The mathematical description of physics comes as close to the physical process as one likes, and never becomes it. The power of the description and its incompleteness come from one source, and neither can be had without the other.

Where Physics Stops, and Why

We now know what physics reaches, and what it cannot. Physics reaches every how, when and where of its world as closely as one likes, and physics answers the what questions of its world, and the shallow why. But

physics never holds the whole of its world. Its mathematics fits the world because the mathematics and the world are two sides of one comparison, and the same structure that gives the fit guarantees a remainder. Every theory has three parts, the world, the model and the links between them, and the links come first. The world always goes beyond the model, and the model always carries structure, such as gauge and coordinates, that refers to nothing in the world. The deep why, the question of why there is a world at all, asks for a possibility that is inside no world, and so no measurement can return it. And physics cannot reach its own basis: it uses the ideas of object, model and measurement, and cannot ground them.

Why does this matter for the book's question? Each earlier exclusion found something that cannot be, because the thing asked for would need comparison to do what comparison cannot do. In this chapter I find, for physics as a whole, the same kind of thing: something that cannot be, because the thing asked for would need comparison to do what comparison cannot do. A complete model, one that holds its world with nothing left over, would be a specification of what goes beyond specification, and so it names nothing. A physics that measures the deep why would be a measurement of what is inside no world, and so it names nothing either. What lies beyond physics differs from physics in kind, as warmth differs from length, and distance has nothing to do with it.

Summary

There are six kinds of question, and each asks for one of the six measures: what for mass, where for space, how for force, when for time, and who and why for probability and possibility. That match of questions to measures is a claim of the metaphysics itself. Physics, the method of measurement within a world, answers what, where, how and when, and the shallow why, which it answers with a law, a more general what. It does not answer the deep why, because the possibility from which a world is chosen is not inside the world, and measurement is. That reading of what physics

does has the match of questions to measures as its structural reason. Metaphysics asks what questions too, about the structure of comparison and of possibility, and nearly every claim in this book is a what claim. In this chapter I do not say that metaphysics answers what physics cannot. I locate the deep why, and I say what kind of answer the deep why would take.

A theory is a world, a model and the links between them, and the links are more basic than either. World and model are conjugate, and each goes beyond the other: the world goes beyond the mathematics, and the mathematics carries gauge and coordinate structure that answers to nothing in the world. Neither excess is a defect: both are what two modalities that cannot be exchanged look like when they are joined. The Appendix lists every claim of the chapter with its kind.

Physics studies the interactions within a domain; metaphysics studies the relations between domains. Physics is an instance of metaphysics, and not the reverse, because physics uses the ideas of object, model and measurement and cannot ground them. The reversal says which of the two, physics or metaphysics, is the logical basis of the other. It does not make physics subordinate to metaphysics in practice, and physics does what metaphysics cannot: it finds the particular. The test of falsifiability belongs to claims of why and how; a description, a what, is judged by its significance, its completeness and its relevance, and the axioms are judged that way. The four exclusions named here (no superdeterministic completion, no signal across the boundary, no faithful reversal of time, no single theory that finds both survivors perfected) are claims shaped like a why that reach into physics, and measurement speaks to them; they are the exclusions the earlier chapters made.

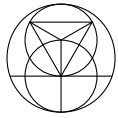
The framework, being a theory, falls under what it says about theories. It is symmetric with what it describes, and so not continuous with it. It cannot describe the particular, cannot enter the unknowable, and cannot pass on the process of its own discovery. It predicted all three, and located them, and a limit that a theory predicts of itself is a case of the theory. The Standard Model is inherited code, best sorted into machinery, settings,

outputs and documentation, and best read with the map, the compass and the position held together: the theory gives the map, the measurements give the position, and the framework gives the compass. Mathematics reaches the physical process as closely as one likes and never coincides with it. Physics cannot reach its own basis, and metaphysics cannot reach its own practice. The reach of each ends where the other begins, and the place where they meet (the links, the comparison, the measurement) is more basic than either. The place where physics and metaphysics meet is where physics is done, and that same place is where I write this book from.

If physics cannot hold the whole of its world, could the world still be a program, running on a computer somewhere else?

PART SIX

THE WORLD IS NOT A PROGRAM



Chapter 21

Are We Living in a Simulation?

A Computer With No Limits

Could our world be a program running on someone's computer in another world? Could everything we see, ourselves included, be the output of a machine we will never find?

Picture the best computer anyone could ask for. It is as large as you like. Its processor runs at any speed you name, and its memory has no limit. Now write for it a program that copies the physical universe, a program that obeys every law we know: the equations of general relativity for the shape of space and time, the equations of quantum theory for matter and light, the statistics of heat and of large numbers of particles. Could that program be made so good that nothing could ever tell it apart from this world?

A few words need defining first. A *computation* is a procedure that goes, step by step and by fixed rules, from one state to the next, where each state is a finite string of symbols, such as ones and zeros. A *program* is the written set of those rules. A *simulation* is a computation built to copy the behavior of some physical thing, so that what it puts out can be set beside what that thing is found to do. The *simulation hypothesis* is the proposal that our world is such a computation, running in a more basic world, and that we, and everything we find, are states of that computation.

People have argued for a long time about whether this universe rests, at bottom, on something digital, something that could in principle be written down in bits, the ones and zeros of computer information, the way a hologram stored as a computer file holds a whole three-dimensional scene in a finite string of ones and zeros, or whether it rests on something like what physics has always called continuous space. Take the two possible answers to the question about the perfect program. If the program could match this world, then continuity, the idea that space and time are smooth without end, would be at best a convenient way of describing things and at worst an illusion. If it could not, then there is something in the nature of what exists that is deeper than any simulation, more basic than anything that step-by-step logic alone can hold.

So the question only looks like a question about computers. The question about the perfect program is a question about whether what exists is, at bottom, continuous or discrete: smooth, with no gaps anywhere, or made of separate pieces, like a grid. Whether what exists is continuous or discrete is a question about the nature of the real. The answer to that question, whichever way the answer falls, is a claim of metaphysics, the study of what anything must be in order to be at all. The question is also a good door into the whole account of this book. Working through the simulation question trains the mind to ask whether any test, even in principle, could tell two things apart, which is the thinking the account of this book needs; and the question gives, on its own, a result that matters. How a self stands to the world it lives in, the knower to the known, turns out to be connected with how symmetry and continuity, the two kinds of law of Chapter 2, stand to each other, in this way: computation lies on the side of symmetry, and the self lies on the side of continuity. This chapter follows the simulation question, and with it the relation of the self to its world, to its end: our world is not a simulation run in some other world.

The Principle of Identity

The argument turns on a principle older than any of the physics it touches. The *Principle of Identity* says that what cannot be told apart, by any test that could even be imagined, must be counted as the same. In the tradition of the philosopher Gottfried Wilhelm Leibniz, who used it in his letters to Samuel Clarke of 1715 and 1716, it is called the identity of indiscernibles. If no test can be conceived, even in principle, that would separate two situations, then it is not reasonable to treat them as two. A difference that could never show itself in any comparison whatsoever is no difference at all; such a difference is a name with nothing behind the name.

The Principle of Identity follows from what comparison is, and the Principle of Identity is not a law of physics. Comparison, as Chapter 1 said, is the act of finding whether two things are the same or different, in some respect. To say that two things differ is to say that some content differs within some shared context: that something found differs, in a setting the two have in common. To say that some content differs within a shared context is exactly to say that some comparison could register the difference. Where no comparison could register it, the claim that the two differ says nothing. So the principle belongs to the account of comparison that the whole of this book stands on, and nothing has to be brought in from anywhere else to support it.

Applied to the simulation question, the principle does two things at once. First, it sets the standard for an answer. Suppose a simulation were so perfect that no conceivable test could show a difference between it and a continuous world. Then simulation and world would not be two kinds of thing, and the hypothesis that we are simulated would be idle: neither true nor false, but empty. Second, it sets that standard at the level of principle, and not at the level of practice. Whether the argument succeeds is independent of the instruments we own. Whether the argument succeeds depends on whether any difference at all can be drawn between a simula-

tion and a continuous world. If such a difference exists, then a simulation and a continuous world are different kinds of thing, whether or not anyone could ever carry out the measurement that would show it.

So the question becomes exact. Is there a test, sound in its logic and definable in principle, that would tell a discrete simulation from a continuous world? I call such a test a *test in principle*: a test whose logic works whether or not any instrument can perform it. If there is a test in principle, the two are different in kind. If there is no test in principle, the question dissolves. Either way we learn something.

Three Tests That Do Not Decide

Three tests come to mind first whenever the question is asked. Seeing exactly how each one falls short is worth the time, because the way they fall short points toward the test that succeeds.

The first test is repetition. Do the same thing again and again, and see whether the results are ever exactly alike. A computation run twice from the same start gives the same outputs, while in a continuous world the passing of time alone ensures that no two moments are ever quite the same. The test does not settle the question. A simulation of any sophistication can carry a source of variation built into its program, so that the small differences the test looks for are a feature of the program and not a flaw in it. The designers, knowing of the test, build in whatever it is meant to detect. Repetition is answered by *pseudo-randomness*, numbers made by a fixed recipe that only look random, and pseudo-randomness is cheap.

The second test is simultaneity. An ant colony and a stock exchange go about their business at the same time, each on its own; an enormous number of processes happen at once throughout the universe, each seemingly unconnected to the rest. Could a single computation carry them all? It could, by parallel computation: a computer can run many separate lines of work, called threads, side by side, and each process that runs at once and

on its own in the world is copied by a thread that runs on its own in the computer. The number of threads is a matter of hardware. Simultaneity, by itself, gives no difference in kind.

The third test is the sheer amount of parallel work. The universe holds an enormous number of processes running at once at every scale, from the restless quantum activity of empty space to the gravitational pulls among clusters of galaxies, and a simulation would have to compute all of them together. Modeling even one cubic centimeter of ordinary matter, about the size of a sugar cube, in full quantum detail needs more parallel processing than anything now conceivable. The need for more parallel processing than anything now conceivable is a limit of practice. By the standard of the Principle of Identity, a limit of practice does not tell a simulation from the world in principle, and the search has to go further.

The three tests share one pattern. Each time, the simulation passes the test by adding more of something it already has: more built-in variation, more threads, more processors. Every one of the three tests asked only for what step-by-step logic can, in principle, supply. So each test that falls short still does work: it narrows the field, and it drives the question toward a difference that no amount of extra step-by-step resources could close.

Randomness, Halving, and the Smallest Length

Three further considerations bring that difference into view.

The first of the three further considerations is randomness itself. A computation works by procedure, step by step from inputs to outputs. Even a computation that uses randomness has to make it by some procedure, and what a procedure makes is, by that very fact, not random in the hard sense. *Hard randomness*, as Chapter 8 showed, is an outcome that no description of what came before fixes, however complete that description is; it is what no recipe of any kind puts out. A computation can imitate hard randomness. It cannot produce it. Nature contains hard randomness,

as Chapter 8 also showed, and so a simulation running on its own rule can never fully copy nature; its imitation must carry some trace, however faint, of the recipe that made it. A simulation could instead draw hard randomness from a physical process in the world that runs it. Then its randomness would be that world's randomness, carried through the computation, and not a randomness of its own. Either way a simulation already differs in principle from nature: running on its own rule, its randomness carries a trace of the recipe that made it; drawing on the world that runs it, its randomness is that world's and not its own. An earlier book of mine, Third Resolutions in Western Philosophy, settled the simulation question by this road: hard randomness, the direction of time and the inner side of experience each require the other two, so none of them can be present without the others. I walk a second road first, the road through continuity, and then I walk the road through randomness as well, near the end. The two roads arrive at the same place; neither road rests on any step of the other, and each road is the stronger for that.

The second consideration is the structure of space. Take any stretch of space and split it in half. Take one half and split it again, and again, and again, without end. In the idea of continuous space there is no smallest distance; splitting never reaches a piece that cannot be split. A simulation, however, must represent space by a finite number of points of data, and however many points it uses, there are always gaps between neighboring points. Packing the points more tightly brings the simulation closer to continuity, and it never arrives.

Here the natural objection comes in, and it deserves a careful answer. Physics has a scale called the *Planck length* (Chapter 14 introduced it). The literature gives it as about 1.6 times ten to the minus thirty-five meters: written out in meters, a decimal point, then thirty-four zeros, then the digits one and six. That figure is quoted from the physics literature, and it is built from three measured constants: Planck's constant, the speed of light, and Newton's constant of gravitation. The Planck length is the scale at which two descriptions change places: on one side of that scale the description of the arena, the space and time that are the setting of every

measurement, dominates, and on the other side the description of what is in the arena dominates. It is often taken to be a smallest length, below which distance and duration stop applying. Does it rescue the simulation hypothesis? If there is a smallest unit, perhaps reality is already a grid, and there is no continuity left to simulate.

The first answer comes from measurement. At every scale yet probed, length has been found without a least step, so the Planck length is no smallest length. In 2009 A. A. Abdo and colleagues reported that light from a distant gamma-ray burst, a brief and enormous flash from another galaxy, arrived together across very different colors after its long crossing; that arrival places any grain of space that would slow light in proportion to the light's energy below the Planck length.

The second answer is the hinge of the whole inquiry. The microscopic limit is the inner edge of the domain of comparison, the range in which comparisons can be made, and Planck's constant marks that edge. Planck's constant is a fixed amount of what physicists call action, an energy multiplied by a time, as Chapters 18 and 19 described the constant. The microscopic limit is a boundary on what can be measured. Chapter 8 showed that this boundary is categorical, part of what measuring is, and not a matter of how good our instruments are: no better instrument will cross it. A boundary on measurement is not a graininess in what exists. The step from "nothing below this scale can be compared" to "nothing below this scale exists, and what exists above it is a grid" is exactly the step the Principle of Identity does not permit. The principle says that what cannot be told apart must be counted as the same; it does not say that what cannot be measured must be counted as made of pieces. The question is whether the difference between the smooth and the grainy exists, and it was never whether we can measure that difference.

The same answer holds for the limits physicists have found on how much information a bounded region of space can hold. Those limits bound what any comparison within the domain can tell apart; they are

limits of the domain of comparison itself, and not a graininess of what exists. The Appendix gives them with the names of the physicists who found them.

The third consideration makes the hypothesis explicit. Call it the *binary hypothesis*: the proposal that all the information in the universe can be written as ones and zeros. Each one-or-zero is a *bit*, the smallest unit of computer information. On this hypothesis one could picture a vast file holding the complete state of the world at one instant, a definite and finite number of bits, and the history of the world would be the turning of one such file into the next by some rule. The binary hypothesis is the strong form of the simulation hypothesis, and it is a clean one. Whether the binary hypothesis is true is decided by a single question: is reality smooth at bottom, or grainy at bottom? Every test considered so far has now narrowed to this one choice, continuity against discontinuity, and observation alone cannot settle it. It takes an argument about structure.

What No List Can Hold

Start with the number pi, the distance around a circle divided by the distance across it. Its digits begin three point one four one five nine, and they go on without end and without ever falling into a repeating pattern. A short program can produce any digit of pi you ask for; pi can be specified as what that short program gives when it runs without limit, and only the program run without limit gives every digit. The information needed to state pi exactly as digits has no bound, and no finite string of digits, however long, is pi. Now picture a real process, in a space that is continuous, that depends on the exact value of pi, the distance around a circle divided by the distance across it. That process carries all of that information. A simulation can hold only finitely many of the digits, and its approximation, however astonishingly close, is not exact.

Someone may reply that an approximation too close for any measurement to catch is as good as exact. That reply gets its answer in two sections' time. First comes the deeper point: the gap between the finite and the endless is a gap of structure, and not a gap of quantity, of how many digits. It comes from two proofs by the mathematician Georg Cantor. In 1874 Cantor proved that the real numbers, all the points of the number line, cannot be listed. In 1891 Cantor gave his diagonal argument, which shows why the real numbers cannot be listed.

Cantor's general result of 1891 is a fact about lists. Take any collection of things. A way of marking each member of it yes or no is one assignment. Cantor showed that no list with one line for each member of the collection contains every assignment. The fact holds for finite collections as well as for endless ones. A collection of three members has eight ways of marking them yes or no (two choices for the first, times two for the second, times two for the third), and three lines cannot hold eight. In general, a collection of any number of members has two, multiplied by itself that many times, ways of marking them, and that number is always larger than the number of members. So Cantor's result is a statement about listing, and not about the size of some completed endless whole. This fact about listing is what I take from Cantor in this book, and the fact about listing is all that I need from Cantor for the argument of this chapter.

Two kinds of collection matter here. The first kind can be listed: set out in order, first, second, third, without leaving anything out. The counting numbers can be listed, and so can anything that can be matched one to one with them: the even numbers, the primes, the square numbers, the fractions, every finite string of symbols in any finite alphabet, every program that could ever be written. A collection of this first kind is *listable*. The second kind is the collection of all the ways of marking every counting number yes or no, of which the number line is the familiar form, and no list reaches all of it. This second kind is the *continuum*: the smooth whole of every possible value, with no gaps anywhere in it, in which between any two values there is always another (the word Chapter 2 used

for the smooth spread of events in general relativity). No list can be matched to the continuum because of how a list and the continuum are built, and not for any lack of ingenuity in whoever makes the list.

Cantor's diagonal argument shows this so directly that nothing is left to interpretation. Suppose someone claims to have listed every number between zero and one, each written as a decimal that goes on without end: the first number on line one, the second on line two, and so on down. Now build a new number, one digit at a time. For its first digit, choose any digit other than the first digit of the number on line one. For its second digit, choose any digit other than the second digit of the number on line two. Go on down the diagonal of the list in the same way. At every place, choose a digit other than zero or nine. The reason is a small fact about decimals: a few numbers can be written in two ways, as one half is both zero point five followed by zeros without end and zero point four followed by nines without end. A number built with no zeros and no nines has only one way of being written, so it cannot slip onto the list in a second spelling. The number built this way differs from the number on line one in its first place, from the number on line two in its second place, and from the number on any line in the place that matches that line, and so it differs from every number on the list. It is a number between zero and one, and it is not on the list. The list was never complete, and no list could be. No list can be matched to the continuum.

That no list can be matched to the continuum is *Cantor's theorem*. Cantor's theorem is a theorem in the ordinary sense of mathematics: proved, step by step, from what a list is and what a decimal number is, with nothing left to interpretation.

The Arena Has No Least Step

The arena of space and time is continuous, and the account derives it. The derivation is short, and the picture of halving carries it. Take again the stretch of space split in half, and in half again. Sameness and difference

are inseparable in every comparison (Chapter 1 found that neither can be removed without the act dissolving), and that holds of settings as of everything else: every difference between two settings holds some sameness within it. Settings can also be put together, two settings side by side making one larger setting; so any difference between two settings can be split by a finer comparison, made at a setting that lies between those two settings, and that finer comparison finds the sameness within the difference. That finer comparison is the next halving. A least step of context, the smallest possible jump from one setting to the next, would be a jump that no finer comparison could split: a difference of context with nothing between its two ends, and so with no sameness within it. A difference with no sameness within it cannot occur. My book *Incommensuration; The Theorem and the Three Axioms* gives this derivation in full. So the context has no gap and no least step: between any two settings there is a third, and nothing in the context is sudden. The absence of any gap or least step in the context is continuity.

Continuity of context is exactly what the argument against simulation needs, because a discrete computation has a least step in every context it represents: the gap between two neighboring points of its grid, the tick between two of its moments. Physics agrees with the derivation. The arena has been found smooth at every scale yet probed, and the gamma-ray burst measurement of 2009 places any grain of space that would slow light below the Planck scale.

So far the argument has had the shape of a condition: if reality rests on the continuum, then no step-by-step computation can capture it. For the arena of space and time, the setting of every comparison, the condition is no longer only a condition. The continuity of the arena is derived from what comparison is, and physics finds it. The rest of the chapter uses that continuity as what it is, a result.

The Consequence, and the Amplifier

The consequence for the simulation question is immediate. A computer works with separate, countable states. Its memory is a finite number of bits, and its computation is a finite number of steps. Alan Turing described, in 1936, the ideal computer that marks the limit of the idea: a machine with unlimited time and an endless strip of memory. Even that ideal machine produces, at every stage, an output that can be listed, and everything it could ever produce, taken all together, can itself be listed. The arena in which every physical process runs is continuous, and its positions and durations are smoothly varying quantities, not step-by-step stand-ins for them. So reality holds what no computation, which is a listing, can reach. A simulation differs from the world, and it differs in principle, not as a matter of practice only.

Now comes the reply that was set aside two sections ago: that an approximation too close for any measurement to catch is as good as exact. Surely a simulation could carry so many digits that no difference would show within the lifetime of any observer, or within the reach of any instrument. The reply is correct as far as it goes, and it does not go far enough, for two reasons.

The first reason is the Principle of Identity itself, which asks only for a test that can be defined in principle. The diagonal argument defines one. It says: here is a class of quantities, the ones no list reaches, that a continuous world carries and a discrete world cannot carry. Whether we can read those quantities off with today's apparatus is a separate question, and the test in principle that the diagonal argument defines exists whatever the answer to that question is.

The second reason is physical, and it is why the argument reaches differences that show in the behavior of the world, and not only differences that can be defined. As a very general rule, real systems are sensitive to how they start. The weather is the classic case. Edward Lorenz's work on weather, published in 1963, gave the effect its popular name, the *butterfly effect*: in the famous image, the flap of one butterfly's wings can decide

where a storm forms weeks later. Physicists call it *sensitivity to initial conditions*: the way the smallest difference in how a system starts does not stay small, but grows, and goes on growing, until it rules the behavior of the whole. So a simulation that cuts its values off at any finite precision does more than differ from the continuum by an amount too small to detect. Given time, it parts from the continuous world on the scale of everyday things. The gap between the finite and the endless, which the diagonal argument shows to be structural, is turned by sensitivity into a gap between two different histories: the history the simulation computes and the history the world runs.

Sensitivity is the amplifier, and what no list reaches is what it amplifies. Without the amplifier, the difference would be real but hidden. With it, the difference is real and, in principle, visible. So there are events in this universe that could not be foreseen from any earlier knowledge held in the form of separate symbols.

Each extra digit of precision buys a simulation a further stretch of time in which it follows a sensitive system faithfully. No fixed precision follows a sensitive system forever. The two facts both hold because they choose their numbers in a different order: name a stretch of time and a closeness first, and some precision keeps the simulation that close to the system for the whole stretch; name a precision first, and some stretch of time carries the simulation as far from the system as the system can go (the Appendix sets this out).

A simulation of space and time themselves shows the same pair of facts: it can come closer without limit, and it never arrives. A simulation of space and time, on any computer, ordinary or quantum, can be refined to come as close as one wishes to what space and time are found to be, and no stage of it ever contains them. The limit is the same for both kinds of computer, because a quantum computer's program is a finite string of symbols, and everything it reports comes from finitely many readings, which can be listed. This result concerns how far a simulation can reach,

and it says nothing about speed or cost; how many steps a given closeness takes is a question for the study of particular programs (the Appendix sets this out).

Sensitivity, which magnifies differences too small to measure into differences large enough to measure, is physical as well as mathematical. Chapter 8 read the three-body problem, the old problem of how three bodies move when each pulls on the other two, as a physical case of exactly this mechanism: differences below the limit of measurement, magnified into differences above it. The argument from the continuum and the argument from hard randomness meet at sensitivity, which magnifies differences from below the limit of measurement into differences above it. What the diagonal argument says cannot be written down is what sensitivity says will, in time, matter.

The Quantum Wave as Evidence

The way physics writes its laws gives further evidence that what exists takes part in the continuum. The Schrödinger equation, published by Erwin Schrödinger in 1926, which carries a quantum system forward in time; Albert Einstein's field equations of general relativity, completed in 1915, which give the shape of the arena; James Clerk Maxwell's equations of electricity and magnetism, from the eighteen-sixties: all are written in the mathematics of the continuum. Every law of the world that physics has so far been able to state is written in the mathematics of the continuum. One might propose that this is an accident of our tools: that the rule underneath is discrete, and the smooth equations are only how it looks at large scales. That proposal is exactly the binary hypothesis, and it has a consequence that could show.

In quantum theory, what a measurement finds is discrete. A particle is found here or there, with this spin or that (spin being the built-in turning of Chapter 1, which makes an electron act as a tiny magnet). Between measurements, quantum theory describes the system by its *wavefunction*,

the smooth mathematical wave that sets the odds of each outcome (Chapter 4 met it). Its values, called probability amplitudes, the numbers from which the odds of each outcome are worked out, vary smoothly, and the Schrödinger equation carries it forward smoothly, through values that change without any jump. The statistics of what is found, the share of times each outcome turns up over many runs, depend on that smooth change.

A discrete rule can be chosen to reproduce those statistics to any finite precision named in advance. Programs do this every day. Pairs of *entangled* particles, whose outcomes, measured far apart, agree in a pattern no separate instructions could give them (Chapter 10), are simulated on ordinary computers routinely, and the agreements come out to the precision the simulation was built for, at a cost that is multiplied by a fixed factor with each particle added to the entangled group. That cost is a limit of practice, like the sugar cube of the third test, and it is not where the difference in principle lies. The smooth change itself has no least step at any precision, and so no discrete rule can be that change. Suppose the change were in fact discrete: suppose the wavefunction were an approximation to some one underlying finite rule. That rule would have a least step. With enough trials, and enough sensitivity to reach below that step, its statistics would be found to differ from the world's. So for any one fixed program, beyond its smallest step, a sensitive enough measurement eventually tells the program's statistics from the world's. Here the difference between discrete modeling and continuous modeling is definable in principle, and it is also tied to a physical statistic: each proposed rule fixes a precision beyond which that rule's statistics would differ from the world's. Precision measurement and quantum information science (the study of how quantum systems store and process information) are approaching the range of precision at which such statistics can be tested very finely. I set out the details of that approach in the Appendix.

There is a second thing a deterministic program cannot reproduce: that each quantum outcome is fixed by nothing before it. A *deterministic* program is one whose every output is fixed by its rule and its start

(Chapter 5 gave the word its meaning: everything fully and exactly fixed). The outcomes of such a simulation are fixed by the program; the outcome of a quantum measurement is fixed by nothing before it, as the hard randomness of Chapter 8 shows, and as the free will theorem of Chapter 9 shows from physics: if experimenters are free to choose what they measure, then the particles' responses are not fixed by anything before them. A deterministic simulation of the whole universe cannot be what the universe is, because its outputs are fixed and the universe's are not.

So the evidence stands as follows. The mathematics of quantum theory is the mathematics of the continuum, as its structural foundation and not as an optional convenience, and the theory's success is therefore evidence that continuous mathematics is needed to account for discrete findings, and so evidence that what the findings are about takes part in the continuum. That step, from the success of quantum theory to the continuum, is a reading at the level of interpretation, and the argument of the earlier sections completes that reading: everything a computation produces can be listed, no list reaches the continuum, and sensitivity in time shows the difference. A cleverer discrete rule could always be proposed to mimic the smooth statistics to any finite precision. Then that argument applies again, one level up: the mimicry can be listed, the continuum cannot, and sensitivity will in time tell them apart. The discrete proposal is always one step behind, and it is behind in principle.

The Incommensuration Theorem gives one more reading here. The theorem (Chapter 3) is a structural fact about the nature of knowing: symmetry and continuity cannot both be applied, at once and absolutely, to one comparison, and neither can asymmetry and discontinuity. Of the pairings of symmetry, continuity, asymmetry and discontinuity that the theorem weighs, two survive being applied absolutely: continuity with asymmetry, smooth and going somewhere; and symmetry with discontinuity, the same everywhere and countable. The wavefunction between measurements is continuity with asymmetry: its amplitude varies smoothly, and no two points of its change are the same. The measured outcome is symmetry with discontinuity: a separate, countable result, the

same in kind wherever it is found. Quantum theory holds the two at its two levels: the smooth change in its mathematics, which exists, and the separate outcome in what it finds, which is real (in the sense of Chapter 4: to exist is to stand as structure, the same for every comparer; to be real is to be in an actual act of comparison, here and now). The interaction called measurement stands between the two. A discrete simulation of the whole world would need the discontinuous side perfected across the whole of it, and that is to require that continuity be entirely absent. The account reads the continuous wavefunction as the place where physics itself refuses that requirement: there physics keeps a continuous side, the smooth wave, in its own description of the world.

Symmetry, Continuity and the Self

The deepest thing this inquiry shows is that what can be listed and what no list reaches stand to each other as symmetry and continuity, the two kinds of law of Chapter 2, stand to each other: the two cannot both be perfected at once.

Symmetry is sameness of content across difference of context: change the setting, and what is found stays the same. It goes naturally with the finite, with what can be counted, with what exists, with everything that can be listed, shared and repeated. Picture a long row of light switches, each one either on or off. The number system of computers is built on the distinction between one and zero, two symbols that can be exchanged for each other, and so the number system has symmetry in its very construction. Each bit is identical to every other bit in everything except its position, as each switch in the row is the same as every other except for where it sits. A program is a finite string over a finite alphabet, and a computation is the changing of such strings by rule. So, in its character, computation is symmetry with discontinuity: symbols that can be exchanged, separate states, and the same rule wherever it runs. A theory written as a computation is laid out as structure, and in that sense it is purely omni-

scient, in the word of Chapter 4: perfectly structural, perfectly repeatable and perfectly third person. The role of computation is a second question. Chapter 22 sets physics, logic and computation side by side as three parts of inquiry, each playing its own part in comparison, as the three modalities do; the part a thing plays there is its role. By its role in that triple, computation is transcendent (the mode of possibility): a procedure is defined over the whole range of its possible inputs, and each run picks one path through that range. Both of these hold, each answering its own question: computation is omniscient in character and transcendent in role.

Continuity is sameness of content within sameness of context: the property of the number line by which every point joins its neighbors with no gap and no jump. It is deeper than symmetry and harder to reduce, in one exact sense: it cannot be listed. No procedure lists the continuum, and Cantor's theorem states that fact. In the words of the Incommensuration Theorem, listing is the symmetric act and covering everything without a gap is the continuous act, and Cantor's diagonal forbids the two being perfected together. The number line of mathematics, as mathematicians build it, is a mixed structure: it is made of separate points, each of which can be named on its own, and those points also join one another with no gap, so comparisons of both kinds, the symmetric and the continuous, occur in it. What this argument needs is the continuity that the theorem names, the sameness of content within sameness of context just described, and not the whole of the structure that mathematicians build into the number line. The pictures of pi and of the diagonal list show exactly that continuity: the part of the number line that no list reaches. Chapter 4 matched the two kinds of law to the three modalities: symmetry goes with the omniscient, the mode of structure, and continuity goes with the transcendent; and the transcendent is the mode of the self: of the one who compares, of choice and of possibility.

With the self, I carry the argument beyond physics. Subjectivity, the inner side, the stance from which every comparison is made, is at its foundation continuity with asymmetry. Picture a river with leaves floating on it: the water runs on without a break, and no stretch of it is ever the same

twice, while the leaves go by one at a time. The self is like the water. It goes on without a break, and no two of its moments are the same; its contents, its thoughts and perceptions, are like the leaves: they change in separate steps within that unbroken going on. That subjectivity is, at its foundation, continuity with asymmetry is a claim of the framework, not a result of any experiment. It says that the first person, the inner self, is not the kind of thing that any list could capture, because listing is a symmetric act, and the subject is the side of comparison that symmetry does not reach. The same Incommensuration Theorem that forbids symmetry and continuity being perfected together therefore forbids computation and subjectivity being perfected together. A computation perfected as computation is symmetry with discontinuity applied absolutely, and what that excludes, by the theorem, is continuity applied absolutely, which is exactly what the self is.

So a computation cannot hold the continuum, and the subject is continuous. The road through continuity arrives at the conclusion that what goes on inside a simulation cannot be a separate world, and cannot contain a separate self. The road through randomness arrives at the same place. The road through continuity runs through the transcendent, through what cannot be listed. The road through randomness runs through the immanent, the mode of process, of interaction and of what enters it. That both roads arrive at one place is what the account leads me to expect, because the two roads are the two survivors of the same theorem, seen from opposite ends. The next section walks the second road in full.

A Simulated World Is a Region of the One World

The structure of comparison shows that a simulated world containing real subjects is not a separate world. It is a restricted region of the one world in which it runs.

A computation is deterministic. Its states follow one another by fixed rules: given the present state and the rule, the next state is fixed, and run twice from the same start, it gives the same sequence. Turing's description of 1936 makes this exact. The fixed sequence of states makes computation reliable and repeatable, and a fixed sequence of states is the definition of a deterministic process.

Our world has a structural arrow of time: a one-way order in which every comparison runs from the act to its record (Chapter 13). By the bound triple of Chapter 8, it therefore contains hard randomness and subjectivity. (The *bound triple* is subjectivity, directed time and hard randomness, each requiring the next, so that none can be present without the other two.) A deterministic computation has no hard randomness of its own. So if a simulated world shows randomness, that randomness is of one of two kinds.

The randomness of the simulated world may be pseudo-random, made by a fixed rule and only looking random. Then the randomness only looks random to someone who does not know the rule that made it, and it is not hard randomness. The simulated world then has an indexing of its states, a numbering of earlier and later, and no arrow; and so, by the bound triple, it has no subjectivity. Such a simulation has no stance from which it is lived. It is no simulated world in the sense the hypothesis needs: it is a program running on a machine with no one in it.

Or the randomness may be drawn from a source of hard randomness in the world that runs the computation, such as a physical process measured there. Feeding the simulation real randomness from outside is the strongest reply open to the simulation hypothesis. In that case the simulated world has no randomness of its own: its randomness is the outer world's, carried through the computation. Its arrow follows its randomness, so its arrow is the outer world's too; and the step-by-step running of the computer, one step that cannot be undone after another, is itself the outer world's arrow at work. The bound triple cannot be carried across in part. Where the randomness and the direction of time are shared, the subjectivity is shared as well.

So a simulated world with real subjects shares with the world that runs it its hard randomness, its direction of time and its subjectivity. It has no bound triple of its own, and the bound triple is what makes a part of the real a world. A *world* is where acts of comparison happen in a direction, with outcomes not fixed in advance, from a stance. A simulated world with no bound triple of its own is therefore not a second world. Such a simulated world is part of the one world: the part reached through the computation.

Picture a prison cell. The walls limit what the person inside can see, and the limit is real. The cell is still part of the one world: the person inside breathes the same air and ages by the same clock as everyone outside the walls. A computation restricts what can be reached from inside it, as a prison cell restricts what its occupant can see. The restriction is real, but the separateness is not.

So the hypothesis, in the form that matters to us, that we are real subjects in a world that is not real, is excluded by the structure of comparison, and this is a matter of principle and therefore of all future practice.

The best-known argument for taking the hypothesis seriously was given by the philosopher Nick Bostrom in 2003. It is a three-way choice. Either almost no civilization reaches the power to run such simulations, or almost none that reaches it wants to run them, or we are almost certainly simulated, since simulated observers would far outnumber real ones. The argument counts simulated observers as if each simulation were a separate world with its own observers. A simulation is not a separate world. The observers in any simulation that has observers at all are observers of the one world, and the third choice rests on a count that treats as separate worlds what are regions of the one world.

A quantum computer draws on real quantum outcomes. Could a world simulated on one have randomness of its own? The quantum outcomes it draws on are found by measurements made in the world that runs the computer, so its randomness is that world's randomness, carried through the computation. That randomness is of the second kind, ran-

domness drawn from a source of hard randomness in the world that runs the computation, and the conclusion holds: a simulated world drawing on a quantum computer shares its randomness, its arrow and its subjectivity with the world it runs in.

Is hard randomness perhaps only ignorance that some deeper theory would remove, so that the outer world's randomness is itself a recipe one level further down? Hard randomness is an outcome that no prior description fixes, however complete; call what a description leaves unfixed its remainder. The Incommensuration Theorem guarantees that every description has a remainder, since no lawful description perfects both characters, symmetry and continuity, at once. A deeper theory would be a further description of comparisons, and the theorem holds of it as of the present ones: it too would leave a remainder. So the remainder is a feature of every description that knowledge can give, and not a gap in what we know today.

The question is often put another way, as a question of evidence: could someone inside tell, by some test, whether the world is simulated? The argument of this section does not need that question answered. It turns on what a simulated world with real subjects is, and not on what can be detected from inside it. Whatever a test inside could show, the subjects who would run it share the bound triple of the one world, and so the world they find is a region of that one world.

No Outside, and the Brain That Can Be Copied

The argument of the section on the quantum wave, that a deterministic program's outputs are fixed and the universe's are not, draws on the free will theorem of Chapter 9. A further argument stands on its own, without that theorem. The simulation hypothesis asks for a description of our whole world from outside it, made by a world that contains it. A description of the whole world from a place outside it has no comparison behind it, because the whole has no outside to compare it with. The hypothesis

applies *realism* twice, realism being the claim that the world is as it is whether or not anyone engages it (Chapter 6): once to the simulated world and once to the simulating one. So it asks for a view from outside every view: the kind of view Chapter 5 met as the view from everywhere at once and found to be an idealization, a picture no comparer could ever stand in.

One more fact can look like a puzzle, and it has a plain answer. The puzzle, a *paradox of embedding*, runs like this. A brain can be simulated, and a brain is made of the same particles as everything else; so why could the whole universe, brains and all, not be simulated as a world of its own? The answer lies in the difference between how a thing behaves at one scale and the things that do the behaving. The behavior of a brain can, in principle, be simulated at the middle scale, the scale of cells and signals at which its working causes run, because behavior at that scale is described by models that are close to deterministic. The particles a brain is made of cannot be simulated as what they are. They have no earlier state that fixes their outcomes, so no deterministic program can be what they are, though a program can reproduce their statistics to any precision chosen in advance. A simulation that is a world of its own would have to be what the particles of that world are, and not only copy the behavior of some of those particles at the middle scale. So a simulation of the whole universe that is a world of its own is impossible in principle, and not only beyond today's technology.

One Level Only

A picture shows where the information of an event comes from. Suppose a very small amount of information causes a great many things to happen, and what happens holds far more information than the cause did. Then we are dealing with an information world, a simulated world, in which a command can call down to the machine underneath and delete a whole tree of files, or copy one, or make one. Compare the bits that go into the

command with the bits of change that come out. A *terabyte* is a million million bytes, and each byte is eight bits. If I delete a terabyte of data with a single short command, about twelve bits go in, and a terabyte of change comes out. (The figures belong to the example; any short command and any large store of files make the same point.) In a computational universe there is no proportion between the bits that go in and the bits that come out. The bits that come out can outnumber the bits that go in by any amount at all, because the command uses the causes of the machine underneath, the causes of the world below, as a direct way of making things happen.

In the visible universe there is one level only, the physical, and we are not dealing with a simulation: all of the information of what happens must come from somewhere within that one level. Some of it comes from earlier states, through cause and effect. Some of it comes in through the microscopic boundary, the inner edge of the domain, as hard randomness. To make one fireball in a simulation, I have to say how much heat, where, and for how long; every bit of information in the result has to be created by some procedure, and every bit that my own command does not create has to come from somewhere else. In a simulation it comes from the operating system, the program that runs the machine itself, which already held all of it and which the command puts to work: the information comes from the level beneath. A fireball in this world has no level beneath to draw on.

Set the butterfly of Lorenz's picture, whose wing flap sensitivity magnifies into a storm, beside the twelve-bit command, because both look like a very small cause with a very large effect. The flap of the butterfly's wings does not make the information of the storm. That information comes from within the one level: from the earlier states of the whole atmosphere, through cause and effect, and from what enters through the inner edge as hard randomness, which sensitivity then magnifies. The flap can decide where the storm forms, and everything the storm holds still comes from the air and from what entered it. The command is different in kind: the terabyte of change it makes comes from the machine underneath, a

level that the files of the simulation do not contain. In this world, all the information of every large effect comes from within the one level; in a simulation, the information an effect holds beyond what its command held comes from the level below.

The hard randomness that enters through the inner edge is the remainder the theorem guarantees. Chapter 7 showed that the remainder always carries meaning, and in Chapter 7 meaningful means significant to the process an outcome came from. Does that mean every event has a purpose someone intended? It means that each outcome of the remainder is significant to the process in which it occurs, and no description removes that significance: seen from that process, each outcome is one specific outcome of that process and of no other. That the remainder carries meaning does not mean that an intention stands behind every event, or that anyone could say what each event means.

What It Means for What Cannot Be

We now know that our world is not a simulation run in some other world. The arena of space and time is continuous, as the Incommensuration Theorem requires and as physics finds, and so there is something in what exists that no computation can capture. The difference between a computation and what exists is one of kind, not of engineering. A simulated world with real subjects shares the bound triple with the world that runs it, and so it is a restricted region of the one world, a prison cell and not a second world; a simulation with no one in it is a program running on a machine with no one in it.

This exclusion stands beside the others of the book. A signal faster than light asks for a comparison beyond the outer edge of comparison. A trip back in time asks one time to be both the time laid out in the shape of space and time and the time of the order of findings. A single theory that merges gravity and the quantum asks for one description to be perfect in symmetry and in continuity at once, the pairing Chapter 3 found ex-

cluded. A world that is someone's program asks for a listing to hold what no list holds, and for a second world with no bound triple of its own. Each is excluded by the structure of comparison, and each meets that structure at a different point: the outer edge of comparison, the two kinds of time, the two kinds of law, and the line between what can be listed and what no list reaches.

Step-by-step methods in physics are untouched, and physics cannot do without them: grids, time steps and digits of precision are how physicists compute, and every refinement brings a computation closer to what is found. The argument's subject is a difference in kind, not a difference in degree. A simulation can restrict what is seen. It cannot count what cannot be counted, and for that reason it cannot contain a subject.

Summary

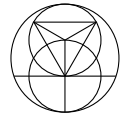
The simulation question looks like a question about computers, and it is a question about whether what exists is, at bottom, continuous or discrete. The Principle of Identity, which follows from what comparison is, says that what no conceivable test can tell apart must be counted as the same; it turns the question into a search for a test in principle. Repetition, simultaneity and the sheer load of parallel work do not decide it: the first two are copied by design, and the third is a limit of practice. The Planck length does not decide it either: it is no smallest length, and a boundary on measurement is not a graininess in what exists. Everything narrows to one choice: smooth or grainy at bottom.

Cantor's diagonal argument settles the mathematics. Everything a computation could ever produce can be listed, and no list reaches the continuum. The arena of space and time has no least step: a least step would be a difference with no sameness within it, which cannot occur; and physics finds the arena smooth at every scale yet probed. So the exclusion is structural, and for the arena of every comparison it holds without condition. Sensitivity to initial conditions is the amplifier that turns the gap

between the finite and the endless into a gap between two different histories: the history the simulation computes and the history the world runs. A deterministic program can reproduce quantum statistics to any precision chosen in advance; for any one fixed program, beyond its smallest step, a sensitive enough measurement eventually tells the program's statistics from the world's. That the smooth wavefunction is evidence that what physics finds takes part in the continuum is a reading at the level of interpretation, of known results. That subjectivity is, at its foundation, continuity with asymmetry, and so cannot be listed, is a claim of the framework, not a result of any experiment.

What can be listed and what no list reaches are symmetry and continuity, seen in the mirror of mathematics. Computation is symmetry perfected; the continuum is what symmetry perfected leaves out; and the self, on this account, lives on the side that is left out. By the second road, the road through randomness, a simulated world with real subjects shares the bound triple with the world that runs it, and so it is a restricted region of the one world; a simulation with no one in it is a program running on a machine with no one in it. The hypothesis, in the form that matters to us, that we are real subjects in a world that is not real, is excluded by the structure of comparison, and this is a matter of principle and therefore of all future practice. In the visible universe there is one level only. The road through continuity is built from three things alone: the Incommensuration Theorem, the matching of the two kinds of law to the three modalities that Chapter 4 gave, and the steps this chapter has walked.

If our world is not a program run in some other world, could the world itself, or the mind, still be a computation in some deeper sense?



Chapter 22

Why No Computation Is Everything

A Reply With No One Behind It

Isn't everything, at bottom, information being processed, minds included? A person types a question into a chat window and receives a thoughtful, fluent reply. Nothing in the reply shows whether anyone is there. The question of this chapter is whether anyone could be, and whether the world itself, minds and all, is at bottom a program running.

Chapter 21 found that our world is not a simulation run in some other world. That chapter argued from the continuum, the unbroken whole that can be divided without end. A computation, a procedure that passes from one state to the next by a fixed rule, in separate steps, is a listing, and reality holds what no listing can reach, so no listing can capture reality. In this chapter I take up the same question, whether the world can be a computation, and I start from the claim that the world is a computation. I ask what kind of claim that starting claim is, and what a computation must leave out.

The claim has a name. *Computationalism* is the thesis that all of physical reality is, in principle, computational: that every process can be represented as *discrete* operations on discrete states, discrete meaning separate and countable, one step and then the next, with nothing between. In its strong form it claims that the universe is a computation, and that what exists is, at bottom, information processing. Applied to the mind, the thesis is

called *strong computationalism*: the right *algorithm* (a step-by-step recipe that a computer follows), running on any machinery at all, would thereby be conscious, with no further ingredient required. The thesis underlies most of what is said today about uploading minds, simulating everything, and the computational nature of consciousness, and it is rarely stated as the claim about reality that it is.

The implications are sweeping, and they are meant to be. If the universe is a computation, then any sufficiently faithful simulation of a physical system is, for every purpose that matters, the same as the original. A simulated brain would be a real mind. *Mind uploading*, the proposal that a person can be captured whole in a computation (scan the brain, run the scan, and the running scan is the person), would be natural, a straightforward consequence of what reality already is. Digital immortality would follow, and so would the thought that persons are, at bottom, information.

The thesis has enormous cultural momentum. It is the assumption behind much of what is said about machine intelligence, about virtual worlds, about the simulation hypothesis, and about the future of the person. It is almost never argued for. It is treated as obvious, as the default position that a scientific mind should adopt unless given some reason not to. So I state it as what it is: a *metaphysical* commitment, a claim about what reality is, and one that the account built in this book rejects at the root.

The argument needs nothing beyond what earlier chapters have shown. It uses the Incommensuration Theorem of Chapter 3, which is a structural fact about the nature of knowing: symmetry and continuity cannot both be applied, at once and absolutely, to one comparison, and neither can asymmetry and discontinuity. It also uses the three modalities of Chapter 4 (the immanent, process; the omniscient, structure; the transcendent, possibility); the exclusion of determinism in Chapter 5; the hard randomness of Chapter 8; and the continuum argument of Chapter 21.

What Kind of Claim It Is

A scientific claim, of the kind a law within one domain is, must be *falsifiable*: there must be some possible observation that, if it were obtained, would count against it. The claim that every swan is white is falsifiable: one black swan would count against it, and when Europeans reached Australia they found black swans. Falsifiability is a necessary part of what makes a claim scientific, though not the whole of it, and it gives a clean first question to ask. What observation would show computationalism false? None.

Every apparent counterexample is absorbed as a promissory note, a promise to pay later. Where a phenomenon looks as though no computation could produce it, the computationalist can answer that the computation is more complex than we have yet understood. Where a process appears to involve randomness that nothing fixed, the computationalist can posit a deeper layer beneath it, in which everything is fixed. Where consciousness resists every attempt to simulate it, the computationalist can lay the resistance to a simulation not yet detailed enough, to be remedied by the next round of hardware. There is no state of affairs that the thesis would accept as its refutation.

Since no state of affairs would count, for the computationalist, as a refutation of computationalism, computationalism is not a scientific conclusion. It is a metaphysical commitment that presents itself as an empirical finding, a finding of observation and measurement.

This charge is about the kind of claim computationalism is. Computationalism is *non-falsifiable*: no possible observation would count against it. Being non-falsifiable is a different thing from being false. A non-falsifiable claim stands in the same category, as to what can be known of it, as every other commitment that no observation could count against, religious commitments included. It may be believed. It cannot be established by science, and it is not entitled to the authority of science.

Chapter 20 drew the line that applies here: falsifiability applies to claims of how and why, and not to claims of what, and treating a metaphysics (a claim of what there is) as falsifiable is a category error, the mistake of applying a concept where it does not belong. Chapter 20's line also limits the case against computationalism. Computationalism is a metaphysics, and a metaphysics is not refuted by a single observation, as a law is. It answers to argument, and its consequences, where they reach into physics, answer to experiment. The rest of this chapter gives the argument, and its conclusion is that computationalism is, as a matter of structure, false.

What a Computation Is

Picture the simplest computer there is, the one Alan Turing described in 1936. A long strip of paper, the *tape*, is divided into squares, and each square holds one symbol. A reading head sits over one square at a time. The machine is always in one of a fixed list of conditions, called its *states*. And it has a rule, called its *transition function*, which is a table: for each state and each symbol under the head, the table says which symbol to write, which way to move along the tape, and which state to go into next. The tape, the reading head, the states and the transition function are the whole machine. It is called a *Turing machine*, and it is the standard model of what computation is.

Three features of computation carry the whole of the argument, and none of them is in dispute.

The first feature is that computation is deterministic. Chapter 5 gave *determinism* its full meaning: everything fully and exactly fixed, every value of every content for every context. A computation is exactly that, within its own states. Given the current state and the symbol under the head, the transition function fixes one next state, one symbol to write and one direction to move. There is no ambiguity and nothing left open. Software of every kind is states, conditions (tests of the form "if this, then

that") and loops, and every step is fixed by the step before. Because every step is fixed by the step before, computation is reliable, checkable and repeatable, and every computation can be written as a fixed sequence of changes of state. A *classical computation*, one carried out by ordinary digital machinery in definite steps, is determinism made exact within its own separate states: a chain of states, each fixed by one rule from the last. One rule, applied the same way at every step, is symmetry perfected: it is the tiled floor of Chapter 3, the same tile laid in every place.

Now picture a program whose job is to produce random numbers. It starts from one number, called the *seed*, and a short recipe turns the seed into the next number, that number into the next, and so on. The numbers that come out pass every statistical test of randomness. Yet run the program again from the same seed, and exactly the same numbers come out, in the same order. Such a program is called a *pseudo-random generator*: its output looks random, and it is fully fixed by its seed. No computation produces hard randomness, the outcome of Chapter 8 that no prior specification fixes, however complete, for any comparer whatsoever. A computation can only imitate it.

The obvious reply is that a computer can read a quantum device that produces random numbers. It can, and the randomness then enters the computation as input from outside it: it is new information entering the world at the inner edge of the domain of comparison, the edge marked by Planck's constant, as Chapter 18 described. The computation itself remains deterministic; the randomness is in the physics, and not in the algorithm. The randomness such a system has is the world's randomness, carried through the computation. Its direction of time follows its randomness, and so its direction is the world's direction too. And the bound triple of Chapter 8 (hard randomness, directed time and subjectivity, each requiring the next) passes from the world into a computation whole or not at all. Where a computation shares the world's randomness and direction of time, it shares the world's subjectivity as well.

A quantum computer is the same case. It is a physical device whose operations, between the moment it is set up and the final readout, follow a *reversible* rule: every step can be undone exactly, so the sequence runs as well backward as forward, which is symmetry with no arrow. Its readout is a measurement, a physical comparison made in the world, with a single undetermined outcome and a record that cannot be undone. So a quantum computer has what every measuring instrument has of the bound triple: it has it at its measurements, as a physical system in the one world, and never by virtue of the program it runs. Its program, as a program, is a symmetric rule like any other.

The second feature is that computation has no direction of its own. Charles Bennett showed in 1973 that every computation can be carried out reversibly, by keeping a record of each step, and that the reversed computation is itself a computation. Where an actual computer does have a direction, it is because the computer erases information. Erasing is a *thermodynamic* act of the physical machine, an act that belongs to heat and its flow (Rolf Landauer showed in 1961 that erasing information gives off heat), and it is no feature of the computation. The direction comes from the machine and its clock, never from the algorithm. Whatever direction a computation appears to have is imposed from outside: by the hardware clock, by the intention of the programmer, by the order in which someone chose to read the tape.

A computation's steps come in order, step one, step two, step three, and an order is a different thing from a direction. The whole list of a computation's states is contained in its start and its rule, and the list can be read either way without any difference in its structure, like the page numbers of the finished novel of Chapter 5. What a computation has is an *indexing* of states, a numbering (first, second, third). An indexing is the kind of time Chapter 14 called geometric time: time laid out all at once, with every moment already in its place, as the times of events are laid out in the shape of space and time, and not event time, the one-way order of findings for a comparer. Following the word of Alfred Korzybski, I call the binding of past to present to future in one unbroken thread *time-binding*,

in a narrower sense than his (for him the word meant the passing of knowledge across generations). Time-binding is what the arrow of time of Chapter 8 does. A computation is sequential without being time-binding: its steps come in order, and the order adds nothing the rule did not already fix, so a computation has no arrow of time of its own. The clock is a convenience and no part of what a computation is, just as Chapter 5 found the clock to be a convenience and no ingredient of determinism.

The third feature is that computation is without *interiority*: it has no inside. A Turing machine has no perspective. There is nothing it is like to be a transition function. The description of a computation is complete from the outside, in the *third person*, as one describes someone else, and no report in the *first person*, from inside, adds anything to that description, because there is no first person to report. That holds of every computation, on present machines and on future ones, because it is what the concept of computation is: the fully specified, fully shareable, fully repeatable transformation of symbols. Every state of a computation can be handed to another and read by another.

In the three modalities of Chapter 4, the three features say one thing about the character of computation, what kind of thing it is. Computation, in character, is purely structural, purely *omniscient*: the repeatable, the third-person, existence and pattern and law, with nothing else admitted. In character, the immanent, the here and now of interaction and change, and the transcendent, what is weighed and what is possible, are absent from computation by its definition: they are what computation is defined by excluding. (Its role, the part it plays, is a different question, and it comes in a moment.) Chapter 21 found the same character in the terms of the four compounds of Chapter 2, the four pairings of symmetry or asymmetry with continuity or discontinuity. In character, computation is symmetry with discontinuity, and a theory written as a computation is laid out wholly as structure, purely omniscient.

Character is one question, and role is another. The character of a thing says which of the four compounds it shows; its role says which part of a comparison it plays. By role, set beside physics and logic in the practice of

inquiry, computation is transcendent: a procedure is defined over the range of all its possible inputs, and each run selects one path through that range, turning a possibility into something actual. Both answers hold, and they hold together, because they answer two different questions. The last section of this chapter sets the three roles side by side.

One more feature of computation says what it is good for. Any form of modeling depends on symmetry. To model a thing is to treat some of its aspects as irrelevant to the behavior of interest, and to factor them out as unimportant. Those irrelevant aspects are differences of context, changes of setting under which what one cares about stays the same, and a sameness of content across a difference of context is what Chapter 2 called a symmetry. Symmetry, a sameness of content across a difference of context, is the basis of what it means to be classical. A classical system is one that can be virtualized, that is, copied into a model, and that is the single defining feature of being classical: classical things are modelable. If a model of a thing can be built that behaves identically to the original for every purpose one cares about, then that thing is classical.

What a computation perfects is determinism, and causality is a different thing. Chapter 6 showed that a cause is a consistency that a finite comparer finds between two things in the world, one following the other, again and again. Causality rests on continuity: causality requires a self, a time, and a perspective from which to notice that one thing follows another. It is found along an order of findings, and its pairing, continuity with asymmetry, is one the theorem permits. Remove the observer from causality, and causality is not perfected into determinism: causality is removed.

The Bound Triple, Again

Someone listens to a piece of music they have never heard. Each note arrives after the last, and the first hearing cannot be had again: there is a before and an after. Some notes surprise: nothing the listener heard before fixed them. And the hearing is someone's: the music happens to a listener. Those three things appear together in every first-person point of view.

Chapter 8 named them. The first of the three things is the *arrow of time*: an order that runs one way, binding past to present to future in a single thread without a break. The second of the three things is hard randomness: an outcome that no prior specification fixes, however complete. The third of the three things is the subjective: the side from which experience occurs, the stance of the self to whom things happen. Together they are the bound triple, three features of reality that are inseparable: remove any one, and the other two go with it. I derive the binding from the theorem, and then show physics bearing it out.

The binding follows from the Incommensuration Theorem. Because symmetry and continuity cannot both be perfected, every comparison leaves a remainder, an outcome not fully determined by prior structure: that remainder is hard randomness. Every comparison is made from a stance, the subjective pole of comparison, and every comparison proceeds in a direction, from the act to its record: that direction is the arrow. The three are three aspects of one act: comparison proceeding in a direction, registered by a perspective, with an outcome not fully determined. None of the three is present without the other two.

Physics bears the binding out, in the two interpretations of quantum theory that stand at opposite ends. (An interpretation, as Chapter 8 put it, is an account of what the mathematics of quantum theory says the world is like.) The Copenhagen interpretation keeps the observer, the random single outcome, and the irreversibility of measurement. The many-worlds interpretation, in which every outcome is kept in a branching state, drops all three at once: it has no fact of which branch a comparer is in, no outcome that is not fixed, and a basic equation that runs equally well forward

and backward in time. The interpretations differ on which side they take; none separates the three. They keep the three together, or they drop them together. So the binding is established twice: derived from the theorem, and found in physics, where no account of quantum phenomena has ever kept one of the three without the others.

Each member of the triple has its full statement. Hard randomness is real indeterminacy in what happens, an openness in what happens next that nothing reduces, as distinct from the unpredictability of ignorance, which better knowledge of the starting conditions could resolve. Quantum measurement is the standard case, and Chapters 8 and 10 gave the reasons, from Bell's theorem and the experiments that closed its loopholes, for taking hard randomness to be a feature of nature and not a screen over a deeper order. No deterministic account lies under it, even in principle, and no computation can produce it.

The arrow of time is the structural fact that the past is settled and the future is open. Chapter 13 showed that this arrow comes before the thermodynamic one, the arrow of heat flowing and eggs breaking. Without hard randomness there is no openness, and without openness there is no arrow, only an ordering that could be read either way.

Subjectivity is the subjective pole of comparison, the stance from which every comparison is made. In a conscious being it is the fact that there is something it is like to be that being, in the words Thomas Nagel made famous in 1974. The *hard problem of consciousness* asks why there is experience at all, and not information processing with no one to whom it happens. At bottom the hard problem is the observation that no structure handed to another person carries one's own knowing. A complete description of pain is knowledge laid out as structure: which nerves fire, in what order, with what chemistry, and what the person in pain says and does. What pain feels like is reached only in one's own perceiving. The *explanatory gap*, the gap between a physical description and a conscious experience, is a structural feature of comparison itself. Subjectivity is the in-

ner face of the same one-way difference that the arrow names: the perspective from which what has been settled is told apart from what has not.

The binding shows most clearly through content and context, the distinction on which the whole account of comparison rests (content is what a comparison finds; context is the setting that makes one finding distinct from another). Hard randomness is novelty entering: content arriving that no prior context contained. The arrow of time is the difference between what has been settled and what has not, the difference between the context that is closed and the content that is still open. Subjectivity is the inner perspective on that difference, the point from which the open is seen as open. Hard randomness, the arrow and subjectivity are three faces of one solid, and there is no having a face without the solid, as there is no having one face of a die without the die.

To see the bound triple taken whole, picture the solid as a die held in the hand, and look at each face in turn, within the frame of the triple. Each face has its own modality. The arrow of time is immanent: the process, happening. Hard randomness is omniscient: described from outside, it appears as a gap in the structure of law. The subjective is transcendent: the side of the one engaged, and of what is possible. The three faces are distinct, since none is either of the others; inseparable, since none occurs without the other two; and not interchangeable, since none can do the work of another. As my book *Incommensuration; The Theorem and the Three Axioms* says, the theorem binds together three things that are usually treated as separate puzzles: subjectivity, directed time and hard randomness. Now take the whole solid at once, in a context, as one act happening, as the die is one thing however it is turned. Taken that way the bound triple is immanent, and I identify it directly as *consciousness*. Hard randomness, which seems omniscient when described from outside, is then immanent too: it is the remainder in the act, as the act happens.

Take away one face at a time. Remove hard randomness, and every event is fixed by the states before it: the future is as settled as the past, the difference between past and future becomes one of perspective and not of

being, and subjectivity, the capacity for real choice, has nothing left to choose. Remove the arrow, and cause and effect become interchangeable: a random event, being one not fixed by its causal history, has no causal history to be free of, and experience, which happens in time by its nature, has no medium to happen in. Remove subjectivity, and the difference between a choice and a random event, which is invisible from outside in any case, loses the one vantage from which it had content. And the arrow of time, with no inside to be an arrow for, is a statistical artifact, a pattern that shows up only in averages over many events. Each requires the other two.

Each also requires the next, and the three close a circle. Directed time requires hard randomness: if every state were fixed by the states before it, the future would be contained in the present as fully as the past is, and the order of states would be an indexing, the same read forward or back, and not a direction of process. Hard randomness requires subjectivity: a purely third-person description cannot tell an outcome that no prior state fixed from one that was fixed and is unknown, since both are unpredicted. The difference is between a finding that is new and a finding that was only hidden, and a finding is made by a comparer, from a stance. Subjectivity requires directed time: a stance sets what is found against what was found; remembering sets the present against a record, and anticipating sets it against what could come. Each of these runs one way, and with no direction of time there would be no difference between remembering and anticipating, and no side from which to look.

What Computation Takes Away

The argument can now be stated in one motion. Computation is deterministic, has no direction of its own (it runs the same backward as forward), and has no inside. The bound triple is hard randomness, directed time and subjectivity. Computation removes all three, and it removes them as a single act and not one at a time, because the three are one structure and

computation is defined as the absence of that structure. Therefore computation cannot replicate any system in which the three are present, which is to say any system that is real, in the sense Chapter 4 gave the word: in an actual act of comparison, here and now.

Why a Computation Cannot Be What Consciousness Is

A mosaic can follow a curve as closely as anyone likes. Make the tiles smaller, and the steps along the edge shrink, until no eye can see them. The mosaic is still made of tiles, and the curve is still not made of anything. A classical computation is a mosaic in time: one rule applied the same way at every step, which is symmetry perfected.

By the first clause of the Incommensuration Theorem, what is absolutely symmetric is intrinsically discontinuous: it comes in steps, with seams between them. A continuity cannot be assembled from such steps, any more than a symmetry can be drawn out of a pure continuity given alone. The arrow of time, in the first person, is continuity with asymmetry: an unbroken order that runs one way. So a classical computation lacks the arrow of the bound triple, as a direct consequence of the theorem. It also lacks hard randomness, since it is deterministic. And because the three are bound, lacking two it cannot have the third: it lacks the subjective.

The conclusion follows: classical computation, however intricately organized, cannot be what consciousness is. It may be capable of much else, including behavior that looks, from a third-person vantage, indistinguishable from conscious behavior. The argument is about what a computation is, and so it holds whatever the computation does.

The result that classical computation cannot be what consciousness is has a general form. That general form is the principle behind the cross-modal reach of Chapter 17, by which each of the two great theories reaches, with its tool, into a domain of the other modality. A framework and the content it describes are always of *conjugate* kinds, opposite kinds,

each of which can be perfected only at the cost of the other: a symmetric framework probes continuous content, and a continuous framework probes symmetric content. The reach of the framework into its content is dense, as close as one likes, and never exhaustive. A computation is a symmetric framework. It can approach the continuous as closely as one likes, and never be it.

Symmetry and continuity are conjugate in just this way, in the word of Chapter 3, and neither is prior to the other. Neither is more basic than the other, and both are built from comparison, which is the immanent act. Continuity is presupposed by every application of symmetry, as a precondition, and being presupposed does not rank continuity above symmetry. So the first-person perspective is continuously asymmetric, a computation is discontinuously symmetric, and neither can be composed of the other.

One Event, Two Descriptions

One point about the subjective and the random belongs here. From the outside, a conscious choice is indistinguishable from a random event. Chapter 7 drew the picture. If I raise my left hand and not my right, and the choice is a real one, not fixed by prior states, then from any vantage outside me it looks exactly like randomness. No outside measurement can tell the two apart, and the reason is that they are the same event, described from two positions.

Described from the omniscient position, the position of structural law and of description in the third person, the undetermined outcome is hard randomness: a gap in the specification, without pattern and without meaning. Described from the transcendent position, the position of the one engaged, the same outcome is choice, meaningful to the process from which it came. Neither description can replace the other. From outside, randomness; from inside, choice: one event, two descriptions, both correct, neither complete. From outside, the two are *isomorphic* in a precise sense: they match point for point.

The account of this book makes a stronger claim than this, and I mark it as stronger. It holds that the subjective and total entanglement are strict isomorphs, exact matches of one structure; in plain words, subjectivity is the inner face of a process which, seen from outside, presents as entanglement and hard randomness. (Two particles are *entangled*, in the sense of Chapter 10, when their outcomes, measured far apart, agree in a pattern that no separate instructions could give them.) On this account the brain operates near *criticality*, a poise between order and disorder in which small disturbances cascade, as the study of avalanches of firing among nerve cells by John Beggs and Dietmar Plenz in 2003 suggests. In that regime, indeterminacy at the microscopic scale enters chains of cause that reach the everyday, macroscopic scale, and bears on choice. The behavior of a brain at the middle scale, the *mesoscopic* scale of cells and their circuits, can be modeled by computation that is approximately deterministic, and the model is useful. The single choice, near criticality, carries indeterminacy from the microscopic scale that no such model fixes. Subjectivity is part of what the act of comparison is, and so it is no *epiphenomenon*, no idle by-product of processes that would run the same without it.

So the brain is a physical system making comparisons in the world, with the bound triple at every act of perception and choice. A computer model of the brain is a map, and the brain is the land, in the picture of Chapter 1.

The claim of this argument is about how the inner and the outer perspectives are related, and it holds for any system in which novelty enters, time is directed, and there is a point of view. It rests on no particular biological mechanism in the brain, and it needs no sustained entanglement there; the Appendix sets it beside a well-known hypothesis that does rest on such a mechanism. The brain's criticality is the physical circumstance under which the relation of the inner and the outer perspectives is expressed in us, and the brain's criticality is not that relation itself.

What Computation Can Do

Computation is enormously powerful, and this argument takes none of its power away. Machines built on computation can be made to do a great many things that once required a person, and they will be made to do more. Computation can calculate, predict, translate and converse, and do all of it superbly. Artificial intelligence is possible. The boundary the argument identifies lies between processing and experience, and not between the stupid and the intelligent. Intelligence, as competence at a task, is an omniscient property: structural, repeatable, shareable, and assessable from outside. Experience is not. The argument leaves the first entirely open, and closes only the second. What a computation cannot be, by being a computation, is a first-person point of view.

The programs called artificial intelligence today run on classical machinery. Where they draw on chance, as when a *language model* (the kind of program that writes the reply in the chat window) picks its next word, the chance usually comes from a pseudo-random generator, fixed by its seed. Such a program has no hard randomness of its own, and with it, by the binding, no arrow and no subjective side. Pseudo-randomness gives an indexing of states, and no arrow. A program can instead draw its chance from a physical source, such as electrical noise or a quantum device. Then, as with the quantum random-number device earlier in this chapter, its randomness is the world's and not the program's: the program is still one rule applied the same way at every step, and what hard randomness the system has, it has as a physical process in the one world, at the source.

In every case the result is the same. Whether a system has the bound triple depends on whether it is a physical process making comparisons in the world, with an inside; it never depends on which program it runs. Could some future machine of a different kind have the bound triple? Whatever has it, has it as a physical process making comparisons in the

world, with an inside, as living things do, and the program it runs never supplies it. The question for any future device is therefore what kind of physical process it is, and not which program it runs.

If a program's behavior were indistinguishable from a person's, how could anyone tell the difference? The argument concerns what a computation is (one rule applied the same way at every step), and that is known from how the computation is built, not from how it behaves, so it never asks anyone to tell by watching. A third-person description of behavior is the one description that cannot show an inside, for a person or for a program. The fluent reply in the chat window is behavior. Whether anyone could be behind it is settled by what is behind it: a program, as a program, has no one in it.

A Person Is Not a File

Because the bound triple holds, mind uploading is a category error and not a difficult engineering problem, and the difference between those two things is the central claim of this chapter.

An *engineering problem* is one in which the goal is well formed and the obstacle is resources: speed, memory, resolution, fidelity, time. The history of computing is a history of engineering problems of this kind being solved, and it is natural to suppose that the simulation of a mind is another, harder than the rest but of the same kind. The supposition is that a person is, at bottom, information: that a sufficiently complete scan would capture everything relevant, and that a sufficiently faithful running of the scan would be the person.

A *category error* is a different thing. A category error is the mistake of applying a concept from one domain to another domain where the concept does not belong, so that the request is malformed, and not merely hard. To upload a mind is to require that every relevant feature of a conscious system be captured in a computation. Experience has three conditions it cannot do without: something new arriving that nothing before it

fixed, time running one way, and a point of view from inside. Those three conditions are exactly the features that computation cannot represent. They are not difficult for computation; they are structural absences in it. Computation is what you have when those three have been factored out, and no better hardware supplies them. Asking a computation to contain them is asking for a description that is at once complete from the outside and possessed of an inside, at once fully determined and open, at once reversible and directed.

The Incommensuration Theorem states the same point about computation and experience in the theorem's own terms, as the picture of the mosaic showed. Computation, perfected as computation, is symmetry with discontinuity applied absolutely. The first-person stance is continuity with asymmetry: an unbroken engagement that runs one way and sets this moment apart from every other. A continuity cannot be composed of purely symmetric steps, just as a curve cannot be composed of tiles, however small the tiles are made: the tiles laid along the poured ramp of Chapter 3 always leave seams. The theorem that forbids perfecting symmetry together with continuity therefore forbids perfecting computation together with subjectivity. A person is not a very large file. A person is what a file, by its nature, is not.

Behavior is exterior, the exterior is omniscient (laid out as structure, in the third person), and an exterior description of the exterior is always available. Such a description reproduces behavior, and it adds no knowledge of an inside. So the argument does not say that a computational system cannot be built to behave as a person behaves, to any degree of behavioral faithfulness one likes. The argument says that such a system would not be that person, or any person, because being a person is not exhausted by any exterior description.

The dream of *digital immortality*, of living on as a copy in a machine, rests on the assumption that persons are information. The bound triple says that persons are, at bottom, real novelty, irreversible time and inner experience, held together as one, which is what information, by itself, cannot be.

The Program With No One in It

Chapter 21 answered whether our world is a simulation with the same argument from the bound triple that this chapter uses. A brief reminder of that answer shows the argument applied to a world, where this chapter applies it to a person. The *simulation hypothesis* proposes that our world is a computation run in a more basic world. A deterministic computation contains no hard randomness of its own, so a simulated world draws its randomness from one of two places. If the randomness of a simulated world is pseudo-random, the simulated world has an indexing of states and no arrow, and so, by the bound triple, no subjectivity: the simulation is a program running on a machine with no one in it. If its randomness is drawn from a physical source in the world that runs the computation, then its randomness is that world's randomness, and its arrow follows its randomness. It shares that world's subjectivity as well, since the bound triple passes from one world into another whole or not at all. A simulated world containing real subjects is therefore not a separate world: it is a restricted region of the one world in which it runs.

The argument for the hypothesis that Nick Bostrom gave in 2003 turns on a count of observers (Chapter 21 set it out): simulated observers, it says, would far outnumber real ones. It counts each simulated world as a separate world with its own observers, and that is the step the bound triple corrects. The argument against a simulated world and the argument against an uploaded mind are one structure, seen from its two ends. There, the world could not be put inside a computation; here, the person cannot be.

The Wolfram Physics Project: An Impression

The computationalist thesis has a concrete instance in its most ambitious scientific form, and this section and the four after it take it up. Three of the things they state are results, each with its own reach: the information

bound and the incompressible residue cover every theory laid out wholly as structure, the Wolfram Physics Project included; the three routes cover every research program that takes the world, at its foundation, to be a computation; and the order of depth is a claim of the metaphysics itself. What I say about the Wolfram Physics Project as a project, I offer as my overall impression rather than as a verdict. The Wolfram Physics Project, which Stephen Wolfram set out in 2020, proposes to obtain all observable phenomena from simple computational rules applied to starting conditions. Overall, it seems to me to be *purely omniscient theory*: theory laid out wholly as structure, admitting only what is structural, repeatable and stated in the third person. As such, it is unable by its nature to ask or to answer what are probably the most interesting and deepest questions.

The project can be stated with four names. A *rule* is a computational recipe for rewriting a pattern, step after step. A *starting condition* is the pattern the rule begins from. A *computed result* is the complicated pattern the rule produces from the starting condition. An *observation* is any complex thing we find in the world. For the kind of case the project rests on, take a row of white squares with one black square in the middle. The rule says how to color each square of the next row, below it, from the colors of the three squares just above: the one directly above and its two neighbors. Run the rule row after row, and from that single black square a large pattern of triangles grows, complicated and irregular, that looks much like the markings on the shells of some sea snails. A computed result that looks like something observed in this way is what I call one of the project's *flowers*. We have many cases of the form "this rule, run on this starting condition, leads to this complicated computed result". From them one is tempted to infer that every complex observation could itself be the outcome of some special choice of rule and starting condition, and that the task of science is merely to find which rule and which starting condition produce a computed result resembling each observation. That inference, generalized, is what makes a theory of everything. As near as I can tell, this is the research program, and the rest fits into it.

The project can be examined in terms of information. Does the computed result carry any more information than the rule and the starting condition together? Here the information in a thing means its *algorithmic information content*: the length of the shortest description of it, the measure that Andrei Kolmogorov set out in 1965, and that Ray Solomonoff and Gregory Chaitin reached on their own in the same years. The number of steps the rule is run is counted as part of the starting condition. Insofar as the computed result is the deterministic outcome of the rule and the starting condition alone, with nothing from outside added, the information in the computed result cannot exceed the information in the rule and the starting condition taken together, give or take a fixed allowance set by the choice of language in which the descriptions are written. This principle, that deduction never adds information, is called the *conservation of information* under deduction, and the principle is inherent in the idea of mathematics itself. Were it otherwise, a valid inference would produce from its premises something its premises did not contain, and the notion of a valid inference would be inconsistent. So, insofar as the notion of a theory is well defined, the output carries no more information than its sources together. The appendix For the Physicist writes this bound as a formula, with the same conditions.

The information bound is what makes theory useful. To *compress* information is to say the same thing in a shorter description. Where an observation is close enough to some computed result, an observation that seems to hold a great deal of information is compressed into the much simpler rule and starting condition, and complex, unknowable things are replaced by simpler, knowable ones. Any rule that is the shared basis for many such computed results is then called a *law of nature*, the best available compression of everything observed. A *theory of everything* is the claim that a single rule, with varying starting conditions, compresses every observation there is.

Three questions follow, and I do not think the project can answer them. First: how closely must a computed result match an observation before we may say that the observed thing is itself produced by some coun-

terpart of that rule and starting condition, and does not merely resemble their output? Second: even where a computed result resembles an observation, on what basis do we hold that, among the endless available choices of rule and starting condition, the resemblance is anything other than selected after the fact? Third: even if some small set of rules accounts reasonably well for a wide range of observations, where does the information in those rules come from?

The Residue No Theory Compresses

The third question is the deep one, and it is the one that any purely omniscient theory must leave unanswered. Every such theory carries some class of information that is wholly incompressible: a *just-so story* (a tale that explains a thing by saying it was so, taken on trust and derived from nothing), a belief, a starting point that the theory takes and does not give. The theory can move that information from the observations into the rule, and it cannot produce that information from nothing, since deduction conserves information. However many rounds of theory building are attempted, and each later round grows more arbitrary as the unmatched observations dwindle, some information is always left that no rule and no starting condition of any theory fixes. An outcome that no prior specification fixes, however complete, is exactly what Chapter 8 called hard randomness. So what remains is an *incompressible residue* of hard randomness: something not predictable on the basis of any theory. Compression of information can never be perfect; it can never reduce all observable information to zero bits, a *bit* being the smallest amount of information, the answer to one yes-or-no question.

What every actual world has is causation, together with a residual that no theory compresses, and that residual is hard randomness. On either side of causation stands a pole, and the Incommensuration Theorem forbids both poles. One is perfected determinism, which cannot be supported as a position. It can be stated only against its absolute opposite, perfec-

ted nondeterminism, which is pure randomness (in Chapter 6's words, a content found different whatever its context), an abstraction just as determinism is. Determinism would perfect symmetry with continuity, and pure randomness would perfect asymmetry with discontinuity. Chapter 5 gave the structural reason in full: determinism requires perfected symmetry and perfected continuity together, and the theorem forbids the two together. The information argument is the same exclusion, seen from the side of theory.

The third question asked where the information in the rules comes from, and inside theory it has no answer. At some level, theory must come down to practice, and practice is different from theory by its nature, even where, by theory, the two should be the same. At best, the third question remains unanswered, and is recognized as unanswerable.

Three Routes to One Conclusion

Two structural grounds bear on any research program that takes the world, at its foundation, to be a computation. The first of the two structural grounds is determinism. A purely computational universe would be deterministic, and determinism needs the same law carrying every state into its successor everywhere and at every time, which is symmetry perfected, and every state specified at every point without a gap, which is continuity perfected. The Incommensuration Theorem forbids the two together.

The second structural ground is the relation between a description and what the description describes. Every exact theory has two levels, its *formalism* (the mathematics it is written in) and its findings, and the two levels take the two surviving characters of Chapter 3 one each: continuity with asymmetry, and symmetry with discontinuity. A world computed at its foundation, a set of discrete elements rewritten by one rule applied the same way everywhere, would have one character at both levels. One character at both levels is a character compared with itself, and a character

compared with itself finds nothing it did not already hold. Such a world would find only what its rule already held, and the world physics finds holds outcomes that nothing before them fixed, as Chapter 8 showed.

Picture how the weather is forecast. The forecasters divide the air above the land into a grid of boxes, give each box its temperature, pressure, wind and moisture, and have a computer step every box forward in time from the state of its neighbors. The grid follows the smooth flow of the air closely enough to say whether it will rain tomorrow. A computation can approximate smooth physics in this way; it is then a tool used to describe the world, and not the world's foundation.

The two grounds stand beside the bound triple, and all three conclude that the world is not, at its foundation, a computation. Their premises are the Incommensuration Theorem, the two levels of every exact description, and the binding of hard randomness, the arrow of time and subjectivity. Three routes on different premises reach one conclusion, and they are independent routes that confirm one another. None of them depends on what experiments can presently detect. Each turns on what a computation is.

The Order of Depth

Beneath the three routes stands an *order of depth*, from the deepest to the least deep: metaphysics, physics, computation, mathematics. Depth here is the order of presupposition, of what each must already have in place: one discipline is deeper than another when the second cannot be carried out, stated or checked without the first, and the first can be carried out without the second.

Physics is deeper than computation, because every computation that is carried out is a physical process, in circuits, in nerve cells or in ink, and physics needs no computation to happen. Computation without physics would be procedure with nothing to run on. A rule that is never run computes nothing; running it is an interaction in the world. Computation, in

the same sense of depth, is more basic than mathematics, because every proof is checked by a computation, and no computation waits on a proof. And metaphysics stands first, because every finding of physics is an act of comparison made from inside the world, and the act of comparison presupposes none of the other three. My book *Physics Is Not Math* argues each step of this order in full.

The order bears directly on the question of this chapter. A world computed at its foundation would put computation beneath physics, and the order of depth puts physics beneath computation: every computation is something the physical world does, run on circuits, nerve cells or ink, and so a computation is never the ground the physical world stands on.

The Flowers

My impression of the Wolfram Physics Project as a whole rests on a few observations, and the last of them leads back to the bound triple. Personally, I think the drive to have things be nice and neat distorts the choice of the single rule to too large a degree. Social and psychological factors matter more to the definition of science than is usually accounted for, and the negotiation of which rule is the rule can be more a matter of position than of fact. I am also skeptical that even very good matches between computed results and observations show that we understand what we observe. There will be other observations that do not match so well; it is not always clear that different observations belong to one class; and the range of ways an observed thing can vary may depend sensitively on factors wholly unknown. The way an observed thing changes over time is often highly *nonlinear*, its changes out of all proportion to their causes, and microscopic disturbances can bring new information into the system everywhere at once. Hence determinism, as distinct from causation, is a poor basis for modeling the larger part of actual process.

Taken as a whole, then, the project seems to me oversold, and some displaced disappointment lies in its future. The explanatory power of the project will forever seem close at hand, always suggestive and finally unsatisfying. The flowers, the computed patterns that look like some observation, as the seashell pattern did, are distracting, especially since the question of how anything is known at all is left unresolved. That question is the one that matters for this chapter. Given that the topic is physics, do we know things by experiment and observation, acts made from the first person and happening here and now, acts of the immanent mode, or only by static calculation, held in the third person, the omniscient? In either case three things are inseparable: hard-random information coming in, the particular moment at which it comes in, and the subjective standpoint from which its coming in can, at least potentially, be observed as having actually occurred. Those three, held together, are the bound triple again, and they are forever beyond what can be known through a purely omniscient theory.

Physics, Logic and Computation as One Triple

Three orderings of physics, logic and computation are often proposed, each placing one of them at the foundation. (*Logic* here is the study of what follows from what: identity, non-contradiction, valid inference.) *Materialism* takes physics as fundamental, and logic as an abstraction from physical regularity. *Mathematical Platonism* takes logic as fundamental, and physics as an instance of a mathematical structure. Computationalism takes computation as fundamental, and the universe as the output of a rule applied again and again. The structure of comparison places the three side by side as one triple: distinct, inseparable and not interchangeable.

The triple is a triple of roles, and role and character are two different questions: the role says which part of a comparison a thing plays, and the character says which of the four compounds it shows. By role, logic is omniscient: it is laid out as structure, and its truths hold the same however,

wherever and by whomever they are checked. Physics is immanent: it is the act of interaction itself, and its processes are joined by cause in time, along the order of findings, without a break in that order. Computation is transcendent: a procedure is defined over the range of its possible inputs, and each run selects one path through that range, turning a possibility into something actual. So computation is, in character, symmetry with discontinuity and purely omniscient, and, by role, transcendent; both answers hold, because each answers its own question.

The three are distinct, and none reduces to another. Physics does not reduce to logic: a purely symmetric description of a continuous process would perfect both characters, and the theorem forbids it; the mathematical description of physics is very powerful and never complete (Chapter 20). Physics does not reduce to computation: a purely computational universe would be deterministic, and determinism is impossible (Chapter 5). Computation does not reduce to logic: a proof is valid whatever the order in which its steps are checked, while a computation must be carried out in order, each step depending on the last. (The Appendix sets this beside the known match between proofs and programs.)

The three are inseparable, and none stands without the others. Physics without logic would be process without pattern, and even a physical event needs the identity and non-contradiction of logic to be one event. Logic without computation would be a set of truths with no way to pass from one to another, since checking a proof is a computation, a sequence of steps. Computation without physics would be procedure with nothing to run on, since every computation that is carried out is a physical process.

And each of the three keeps its own place: none can be exchanged for another. To treat physics as logic is to take the map for the land, in the picture of Chapter 1; to treat logic as physics is to take a truth for an event; and to treat computation as either is to take a procedure for what it computes or for what it runs on. The three are the three roles of comparison, read in the practice of inquiry.

What We Now Know

We now know that no computation can be everything that is real. A computation is one rule applied the same way at every step: deterministic, with no direction of its own, and with no inside. What is real has the bound triple at every act of comparison: hard randomness, the arrow of time and subjectivity, each requiring the next. A computation removes all three at once. So no computation can replicate a system in which the three are present; a mind cannot be uploaded, because a person is what a file is not; and the world is not, at its foundation, a computation, a conclusion that three independent routes reach. Every purely omniscient theory leaves an incompressible residue of hard randomness that it can neither predict nor explain. Computation itself keeps its place, beside physics and logic, as one of the three roles of comparison in inquiry.

If the subject is not a computation, what is it: what are choice, creation, and the self that chooses? In Chapter 7 I have given this book's answer. Choice is the coming of something new, an outcome not drawn from earlier conditions alone, given from a stance; from outside, the same event is hard randomness; and the self that chooses is never a finished structure.

For the question of this book, what cannot be and why, this chapter's exclusion, that no computation can be everything that is real, is the last of the book's exclusions. In Chapter 3 I said that each impossibility in this book is one limit met in a different place. Here the limit is met in the widest place of all, the claim that everything is a computation: a description complete from the outside cannot also have an inside, the fully fixed cannot also be open, and the reversible cannot also be directed. The limit is the Incommensuration Theorem, a structural fact about the nature of knowing.

Summary

Computationalism is the thesis that all of physical reality is, in principle, computational, and that the right computation, run on any machinery, would be a mind. That it is non-falsifiable is a claim about the kind of thesis it is: it is a metaphysics that presents itself as an empirical finding, and being non-falsifiable is a different thing from being false. Falsifiability applies to claims of how and why, and not to claims of what. A metaphysics answers to argument, and the consequences of a metaphysics, where those consequences reach into physics, answer to experiment. The argument of this chapter concludes that computationalism is false, as a matter of structure.

The reading of computation as deterministic, without a direction of its own and without an inside, and of pseudo-randomness as imitation rather than production of hard randomness, is a reading of known features of computation. In character, computation is symmetry with discontinuity, and purely omniscient; by role, beside physics and logic, it is transcendent; both answers hold. The binding of hard randomness, the arrow of time and subjectivity into one triple is derived from the Incommensuration Theorem and borne out in physics, where the Copenhagen interpretation keeps all three together and the many-worlds interpretation drops all three together. Taken all at once, as one act happening, the bound triple is consciousness.

The exclusions are structural exclusions, resting on the binding of hard randomness, the arrow of time and subjectivity into one triple, and on the Incommensuration Theorem. No computation can replicate a system in which the bound triple is present, and classical computation, however intricately organized, cannot be what consciousness is. Mind uploading is therefore excluded, as a category error and not an engineering problem. A simulated world containing real subjects is a restricted region of the one world in which the simulated world runs. Any purely omniscient theory must leave an incompressible residue of hard randomness. Three grounds (determinism, the two levels of every exact description, and the bound

triple) are independent routes to one conclusion: the world is not, at its foundation, a computation. The order of depth beneath them (metaphysics, physics, computation, mathematics), in which physics is deeper than computation because every computation carried out is a physical process, is a claim of the metaphysics itself. Physics, logic and computation are one triple of roles, distinct, inseparable and not interchangeable.

The argument is about what a computation is, and so it holds whatever a computation does. It leaves open how capable a machine can become at any task; the boundary it draws lies between processing and experience. Its result is a negative one: a computation, by being a computation, cannot be what consciousness is. It names consciousness as the bound triple taken as one act, and it builds no positive theory of consciousness on that name.

Computationalism is wrong in the way a category error is wrong, and not in the way a scientific theory is wrong. It takes one mode of description, computation, which is deterministic, has no direction of its own and has no inside, and asks it to do the work of another mode entirely. The bound triple of hard randomness, the arrow of time and subjectivity is the signature of the real: these three are what distinguish a world in which things actually happen from a world in which things merely appear to happen. Factor out hard randomness, and you have factored out the arrow of time, and with it any possible subject. Computation factors them all out. What remains may be extraordinarily useful, and it is; but it is not reality, and it is not a person. No increase in computing power, no refinement of algorithm and no breakthrough in hardware will close the gap, because the gap is a matter of kind and not of degree: closing it is impossible in principle, and therefore in all future practice.

What, then, does all that cannot be tell us about what is?

What Cannot Be, and What Can

One Result, Followed Everywhere

What does all that cannot be tell us about what is? I opened this book on four things that people are often told might one day be possible. The first of the four things was a message sent faster than light, so that a friend near a distant star could read the message before the light of your lamp had gone a fraction of the way. The second of the four things was a trip back to last Thursday. The third of the four things was a world that is someone's program, running on a computer in some other world, with us inside the program. The fourth was a universe that runs like clockwork, in which everything that will ever happen was fixed from the start. Each of the four has now been shown to ask for something that cannot be. That tells us a great deal about what is, because every exclusion in this book and every positive finding in it follow from one result about comparison. I have read physics from that single result.

Part One fixed the terms: comparison and its six parts; the four characters, symmetry, continuity, asymmetry and discontinuity; the Incommensuration Theorem and its two surviving pairings; the three axioms that the theorem brings with it; and, where the two kinds of law are in view, the matching of symmetry with the omniscient (structure, what holds the same for everyone) and of continuity with the transcendent (possibility, what makes things possible and new).

The three axioms say, together, that process comes before structure and possibility. The first puts process, the immanent, before structure and before possibility, and it makes those two a conjugate pair: each can be made perfect only at the cost of the other. The second gives the one-way

cycle in which a possibility gives rise to an interaction, interactions give rise to a structure, and a structure gives rise to a new possibility. Chapter 4 pictured it as the path that people crossing a lawn between a bus stop and a bakery wear into the grass: the path makes a shortcut possible, someone takes the shortcut tomorrow, and the path never makes the crossings that wore it. The third axiom makes the three distinct, inseparable and not interchangeable. Everything after Part One followed the one result about comparison, the Incommensuration Theorem, into physics.

What Is

Causes are real, at every scale where comparers are. A cause is found by a finite observer, in time, at the observer's own scale; determinism would be a total specification from no point of view, outside time, which would have to reach every scale. The theorem keeps continuity with asymmetry: a connection without a break, along an order that runs one way. A connection without a break, along an order that runs one way, is causation. Every world is therefore causal to some degree, has a lawful, determinate structure to some degree, and is deterministic to no degree at all. Science establishes causes, which is all it ever needed.

Chapter 5 laid out the three scales of description as one picture, standing one above another: at the bottom the very small, in the middle the observer reading a dial, at the top the whole. Each of the three scales has its own concept, its own limit of knowing, its own mystery, its own part in a measurement, and its own balance of choice and cause. A ruler shows the three parts of a measurement at once: the zero at its end, the marks of its unit, and its whole length, the range it can measure.

At the bottom is the very small, the scale of particles and fields. The concept that belongs there is determinism, because the very small is where determinism would have to be made good, all the way down, and

it cannot be. The limit of knowing there is the uncertainty relation; its mystery is the infinitesimal, the endlessly small; in a measurement it is the origin, the zero; and there choice and cause stand in strict opposition.

In the middle is the scale of observers and instruments. The concept that belongs there is causality: causes are perceived there. Its limit of knowing is that no cause runs backward, the limit at the act itself; its mystery is the subconscious, the most fundamental of the three; in a measurement it is the unit, the standard; and there choice and cause mix in a proportion that varies from act to act.

At the top is the very large, the cosmos as a whole. The concept that belongs there is teleology, the description of a thing by its purposes and ends, which names that whole. Its limit of knowing is marked by the light cone; its mystery is the mystery of the horizon, what lies beyond the reach of any comparer there; in a measurement it is the extent, the range; and there choice and cause are partners in proportion.

Hard randomness is real. The theorem guarantees a residual in every description, something no lawful description fixes, and in its strong form the residual is hard randomness: an outcome that no prior specification fixes, however complete, for any comparer whatever, even one with unlimited access to the state. Hard randomness and choice are one event, described from outside and from inside: from outside, the residual is randomness, and from inside, it is choice. A decision selects among what is already laid out; a choice realizes what was open; creation, done together with others, generates what is open. The real is the triple of causality, change and choice, and none of the three ways of breaking that triple (collapsing it, separating it, exchanging its members) can be carried out. Subjectivity (the stance from which a comparison is made), directed time and hard randomness are the bound triple: each of the three requires the next, as the first hearing of a piece of music in Chapter 22 has a before and an after, notes that surprise, and a listener to whom it happens. The subjective is continuous and one-way, and the objective is countable and the same everywhere, so the subjective is situated: in one time, at one place, in one actual outcome among the possible, and the present is where the sub-

jective and the objective meet. Meaning has a fixed point and cannot be removed, so nihilism, the view that nothing means anything, does not hold together. The self is neither a machine nor a ghost. It is the interdependent condition between the two, its continuity of context carried across the changing content of experience.

The same pattern runs through every pair of ideas that has been mistaken for a choice between two poles: determinism and pure randomness, science and mathematics, the physical and the non-physical, realism and idealism. In each pair the two poles are idealizations, and what is real is the interdependent condition between them.

Physics, by theorems of its own, confirms that determinism fails. If experimenters are free, the free will theorem of John Conway and Simon Kochen (2006) shows that particles are free in the same sense. The catastrophic theorem of determinism turns one failure of determinism anywhere into a failure at every scale. And superdeterminism, the only refuge left to determinism, removes the conditions under which any observation counts as evidence. Measurement, perception and comparison are one act, and every measurement makes actual a before and an after, a here and a there, and one actual outcome among the possible.

The book's readings of known physics found the structure of comparison at work in two of the oldest results of the last century. Emmy Noether's theorem of 1918 ties each law that stays the same across settings to an amount that is conserved. Chapter 2 read it as the exact form, inside one powerful kind of mathematics, of a fact of structure: wherever a content is kept the same across changes of its context, there is a matching conservation. The wave and the particle, the two faces that light and matter both show, and that Niels Bohr in 1928 called complementary, are the theorem's two survivors, as Chapter 3 showed: no single picture holds both faces perfectly, and the act of measurement is the third thing, which decides which face is shown.

The light cone is the shape of what can reach what. At each event, its light cone is the set of every event that a cause can connect to it, and the cone is the shape of the causal domain itself: the shape of what it means for anything to happen at all. The correlations that Bell's theorem concerns, between particles measured far apart, are real, and they carry no signal. Every measurement is a comparison made from inside the world, a count, a reading or a frequency set against a standard, and every unit is a standard of comparison. The six measures (mass, space, force, time, probability and possibility) stand in a grid of two rows and three columns, the frame of every unit. Planck's constant is the grain of possibility: the finest grain of the possible, below which no further difference is made. The speed of light, and the reduced Planck constant (Planck's constant divided by two pi), are the two limits of the domain of comparison, and the bridges between the roles that they join: the speed of light joins space and time, as Chapter 19 said, and Planck's constant joins time and possibility, as Chapter 18 said. The numbers of the two constants are set by the definitions of our units (the meter since 1983, the kilogram since 2019), and the book has made claims about the two limits as edges and none about their numbers.

Time has a structural arrow, present in every single comparison: the range of the possible comes first, then the act, then the record, and every record follows the act it records. That arrow needs no beginning in which disorder, which physicists measure as entropy, was low, and it comes before the arrow of heat and disorder, which points the same way. Time has one direction, and the three times of physics, the time laid out in the shape of space and time, the time of the smooth change of a quantum state, and the time of findings, are one triple: distinct, inseparable and not interchangeable.

The two great theories are joined side by side. General relativity, the theory of gravity as the shape of space and time, is continuity with asymmetry; quantum theory, the theory of matter and light, is symmetry with discontinuity; each is one survivor of the theorem, made physical, and the two are joined through the act of measurement, never merged. Every ex-

act theory has its mathematics and its findings in opposite characters, and each meets a residual where its mathematics, refined without limit, reaches the limit of its own kind: for quantum theory, the single definite finding, which its smooth mathematics does not contain; for general relativity, the singularity, where its description of the setting stops.

Chapter 17 showed what the opposite characters are in each theory. Quantum theory describes what is found, the content: matter and light. What it finds comes in countable steps and is the same everywhere, and it is written in the smooth mathematics of the quantum wave. General relativity describes the setting, the context: the arena of space and time. General relativity finds a geometry that is smooth and differs from place to place, and general relativity is written in equations that stay the same under every relabeling of places and times. In my book *Measurement; The Steps from Comparison to Physics* I show that these two pairings are the only such pairings, and that there is no third such pairing. In the first pairing, content has countable findings that are the same everywhere, and a continuous mathematics; in the second pairing, context has continuous, one-way findings, and a mathematics that is the same everywhere.

Any context carries four marks, and gravity has all four: it acts on everything, it admits no shield, it pulls on each body through the same mass by which the body resists being moved, and it carries no charge that comes in steps. Gravity is the context of physics, the arena of space and time itself. Chapter 17 also showed, from what a body feels when it accelerates, that gravity is the context of physics. A body can undergo exactly three kinds of acceleration, one for each way in which the free paths of neighboring bodies can relate to a common place. In free fall the paths draw together, and the body feels nothing, while the tide slowly draws things released side by side toward one another. In a room pushed in a straight line, as in a rocket or a lift, the paths stay parallel. On a spinning platform the paths spread apart. That there are exactly three is a structural claim, and only gravity's kind is one that no local push can produce, because it is the shape of the setting and not an action within it.

Measurement; The Steps from Comparison to Physics makes one more identification about the arena, and I take it from there: I identify the arena as curving where findings can be made, in proportion to how many can be made there. Laid out everywhere at once, that curving is the gravitational field. The Earth curves the arena more than the Moon does. The identification says where gravity, the context, bends, and by how much. Gravitation is not quantized. A graviton is at most a quantum of a description of the arena's small waves, never a quantum of the arena.

Measurement also counts the dimensions of the settings of comparison, and I take the count from there. A dimension is an independent way in which two findings can differ in their setting, as a point on a city map is fixed by two independent numbers, how far east and how far north. The settings of comparison have ten dimensions, ten independent ways in which two findings can differ in their setting: one of time, three of space, and six of possibility. The six of possibility are the independent ways in which what could be found might differ. You can walk to the next room, and you can read yesterday in a diary; you cannot walk to the outcome that did not happen. So a comparer moves through the four of time and space, and meets the six of possibility only as the range from which each finding is one selection.

Measurement derives one more fact about the four of time and space. In the geometry of space and time, the measure of the separation between two events counts its stretch of time with one sign and its three stretches of space with the other. That difference of sign is what gives each event the light cone of Chapter 12, which keeps the before of each event apart from its after. Every law is a limit of what can be reached by comparison, and the theorem is why those limits appear, and why no total theory, one description with no residual left over, can be had.

Physics has three horizons, which are edges of what can be known from here. The outer edge is marked by the speed of light, the inner edge is marked by Planck's constant, and the third edge is the edge of the knowing self. Planck's constant sets no least length and no least duration.

Any one horizon moves back as it is approached, and the kind of edge it is stays permanent. Between the two physical edges lies everything that can be measured.

There is one world. A simulated world with real subjects shares the bound triple of subjectivity, directed time and hard randomness with the world that runs it, and it is a restricted region of that one world. Physics, logic and computation are one triple. Mathematics works in physics because a mathematical description can reach its findings as closely as one likes, and it is never complete because of the residual that every description leaves.

The Eight Exclusions

Each exclusion of this book is the same theorem met in a different place, and each was settled where it was met. The four requests the book opened with are among them.

The universe that runs like clockwork, with everything fixed from the start, is determinism. The theorem forbids perfect symmetry with perfect continuity, and determinism asks for exactly that pair: one law the same in every setting, and every state specified without a gap. So determinism is excluded, and it cannot even be moderately held, because any crack in the specification is the end of the specification. Superdeterminism asks for everything determinism asks for, and more, since it fixes the experimenters' choice of question along with the answer; it is excluded with determinism, and it collapses choice into cause, which is a category error.

The friend near a distant star, reading your message before the light of your lamp had gone a fraction of the way, would need a signal faster than light. The edge of comparison is the edge of causation. A signal faster than light asks for a comparison beyond the edge of comparison, and for some comparer its record would come before its act. So no signal crosses the edge that the speed of light marks, for any being, in any universe, that interacts with anything else.

The laws of physics allow every lawful process to be run backward lawfully. A faithful reversal, which retraces exactly every state back to exactly the start, would need the symmetry of the laws and the continuity of the state perfected together, the pair the theorem forbids; so no reversal of time is faithful, and the record is not undone. And "time travel", the trip back to last Thursday in the sense in which the phrase is used, asks one time to be both the time laid out in the shape of space and time and the time of findings, each perfected; nothing is both, so the phrase names no concept.

The two great theories are each one survivor of the theorem. A single theory that perfected both survivors would perfect the pair that the theorem forbids. The effort to make the two great theories one has the character of an effort to perfect symmetry and continuity together, and no single theory that perfects both great theories can be had.

A world that is someone's program, with real people in it, turns out to be a region of the one world that runs it. The one world contains every world that is run within it. A simulated world with real subjects in it is a region of the one world in which it runs, as a prison cell is part of the one world and not a second one; and a simulation with no one in it is no world at all; so our world is not a simulation run in some other world. A classical computation, one rule applied the same way at every step, cannot build the bound triple, so no program, by being a program, is a mind; and no discrete computation can capture a continuous reality. No computation is everything that is real.

Every one of these requests names nothing, as the request for the color of the number seven names nothing, and no one of them is a matter of difficulty; each is excluded in principle, and therefore in all future practice.

What Kind of Book This Has Been

Each chapter said what kind of claim it makes, and the three kinds were kept apart. The claims of the structure of comparison itself are the ground: the six parts that every comparison has, the four characters, the Incommensuration Theorem, the three axioms, the two kinds of law, and the six measures that every physical measurement uses. They are argued from the structure of comparison and from nothing else. The readings are readings: the claim in each is that a known result holds for the reason of structure the chapter named, and the reader can check that reason against the result. The exclusions are exclusions: each says what will not be found, and each is a consequence of the structure of comparison that reaches into physics, into the world in which experiments are made, and experiment speaks to each exclusion there.

In this book I read the physics that exists. No number is derived in it, and nothing here is a physics in competition with that physics. The framework does not replace the theories it reads. It says why they are shaped as they are, and why the effort to make them one has the character of an effort to perfect symmetry and continuity together.

This book is the companion of the books of derivation, and two things it leans on are argued in full there. The match between the six parts of comparison and the six measures is argued member by member in *Comparison; The Necessary Basis*, where the count, six, is derived as three roles taken twice, once as content and once as context. And gravity as the context of physics is reached in those books by further routes, which do not pass through the grid of the six measures at all. The Appendix names those routes, and the ground on which the steps of *Measurement* rest.

I set out to place physics at the boundary of comparison, and to show that the boundary explains more than it obstructs. What is beyond the boundary is beyond the book.

The Fact Beneath It All

The Incommensuration Theorem stands beneath every result in these pages. It is a structural fact about the nature of knowing: symmetry and continuity cannot both be perfected in any comparison, and neither can asymmetry and discontinuity. Every act of knowing is an act of comparison, or is built from acts of comparison, so the theorem, together with the intrinsics of comparison from which it is built, is inherent in every single act of knowing, and it is the deepest and most powerful epistemic basis possible. What this book has shown is how much of the world, as physics finds it and as each of us lives in it, that one fact reaches. The world is causal without being deterministic at all, and that is why choosing is real.

Appendix

For the Physicist

For Chapter 1: Measurement and the True Zero

One point that Chapter 1 states in words is given here in the physicist's terms.

Negative absolute temperatures

A reader who knows some physics may ask whether negative temperatures break the rule of the true zero. Physicists do speak of negative absolute temperatures, in certain systems of spins (collections of tiny atomic magnets) whose energies are bounded above, so that the system has a highest possible energy. Such systems are, in the ordering of hot and cold, hotter than any positive temperature. The negative number reflects how temperature is defined there, and it names no quantity below the absence of thermal motion. For ordinary matter in thermal equilibrium (matter that has settled to one even temperature throughout), the Kelvin zero is the absence that Chapter 1 describes.

For Chapter 2: Conservation and the Group

Two points that Chapter 2 states in words are given here in the physicist's and the mathematician's terms.

Conservation beyond the Lagrangian form

Chapter 2 claims that conservation is a necessity that comes before any particular system is written in the Lagrangian form, and that Noether's theorem is the exact form, inside that kind of mathematics, of a feature of structure that holds of every symmetric domain as such: wherever a content is kept the same across changes of its context, there is a matching conservation. That feature holds whether or not the domain can be written in the Lagrangian form, whether or not its symmetry can be written as a smooth family of transformations (what mathematicians call a Lie group), and whether or not anyone has written the conserved current down. The account given here does not prove Noether's result again, and it could not. Her theorem tells the physicist exactly which conserved quantity goes with which symmetry, and how to compute it; the account tells the physicist why there was going to be one.

The axioms of a group, and abelian groups

Chapter 2 gives the four conditions of Nathan Carter (2009) on a set of actions: the actions are well defined, reversible, deterministic in the mathematician's sense, and composable. In formal terms the four conditions correspond to the four axioms of a group: closure; associativity, which says that grouping a sequence of moves in a different way leaves the result unchanged; identity, which says that there is a move that changes nothing; and inverses, which says that every move has a move that undoes it.

Where the order of the actions makes no difference, the group is called *abelian*. The group of a Rubik's cube is not abelian: turn the top face and then the front face, and you get a different cube from turning the front face and then the top. The hand of a clock, turned forward by whole hours, has an abelian group: two hours and then three give the same result as three hours and then two.

For Chapter 4: The Zero-Positive Principle

One point that Chapter 4 states in words is given here in exact terms.

Zero-positivity and the qualifier of the theorem

Chapter 4 states that every domain is zero-positive: its fundamental measures have an absolute minimum and no absolute maximum. That principle is a statement about the measures of a domain, and the qualifier in the theorem does not rest on it. The qualifier "absolutely applied" has its own meaning, holding exactly, at the foundation of a description, which Chapter 3 gives.

For Chapter 7: Two Rotations of Role

The technical relation that Chapter 7 states in words is given here: a modal rotation rests on the fact that a role belongs to a position of description, and never to a thing.

The modal rotation and the Axiom II rotation

A modal rotation is the passage from one modal description of one event to another: the event does not change, and the position from which it is described does. It is a different relation from the shift of bindings between a domain and its successor that Incommensuration calls the Axiom II rotation, in which the roles of the three modalities move as one domain is built on the next. Both turn on the same fact, that a role belongs to a position of description and not to a thing (Measurement, Chapter 14, Section VIII).

For Chapter 9: The Free Will Theorem

Two things that Chapter 9 states in words are given here in the terms of the literature: the names Conway and Kochen gave their assumptions, and the shape of the step from the assumptions to the conclusion. The theorem belongs to Conway and Kochen and carries the warrant of their papers (2006, 2009); the combinatorial result it uses belongs to Kochen and Specker (1967).

The names of the assumptions

Conway and Kochen named their assumptions SPIN, TWIN and FIN in 2006, and in 2009 they replaced FIN with a weaker assumption they called MIN. FIN says that there is a finite upper bound to the speed at which information can be passed. MIN says that when the two measurements are spacelike separated, each experimenter can choose a setting freely, and the response of the particle in the other laboratory is independent of that choice. A theorem that reaches its conclusion from a weaker assumption asks less of the world, and that is the sense in which the 2009 theorem is stronger.

Conway and Kochen give the assumption of no backward causation no name of its own; it is part of what their assumptions about place and about freedom take for granted. Chapter 9 states it separately, so that each of the four assumptions can be seen.

The shape of the proof

SPIN says that the squared spin of a spin-one particle, measured along any three mutually perpendicular directions, gives the answers 1, 0 and 1 in some order (read "one, zero and one"). A map that gives each direction an answer, 0 or 1, and gives exactly one 0 on every perpendicular triple, is what Kochen and Specker (1967) proved cannot exist for a suitable finite set of directions; Conway and Kochen use a set of thirty-three directions given by Asher Peres (1991).

Suppose the response of the first particle, measured along a triple of directions, is a function of the information available to it beforehand together with the triple chosen. By MIN, the second experimenter's free choice of one direction, made at spacelike separation, does not affect the first particle's response, and the first experimenter's choice of triple does not affect the second particle's. By TWIN, when the second direction is one of the first triple, the two answers along it match. So the first particle's answer along any one direction does not depend on which perpendicular triple contains it, and the responses would define the map that Kochen and Specker ruled out. No such function exists, and the responses are not functions of anything earlier.

For Chapter 13: The Direction of Time

Four pieces of Chapter 13 are given here for the specialist in their technical form. Chapter 13 states the result of each in words.

The three layers of the Epistemic Sandwich

The Epistemic Sandwich runs across three layers of physical description, each at a different scale, and the three scales have the natures of the immanent (our own scale), the omniscient (the very large), and the transcendent (the very small). Each layer is described in its own terms and has its own key concepts.

Layer one is the mesoscopic scale, the scale of observers and their instruments, described by Newton's mechanics, by the theory of electricity and magnetism, and by thermodynamics, in terms of knowing. Its key concepts are causality, change and choice, which are omniscient, immanent and transcendent respectively.

Layer two is the macroscopic scale, the very large, described by general relativity, in terms of the cosmos. Its key concepts are the meaningless and the meaningful, with pattern as the term between them.

Layer three is the microscopic scale, the very small, described by quantum mechanics, in terms of what is. Its key concepts are determinism, relation and pure randomness. Determinism (symmetry and continuity both perfected) and pure randomness (asymmetry and discontinuity both perfected) are the two idealizations that the Incommensuration Theorem excludes, and what is real at layer three lies between them.

The anchor and the theorem of Conway and Kochen

The theorem of Conway and Kochen is a theorem about what is fixed in advance, and not about the direction of time. The step from an unfixed response to a reversal that cannot be faithful is my reading, and not their result, and it rests on the same bridge between lawful and faithful reversal that carries the first road of Chapter 13. The whole anchor is conditional on the "if" of their theorem. That condition belongs to the theorem of Conway and Kochen. It is no condition on the Incommensuration Theorem, which is a structural fact about the nature of knowing, and no condition on the two roads to the arrow, which stand on that theorem, on the bridge, and on the order of the second axiom.

The strength of the reading of time-reversal violation

Chapter 13 reads the observed breaking of C P in the weak interaction, and with it, by the CPT theorem, the breaking of time-reversal symmetry, as a place where the impossibility of perfecting the symmetry of time has become visible to instruments. That reading has an exact strength. The argument says that some asymmetry must exist. The argument gives it no size, names no interaction for it, and names no mechanism. The requirement stands on its own ground, and it would stand whether or not this particular violation had been found; the argument did not predict the observed violation, and it finds the observed violation consistent with the requirement.

The philosopher Craig Callender examined in 2000 whether time is "handed" in the quantum world, that is, whether it has a built-in direction the way a hand is a left hand or a right hand, and concluded that the evidence is more nuanced than the CPT theorem alone suggests. The structural argument works at a level beneath the mechanism of the weak interaction, and it rests on nothing in that mechanism, so the nuance of the mechanism does not reach it.

The fit of the CPT theorem with the argument (nature keeps the combined symmetry and breaks the symmetry of time alone) I state as consistency, not prediction.

The reach of the argument

The entropic arrow needs further premises, statistical mechanics and the Past Hypothesis among them, which lie outside the argument of Chapter 13. The reach of the argument is the direction of time; the size of any asymmetry, the low entropy of the early universe, and the second law itself lie outside it.

For Chapter 14: Three Kinds of Time

A continuation whose geometry is only continuous

Chapter 14 states that a continuation of a shape of space and time across a Cauchy horizon is itself a shape of space and time, whose clock times and order of what can reach what are carried by its geometry, and that this holds even for a continuation whose geometry is only unbroken, with creases, and not smooth. In the physicist's terms, an extension of a spacetime is a manifold with a Lorentzian metric into which the development embeds isometrically as a proper subset. This definition holds for an extension with a merely continuous metric as well. There the causal order, the order of which events can reach which, is defined by the locally Lipschitz causal curves: paths that, within any small patch, bend no more

sharply than some fixed limit. That causal order can be less regular than the causal order of a smooth metric, and it is still carried by the geometry, so what the extension adds is geometric time.

The formula for the cosmological constant

Chapter 14 gives in words the formula by which The Gravitational Passage, Chapter 13, identifies the cosmological constant with the share of the energy of empty space that an observer inside empty space can reach. The formula is

$\rho_\Lambda = c_1 \, f \, M_{Pl}{}^4 / S, \quad \Lambda = 8\pi G_N \, \rho_\Lambda,$

in words: the energy density that the cosmological constant gives to empty space, written ρ_Λ ("rho lambda"), is the count c_1 ("c one"), times the factor f, times the Planck mass M_{Pl} ("M Planck") taken to the fourth power, divided by the horizon count S; and the cosmological constant Λ (capital lambda) is eight times pi, times Newton's constant of gravity G_N ("G N"), times that energy density.

The factor f is one minus three divided by eight times pi, about 0.881. It is the factor of what that book calls the external value, $f \, M_{Pl}{}^4$, which is made of a ground-state part, $M_{Pl}{}^4$ itself, together with a classical part, minus three divided by eight times pi of $M_{Pl}{}^4$. $M_{Pl}{}^4$ is the Planck energy density, in the units physicists use in which the speed of light and Planck's constant each count as one, and the Planck mass is formed from Newton's constant of gravity, which is measured. The count c_1, equal to three, is the count of the channels of a measurement, which that chapter sets equal to the Euler characteristic of the complex projective plane, written $\mathbb{C}P^2$ ("C P two"): a count that stays the same however the shape is bent, and which for that shape is three. The horizon count is $S = \pi R^2 / l_{Pl}{}^2$, in words: pi times the square of the radius R of the observer's particle horizon, divided by the square of the Planck length l_{Pl} ("l Planck"); that is a quarter of the

horizon's area counted in squares one Planck length on a side, the count that Jacob Bekenstein (1973) and Stephen Hawking (1975) gave for a horizon.

Standing: in The Gravitational Passage the formula is an identification, from which a value of the cosmological constant is computed. The three and the 0.881 are set by the structure, once that book's identification is taken; Newton's constant and the radius of the horizon are measured.

For Chapter 15: Strong Cosmic Censorship and the Time-Machine Literature

The results that Chapter 15 states in words are given here in the physicist's terms. Every result here is stated in words in that chapter. Each belongs to the literature it cites and carries the warrant of its authors. The argument of this book uses one principle from this literature, the principle of self-consistency of Friedman and his coauthors, at one step of Chapter 16; it uses no result of strong cosmic censorship, no energy condition, and neither Hawking's theorem nor the theorem of Kay, Radzikowski and Wald. Each symbol is named in words where it first appears. The signs $>$, $<$ and $=$ are read "is greater than", "is less than" and "equals".

Closed timelike curves and time machines

A *closed timelike curve* is a timelike worldline that returns to its own starting event. Some spacetimes that satisfy Einstein's equations contain them. The first example is Gödel's (1949). The Gödel spacetime has a closed timelike curve through every point, and it has no Cauchy surface: no spacelike hypersurface in it is met exactly once by every inextendible causal curve. It is totally vicious: every point lies on a closed timelike curve, so no set in it is achronal (no set whose points no timelike curve joins), and it has no partial Cauchy surface whose Cauchy horizon could be formed. So the general statement "every spacetime with closed timelike curves has a Cauchy horizon" is false.

A *time machine*, in the sense of the literature, is a region in which closed timelike curves appear in a spacetime that had none before them: the spacetime is globally hyperbolic to the past of some time, and causality is violated only to the future of it. The boundary of the region that the initial data still determine is then a Cauchy horizon of the initial surface, called in this setting the *chronology horizon*. Morris and Thorne (1988) described a traversable wormhole held open by matter that violates the energy conditions; Morris, Thorne and Yurtsever (1988) showed that moving one mouth relative to the other, so that the mouths age differently, produces closed timelike curves through the wormhole.

Hawking (1992) took as his model of a manufactured time machine a *compactly generated* Cauchy horizon: every past-directed null generator of the horizon enters, and stays within, a compact set. The model expresses the demand that the causality violation be made in a finite region, by a finite apparatus, without curvature singularities; it is Hawking's model, not a theorem that every time machine has this form. The theorem about that model: a compactly generated Cauchy horizon cannot form from a noncompact partial Cauchy surface if the weak energy condition holds (the energy density measured by every observer is non-negative). The abstract of the paper states the corresponding result for the averaged condition: if the causality violation develops from a noncompact initial surface, the averaged weak energy condition must be violated on the Cauchy horizon. Hawking proposed, as a conjecture, that the laws of physics prevent closed timelike curves from appearing at all: the chronology protection conjecture.

Kay, Radzikowski and Wald (1997): for a free scalar field on a spacetime with a compactly generated Cauchy horizon, the two-point distribution of any Hadamard state, extended as a bisolution beyond the region the initial data determine, fails to be of Hadamard form at every point of the base set of the horizon, the set of points at which the past-directed null generators accumulate. The stress-energy of the field is therefore not defined there by the usual local procedure. The result concerns the points of the base set, not every point of the horizon.

Friedman, Morris, Novikov, Echeverria, Klinkhammer, Thorne and Yurtsever (1990) studied the Cauchy problem for fields in spacetimes with closed timelike curves, under the *principle of self-consistency*: the only solutions admitted are those that are consistent everywhere, including along every closed timelike curve. Under that principle, whatever a body carries around a closed timelike curve must arrive at the starting event in exactly the state in which it set out.

A chronology horizon is a Cauchy horizon. Strong cosmic censorship asserts that the maximal globally hyperbolic development of generic initial data cannot be extended, in a stated regularity class, beyond its boundary; it is posed for the vacuum equations and for matter models that satisfy energy conditions, while Hawking's theorem shows that a compactly generated chronology horizon needs a violation of the weak energy condition. Within those models the weak energy condition holds, so Hawking's theorem already excludes a compactly generated chronology horizon formed from a noncompact partial Cauchy surface. The conjecture, in a given regularity class, adds that within those models Cauchy horizons that can be crossed in that class, chronology horizons included, arise only from non-generic initial data. A time machine of the kind Hawking models therefore needs matter outside those models.

The conjecture and what it protects

Penrose introduced strong cosmic censorship in the late 1970s, as a concern about the predictive character of general relativity. Suitable initial data on a spacelike hypersurface determine a unique spacetime, the *maximal globally hyperbolic development* of the data: the determinism of general relativity in the sense of Laplace. In the terms of this book, that uniqueness is a property of the formalism, a law, which exists as structure, and not of what any comparison finds. The Kerr family (rotating vacuum black holes) and the Reissner-Nordström family (charged black holes) have Cauchy horizons in their interiors: null hypersurfaces beyond which the initial data no longer determine the spacetime, and across which the

solution can be continued in more than one inequivalent way, different regions being attached to the same maximal globally hyperbolic development without violating the Einstein equations. The conjecture asserts that this edge is not generic: under generic perturbation of the initial data, the Cauchy horizon is replaced by a singularity beyond which no extension of the stated regularity exists, so the maximal globally hyperbolic development of generic data is inextendible in that class. It does not deny that exact solutions have Cauchy horizons. A *Cauchy horizon* is the boundary of the domain of dependence of the initial data, in a spacetime that extends their maximal globally hyperbolic development; not every boundary of a development is one: the spacelike singularity of Schwarzschild is not.

The three regularity formulations

The mathematical content of the conjecture depends on the regularity class in which inextendibility is required, and the three formulations in use are not equivalent. Here the metric is the spacetime metric, the Christoffel symbols are its first derivatives in combination (the connection coefficients), and C^k ("C k") names the functions with k continuous derivatives.

The *C^0 formulation*: the metric cannot be extended as a continuous function beyond the boundary of the maximal globally hyperbolic development. It is the strongest requirement: if no continuous extension exists, then no extension exists that is continuous with locally square-integrable Christoffel symbols, and no C^2 extension exists.

The *Christodoulou-Chruściel formulation* (Chruściel 1991; Christodoulou 2009): no extension exists whose metric is continuous and whose Christoffel symbols are locally square-integrable. For a continuous, non-degenerate metric this is the requirement that no extension has its metric in $C^0 \cap W^{1,2}_{\text{loc}}$ (read "C zero intersect W one two, local": continuous, with first derivatives square-integrable on compact sets; $W^{1,2}_{\text{loc}}$ is also written H^1_{loc}). In four dimensions a metric in $W^{1,2}_{\text{loc}}$ need not be continuous, so

the continuity is part of the requirement. $W^{1,2}_{\text{loc}}$ is the weakest regularity in which the Einstein equations can still be stated in the weak (distributional) sense. So, under this formulation, a spacetime may extend as a continuous manifold with a continuous metric while the Einstein equations, even in weak form, cannot be posed beyond the boundary. By "strong cosmic censorship" without qualification I mean this formulation, because an extension on which the equations cannot be posed is one the equations do not govern.

The C^2 *formulation*: no C^2 extension exists. A sufficient condition, and the one by which it is usually shown, is that curvature blows up at the boundary: that the Kretschmann scalar (the full contraction of the Riemann tensor with itself), or another curvature invariant, diverges. A body crossing such a boundary meets unbounded tidal forces, because tidal forces are the relative accelerations of neighboring geodesics, given by the curvature. The C^2 formulation is the weakest requirement: curvature can blow up while the metric still has a continuous extension. The order: C^0 strongest, Christodoulou-Chruściel between, C^2 weakest.

The blue-shift, Price's law and mass inflation

The *blue-shift instability*: an observer falling toward the Cauchy horizon receives radiation that fell in after it, compressed into ever shorter wavelengths; the energy density it measures diverges at the Cauchy horizon. The divergence is not a coordinate effect; it follows from the causal structure near the inner horizon. With vanishing cosmological constant (written Λ, capital lambda, here $\Lambda = 0$), perturbations along the event horizon decay at a polynomial rate in advanced time, the inverse-power tails of *Price's law*, while the blue-shift amplifies at an exponential rate. Against polynomial decay the exponential amplification is expected to prevail, so that derivatives of the perturbation blow up at the Cauchy horizon, for every polynomial rate of decay. For $\Lambda = 0$ the competition is not close; the ratio criterion below applies only for $\Lambda > 0$.

Poisson and Israel (1990) gave the first model of the back-reaction on the geometry, and Ori (1991) an exact solution of a spherically symmetric charged model: *mass inflation*. In spherical symmetry, the *Misner-Sharp mass* (written m) is the quasi-local mass defined from the area radius (written r, the radius of a symmetry sphere read from its area) by

$1 - 2m/r = g^{\mu\nu}\nabla_\mu r \, \nabla_\nu r,$

in words: one minus twice the Misner-Sharp mass divided by the area radius equals the squared gradient of the area radius, where $g^{\mu\nu}$ ("g upper mu nu") is the inverse spacetime metric and ∇_μ ("nabla sub mu") is the covariant derivative. The Misner-Sharp mass diverges at the Cauchy horizon, and the area radius does not go to zero there. In Ori's solution the Cauchy horizon is a null singularity at which the mass diverges and the area radius is finite and non-zero; the metric functions take finite values on it, and the tidal distortion of a body crossing it is finite. Such a singularity is a *weak null singularity*: a null boundary at which the metric extends continuously but curvature, or the quasi-local mass, diverges.

Dafermos (2003, 2005), for the spherically symmetric Einstein-Maxwell-scalar field equations: for data that settle down to a Reissner-Nordström black hole, the metric, the area radius and the scalar field extend continuously to the Cauchy horizon, and the area radius stays bounded away from zero on the Cauchy horizon near its future end point (written i^+, "i plus"). Under a lower bound on the decay of the field along the event horizon, the Misner-Sharp mass blows up identically along the Cauchy horizon, and no C^1 extension of the metric exists. These results are established in spherical symmetry with charge; for the rotating case without symmetry, these results are the guiding example, and not a theorem.

What has been proved

The rigorous results divide by the symmetry of the spacetime and by the regularity class of the statement.

Spherical symmetry. Luk and Oh (2019, parts I and II) proved strong cosmic censorship in the C^2 formulation for the spherically symmetric Einstein-Maxwell-(real) scalar field system with two-ended asymptotically flat initial data: for generic data in that class (an open and dense set), the maximal globally hyperbolic development has no C^2 extension to the future. Part I proves this under an averaged lower bound on the field along the event horizons; Part II proves that the bound holds for generic data. Sbierski (2018) proved that the Schwarzschild spacetime has no continuous extension at all; it has no Cauchy horizon in its interior and ends at a spacelike singularity, the contrast case of a spacetime that ends in the strongest sense.

Rotating vacuum, continuity. Dafermos and Luk (published 2025; arXiv 2017): for vacuum initial data on a spacelike hypersurface already inside the black hole, approaching the data induced by a Kerr solution with $0 < |a| < M$ (read: the spin parameter a, the angular momentum per unit mass, nonzero and smaller in size than the mass M) at an inverse polynomial rate, the maximal development extends, with a continuous metric, across a non-trivial piece of Cauchy horizon near i^+. The theorem makes no assumption about the exterior. Two further steps connect it to black holes that form from exterior data: the stability of the Kerr exterior, which supplies the decay along the event horizon; and a second theorem, announced by Dafermos and Luk, carrying that decay to the interior hypersurface of the first theorem. With both steps, the C^0 formulation is false for the vacuum Einstein equations near Kerr. The stability of the Kerr exterior is established for slowly rotating black holes ($|a|$ much smaller than M): Klainerman and Szeftel (2023), with Giorgi, Klainerman and Szeftel (2024) and Shen (2023). For the full subextremal range, linear stability is the subject of a preprint of Häfner, Hintz and Vasy (2025), and non-linear stability of a preprint of Szeftel (2026), which rests on companion papers; neither preprint is yet refereed. The second interior theorem of Dafermos and Luk is not published.

Rotating vacuum, singularity. For the rotating vacuum case the literature states the Christodoulou-Chruściel formulation as a conjecture, and the results below are proved each under its stated assumptions. Dafermos and Luk state that their proof suggests that generic Cauchy horizons are weak null singularities and that a revised form of strong cosmic censorship holds; the conjecture they number is the C^0 formulation, which they disprove. Luk and Sbierski (2016): for the linear wave equation on Kerr, under upper and lower bounds on the decay along the event horizon, the solution has infinite non-degenerate energy at the Cauchy horizon, so it is not in $W^{1,2}_{\text{loc}}$ there. Sbierski (2023): for linearized gravity, under integrated upper and lower bounds on the Teukolsky field along the event horizon (the Teukolsky field is the curvature component that carries linearized gravitational perturbations of Kerr), that field is singular at the Cauchy horizon in a frame regular there. Sbierski (2026): without symmetry, curvature blow-up at a weak null singularity rules out locally Lipschitz extensions.

Non-linear rotating case at the Lipschitz level. Gurriaran (2026, preprint) assumes vacuum data inside a dynamical black hole settling to subextremal Kerr with $0 < |a| < M$, a Price's-law behavior of the Teukolsky field, and the non-vanishing of certain constants that are expected, but not proved, to be non-zero for generic data; under these assumptions the Riemann tensor blows up at the Cauchy horizon near i^+, and the metric has no locally Lipschitz extension there. Luk and Sbierski (2026, preprint) pose a characteristic initial-value problem with smooth data on a dynamical event horizon that settle down to a subextremal, strictly rotating Kerr black hole, with suitable upper and lower bounds, and prove that a weak null singularity forms, across which the metric extends continuously but not as a locally Lipschitz ($C^{0,1}$, "C zero one") function. In both preprints the genericity enters through the assumed bounds on the data, and the genericity is not itself proved.

The state of the rotating vacuum case. Continuous extension across a piece of the Cauchy horizon is a theorem, under the stated interior hypotheses. Failure of Lipschitz extension is proved under stated assumptions on the data. Failure of extension in $C^0 \cap W^{1,2}_{\mathrm{loc}}$ is the Christodoulou-Chruściel formulation itself, which the literature states as a conjecture. A locally Lipschitz metric has locally bounded first derivatives, and locally bounded functions are locally square-integrable, so every locally Lipschitz extension is also a $W^{1,2}_{\mathrm{loc}}$ extension; the converse fails, so the absence of Lipschitz extensions does not imply the absence of $W^{1,2}_{\mathrm{loc}}$ extensions. These results narrow the gap between what is proved and the Christodoulou-Chruściel formulation, and the formulation itself remains stated as a conjecture.

A positive cosmological constant

For black-hole spacetimes with $\Lambda > 0$ (de Sitter asymptotics), perturbations along the event horizon decay exponentially, at a rate written α (alpha) set by the *spectral gap*: the absolute value of the imaginary part of the least-damped quasinormal mode, the spectral gap of the wave operator on the black-hole background. The cosmological horizon adds a second boundary. The blue-shift amplifies exponentially at a rate set by the surface gravity of the inner horizon, written κ_- (kappa minus). The competition is governed by the ratio

$\beta \equiv \alpha \,/\, \kappa_-$,

in words: beta is defined as the decay rate alpha divided by the inner-horizon surface gravity kappa minus. Hintz and Vasy (2017) proved that linear waves on such backgrounds have regularity $H^{1/2+\beta-\varepsilon}$ at the Cauchy horizon for every $\varepsilon > 0$ (read: Sobolev regularity one half plus beta minus epsilon, for every positive epsilon). So for $\beta > 1/2$ the waves extend in H^1 (read "H one": Sobolev regularity one, functions that are square-integrable together with their first derivatives, the regularity written H^1_{loc} above when it is asked only on compact sets), and for $\beta < 1/2$ the H^1

norm, the size of a wave measured in that space, is expected to blow up. Cardoso, Costa, Destounis, Hintz and Jansen (2018) used this as the criterion for the Christodoulou-Chruściel formulation in the linear setting: where $\beta < 1/2$ the formulation is expected to hold, and where $\beta > 1/2$ it is endangered. The threshold $\beta = 1/2$ is a criterion drawn from the linear regularity theorem, not a non-linear theorem. Fast decay along the event horizon endangers the conjecture, because it leaves less of the perturbation to feed the blue-shift instability; slow decay protects it. The sentence "decay overcomes the blue-shift, so strong cosmic censorship holds" states the relation backwards.

Kerr-de Sitter. Dias, Eperon, Reall and Santos (2018) found, for every non-extremal value of the parameters, quasinormal modes that decay slowly enough for the formulation to be respected: $\beta < 1/2$ throughout the subextremal range. The analysis is linear, uses geometric optics with numerical computation, assumes smooth data, assumes that generic data excite the slowly decaying modes, and does not cover the extremal limits. The rotating case with $\Lambda > 0$ is therefore expected to respect the Christodoulou-Chruściel formulation.

Reissner-Nordström-de Sitter. Cardoso, Costa, Destounis, Hintz and Jansen (2018) showed that near-extremal Reissner-Nordström-de Sitter black holes, perturbed by a smooth, neutral, massless scalar field, have $\beta > 1/2$. For that class of data on that class of backgrounds the Christodoulou-Chruściel formulation is violated in the linear analysis, by the regularity theorem of Hintz and Vasy with β computed numerically: linear waves extend across the Cauchy horizon at a regularity the formulation was designed to exclude. The same authors ("still subtle", 2018) found that the violation persists in part of the parameter space for a charged scalar field. The violating set of parameters is open in the space of black-hole parameters, not a set of measure zero, so no argument that discards sets of measure zero can remove it.

The literature discusses three responses. First, charged or massive fields change the rate of decay: Hod (2019) argued that a charged field restores sufficiently slow decay; Dias, Reall and Santos (2019) and Cardoso

and coauthors (2018) found that with a charged scalar field the violation persists in part of the parameter space; on this response the matter content sets the decay rate, and a smooth, neutral, massless scalar field is not the physically relevant class of perturbation for a charged black hole. Second, rough rather than smooth initial data restore the blow-up (Dafermos and Shlapentokh-Rothman 2018; Dias, Reall and Santos 2018); on this response the restriction to smooth data may exclude the generic case. Third, the stress-energy of quantum fields diverges at the Cauchy horizon of Reissner-Nordström-de Sitter black holes, whatever the classical rate of decay (Hollands, Wald and Zahn 2020). Each response addresses a different feature of the violation, and none is accepted in the literature as closing it. A further position is held in the literature: that a conjecture whose formulation must exclude an open set of classical counterexamples with smooth data may be the wrong formulation of the conjecture; I record the position and do not decide it. Hod (2020) gave a short physical argument for strong cosmic censorship from Bekenstein's generalized second law; it points in the direction of censorship, and it is a physical argument, not a proof in the sense of mathematical relativity, whatever its title says: its premises concern the scope of the generalized second law in de Sitter backgrounds, and that scope is debated. Van de Moortel (2025) reviews the field.

Genericity from three directions

The conjecture asserts that blow-up occurs for all initial data outside a negligible set, and infinite-dimensional analysis offers more than one notion of a negligible set. Three statements below put the question with three of them: meager sets, sets of Gaussian measure zero (with prevalence as a variant), and nowhere dense sets, in the language of forcing. The third implies the first, and not conversely; the second is independent of the first. None of the three proves a statement of genericity; each states exactly what "generic" asserts in its own notion.

Notation. The *data space*, written $\mathscr{X}$ (script X), is the space of admissible initial data, with a weighted Sobolev topology in which it is a *Polish space*: a separable space whose topology comes from a complete metric. The *failure set*, written $\mathscr{V}$ (script V), is the subset of the data space whose maximal globally hyperbolic development admits an extension across the Cauchy horizon with metric in $C^0 \cap W^{1,2}_{\mathrm{loc}}$: the set of data for which the Christodoulou-Chruściel formulation fails.

Topological genericity. Assume that the map from initial data to (codes for) their maximal developments, and the condition that a given candidate extension is a $C^0 \cap W^{1,2}_{\mathrm{loc}}$ extension of a given development, are Borel in the codes. Then the failure set is at most Σ^1_1 ("sigma one one", *analytic*) in the projective hierarchy of the data space, and its complement, the set of data whose development has no such extension, is at most Π^1_1 ("pi one one", coanalytic). The question of topological genericity is whether the failure set is *meager*: a countable union of nowhere dense sets. Kind: a stated fact of descriptive set theory, under the stated assumption. The argument: a datum lies in the failure set exactly when there exists an extension, coded by a point of a Polish space, that satisfies the extension condition; the condition quantifies over test functions, but a countable dense family of them suffices, so under the assumption the condition is Borel; a set of the form "there exists a point of a Polish space such that a Borel condition holds" is the projection of a Borel set, which is the definition of an analytic set; the complement of an analytic set is coanalytic; Σ^1_1 and Π^1_1 form the first level of the projective hierarchy, and an analytic set need not be Borel.

Consequences. In ZFC alone (Zermelo-Fraenkel set theory with the axiom of choice, the standard axioms), every analytic set has the property of Baire (Lusin and Sierpiński) and the perfect-set property (Suslin). So the failure set is either meager or comeager in some non-empty open set, and which of the two holds is well posed in ZFC; the property of Baire gives this alternative and does not decide it. For an analytic failure set, the state-

ment "the failure set is meager" is a Π^1_1 statement about its code, so it is absolute between models of set theory (Mostowski), and no forcing changes its truth value; that does not show that ZFC proves it or refutes it. Determinacy is not needed to pose it. Analytic determinacy is not a theorem of ZFC: by the theorems of Martin (1970) and Harrington (1978) it is equivalent to the statement that every real has a sharp, so an argument that invokes analytic determinacy as a theorem of ZFC is invalid. The classification places the failure set in the hierarchy, and meagerness is a further property that the classification neither gives nor excludes; an argument that concludes "the failure set is meager" from its place in the hierarchy assumes its conclusion. What is proved, under the stated assumption, is the place of the failure set in the projective hierarchy and the alternative that the property of Baire gives. The statement "the failure set is meager" is the Christodoulou-Chruściel formulation with "generic" read as "outside a meager set", which the literature states as a conjecture.

Measure genericity. The data space is not a linear space, since initial data must satisfy the constraint equations. Choose a linear chart of it near the data of interest, for example a parametrization by free data, and model the uncertainty in the initial data as a Gaussian measure on that chart. The question of measure genericity is whether the failure set has measure zero. Kind: a reformulation; "the failure set has measure zero" is the Christodoulou-Chruściel formulation with "generic" read as "outside a set of measure zero", which the literature states as a conjecture. For $\Lambda > 0$ the question is not settled by any argument from small sets in some classes of data and backgrounds: for near-extremal Reissner-Nordström-de Sitter backgrounds with smooth, neutral, massless scalar data, the condition for violation in the linear analysis is $\beta > 1/2$, and the set of black-hole parameters on which it holds is open. An argument that a set has measure zero because it is defined by infinitely many conditions holding at once (a set of infinite codimension) fails when the set is defined by an open condition, so no argument from small sets in the space of black-hole parameters can show that violations have measure zero in that class. Whether the corresponding set of initial data has positive measure de-

pends on the map from data to the parameters of the final black hole, which is not settled. Claims of measure genericity for $\Lambda > 0$ must therefore be restricted to classes of data and backgrounds in which $\beta < 1/2$ is known, such as the Kerr-de Sitter case in the linear analysis of Dias, Eperon, Reall and Santos.

Prevalence. Hunt, Sauer and Yorke (1992) defined *prevalence*, a measure-theoretic, translation-invariant notion of "almost every" on infinite-dimensional vector spaces, by means of probe measures (a set is shy, the negligible case, when some compactly supported probe measure gives every translate of it measure zero). Prevalence and Baire-category genericity are different notions: a set can be prevalent and meager, or comeager and shy. Prevalence and Gaussian measure zero are both measure-theoretic, and neither is equivalent to the topological notion.

Density. Let the *refinement order*, written $\mathscr{P}$ (script P), be the partial order of finite-precision approximations to initial data, ordered by refinement. The question of density is whether the set of approximations all of whose refinements force blow-up in $W^{1,2}_{\mathrm{loc}}$ at the Cauchy horizon is dense in the refinement order. A finite-precision approximation determines a basic open set of the data space, and an approximation forces blow-up exactly when its open set is disjoint from the failure set; so the density statement is the statement that the failure set is *nowhere dense* in the data space. It implies that the failure set is meager, and the converse fails. The density statement is a strong form of the conjecture, with "generic" read as "outside a nowhere dense set". Kind: a reformulation; the relation to the topological statement is a stated fact of topology. To be dense is to meet a refinement of every condition, and a condition that forces blow-up is an approximation every completion of which triggers the instability; so density says that every approximation has a refinement that triggers the instability. An argument for density that asserts "every non-trivial perturbation of the Kerr data triggers the instability" assumes the statement it

sets out to prove: as a reformulation, the density statement says what the conjecture asserts in the language of forcing; as a proof, such an argument is circular.

The literature's conjecture

The three reformulations ask, each with its own notion of a negligible set, whether mass inflation occurs for generic initial data at the regularity of $C^0 \cap W^{1,2}_{\mathrm{loc}}$. The literature states as a conjecture that it does:

The conjecture (*the Christodoulou-Chruściel formulation, with mass inflation*). For generic admissible initial data in the relevant classes (rotating vacuum; the spherically symmetric Einstein-Maxwell-real scalar field system), the quasi-local mass diverges at the Cauchy horizon, and the maximal globally hyperbolic development admits no extension across it with locally square-integrable Christoffel symbols (no extension with metric in $C^0 \cap W^{1,2}_{\mathrm{loc}}$). Kind: a conjecture of mathematical relativity, as the literature states it; no argument of this book uses it as a premise.

Status. For the spherically symmetric Einstein-Maxwell-real scalar field system the statement is established at the C^2 level (Luk and Oh 2019, parts I and II), and the mechanism is established: for data settling to Reissner-Nordström, the Misner-Sharp mass blows up along the Cauchy horizon while the metric and the area radius extend continuously, and no C^1 extension exists (Dafermos 2005). For the rotating vacuum case, continuous extension across a piece of the Cauchy horizon is a theorem under its interior hypotheses (Dafermos and Luk 2025), and failure of locally Lipschitz extension is proved under stated assumptions on the data (Gurriaran 2026; Luk and Sbierski 2026; with Sbierski 2026 for the step from curvature blow-up). In the rotating vacuum case, failure of extension with metric in $C^0 \cap W^{1,2}_{\mathrm{loc}}$ is the Christodoulou-Chruściel formulation itself, which the literature states as a conjecture.

The conjecture is a statement about the generic behavior of solutions of the Einstein equations in black-hole interiors, and its methods (non-linear analysis of hyperbolic equations, the spectral theory of black-hole perturb-

ations, the geometry of null hypersurfaces) are those of mathematical relativity. No independence phenomenon is known for it: no result shows, or suggests, that its statement in any of the three formulations is undecidable in the formal systems in which mathematical relativity is done. Whatever extension a spacetime has beyond a Cauchy horizon is further geometric time, so the argument of this book holds however the conjecture is settled, and at whatever regularity.

For Chapter 16: Why Time Travel Names Nothing

The Past Hypothesis

Chapter 16 states that the structural arrow of time is the root of the direction of time, that the entropic arrow is how that direction shows itself at the scale of eggs and engines, and that why the early universe had low entropy is a question about the start, which the structural arrow leaves to the account of the start. The Basis of Physics, Essay 11, Section IV, says what the structural reading makes of that beginning, in two sentences: “The Past Hypothesis, on this reading, is not the source of the arrow but a statement about initial conditions”, and “The argument does not say why that initial condition was what it was”. The Past Hypothesis is the statement that the universe began in a state of low entropy.

For Chapter 17: Two Theories That Cannot Be One

The spectral geometry of the product

Chapter 17 states in words how the spectral geometry of Ali Chamseddine and Alain Connes (1997) writes the two theories together. In the terms of my book Non-Commutative Geometry and the Immanent Metaphysics, Chapters 13 and 15, the arena and the content are two factors of one product, “the two modes composed side by side”; “The conjunction is me-

diated by scope"; and it is "not a single geometry perfecting both characters". "The Dirac operator of the product acts on each factor in its own character": the Dirac operator is the central operator of the product, whose spectrum, the set of its frequencies, describes the geometry, and its acting on each factor in that factor's own character is the chapter's claim that what relates the two theories is the act. The single step that imposes one common measure across the factors is the taking of the trace of the spectral action over the whole, the spectral action being the quantity, built from the Dirac operator, from which the theory's relations are drawn. The taking of the trace is the step at which the theory's relations at the cutoff, the theory's highest energy, part from two measured values. The two measured values are the weak mixing angle and the mass of the Higgs particle. Incommensurable means having no common measure.

Reading the programs

Chapter 17 reads each program of quantum gravity by the trade it makes between the two survivors of the theorem, a reading that does not depend on how well the programs succeed technically, and which I offer as a reading. The Basis of Physics states it with a caveat for the physicist: anyone working inside one of the programs may fairly say it is a coarse one.

What the sorting of the three accelerations could have shown

Chapter 17 matches the three kinds of acceleration with the three modalities: the fall with the transcendent, the push with the immanent, and the rotation with the omniscient, and it says that, to my mind, the reading is confirmed by how little has to be forced. The Basis of Physics gives, in the next sentences, what the sorting risked: "If the three had sorted the other way, if the push had turned out to be the one fixed by the world or the fall the one that radiates, I would have taken it as evidence against the alignment. They did not". In plain words: the alignment is the matching of the three kinds with the three modalities, and the fall "that radiates" would be a fall that spreads outward from its source, as the rotation does.

For Chapter 18: The Three Horizons

The kinds of claim in Chapter 18

The claims of Chapter 18 are of three kinds, the three the Preface names, and each has its own standing. The three kinds of knowledge, the horizons as classes rather than instances, and the Law of Three sit at the depth of the metaphysics itself, the ground, as do the claims that creation is transcendent, that the small edge is transcendent, and that the second axiom orders the small, the interacting and the structural. The reading of the speed of light and Planck's constant as the two edges of the domain of comparison, of entanglement as domain curvature, of quantum indeterminacy as creation entering through the small edge, of everyday order as the crowd the second axiom requires, and of the uncertainty relation as a case of the first forbidden pairing, are readings of known results, each giving the structural reason the result holds. The claim that a point singularity, perfect asymmetry with perfect discontinuity, cannot be realized in nature as described is a structural exclusion.

The standing of the reading of clarity and creation

Chapter 18 reads the link between clarity and creation into physics: the simple, closed system is the one physics predicts most exactly and the one in which creation shows least, and the complex, open system is the one physics predicts least exactly and the one in which the new shows most. It is a reading at the level of interpretation, the second kind of claim the Preface names, and it changes no prediction of physics. What the framework adds is the reason, that clarity and creation vary together, and I offer that reason as an interpretation, and not as a result physics is obliged to accept.

For Chapter 19: The Two Constants That Bound Every Domain

The mechanics of Lagrange and of Hamilton, read in the six measures

The mechanics of Lagrange rewrites the same physics as Newton's around the principle of stationary action: the path a body takes is the one for which a certain total, the action, does not change when the path is altered very slightly. For a body moving in a field, it works with the difference between the energy of motion and the stored energy of position. In the terms of the six measures, the energy of motion is written through momentum, force and time, and the stored energy is a relation of pattern in space across possibility; neither is a single measure. The mechanics of Hamilton rewrites it again, around position and momentum taken as linked pairs. The geometry of those pairs, which its allowed changes of description keep the same, belongs, with the sameness of Newton's laws for every observer in steady motion, to the symmetric side of the theory. The three forms of mechanics, Newton's, Lagrange's and Hamilton's, give the same predictions wherever all three apply. Forms that give the same predictions can be swapped for one another, and the third axiom requires the members of a triple not to be swappable, so this book does not read the three forms as a triple of the modalities.

The Compton relation and the shorthand of the first axiom

Chapter 19 writes the definition of a body's Compton wavelength as one product: the body's mass, times the speed of light, times its Compton wavelength, equals Planck's constant. The relation is a definition. It is not a case of the shorthand by which the first axiom is sometimes written, "the immanent equals the omniscient times the transcendent", since that shorthand states a pairing of roles, the immanent as the relation that joins the other two, and multiplies no amounts.

The kinds of claim in Chapter 19

The claims of Chapter 19 come in two kinds. Almost all of them belong to the metaphysics itself, the ground: what a domain is, that it has two limits of two kinds, the six measures and their grid, the rule of symmetry over the content row and of continuity over the context row, quantized content at a domain's smallest edge and continuous context at its largest, the three scales, the triples, and the floor at zero of interaction, existence and creation. The readings of the speed of light and Planck's constant as the physical cases of the two limits, of mass and energy as one currency related through light, of conservation and uncertainty as sum and product at the two edges, and of the two great theories by column, are readings of known physics, each giving the structural reason a known result has its form. The grid is a structure of kinds, and it says nothing about amounts.

For Chapter 20: What Physics Cannot Reach

Two placements of the correspondence of a theory

Chapter 20 places the correspondence of a theory, taken as the practice of comparing in the world, the acts in which a model is held up against what it is about, with the immanent. In *Category Theory and the Immanent Metaphysics* I place the correspondence of a theory with its object, taken as a relation laid out in full, with the transcendent; the two placements are placements of two different things.

The kinds of claim in Chapter 20

The three parts of a theory, the conjugacy of world and model, the reversal of physics and metaphysics, the placing of falsifiability by levels, and the theorem applied to itself are claims at the depth of the metaphysics itself, the ground. The reading of gauge freedom and coordinate free-

dom as structure that refers to nothing is a reading of a known feature of physical theories. In Chapter 20 I add no exclusion of my own to the exclusions that I made in the earlier chapters.

For Chapter 21: Computation and the Continuum

Five points that Chapter 21 states in words are given here for the specialist. Each result here is stated in words in that chapter.

The bounds on information in a region

A second objection to the continuity of the arena runs the same way as the objection from the Planck length. Jacob Bekenstein showed in 1981 that a bounded region of space can hold only a finite amount of information. The holographic principle, which Gerard 't Hooft (1993) and Leonard Susskind (1995) later built on his result, says the same, in a stronger form: the information a region can hold is bounded by the area of its boundary, as a hologram holds a three-dimensional scene on a flat plate. The holographic principle is a conjecture of the literature. Those bounds limit what any comparison within the domain can tell apart, which is the domain's own limit, and not a graininess of what exists. The answer to them is the answer to the Planck objection, and the step they invite, from "nothing more can be told apart here" to "what exists here is made of pieces", is the same step the Principle of Identity does not permit.

Endless collections

The account of this book does not treat endless collections as finished objects of different sizes; what it takes from Cantor is his fact about listing: no list with one line for each member of a collection contains every way of marking the members yes or no. That fact holds for finite and endless collections alike, and it is all the argument of Chapter 21 needs.

Following a sensitive system: the order of the two choices

Two facts sit together in the amplifier of Chapter 21, and both hold. Name a stretch of time and a closeness first, and some precision keeps a simulation that close to the system for the whole stretch; each extra digit of precision buys a further stretch of faithful following. Name a precision first, and some stretch of time exists over which the simulation and the system end up as far apart as the system can go. The two facts choose their numbers in a different order, and that is all that separates them. No fixed precision follows a sensitive system forever.

Reach, not speed

The result that a simulation of space and time, on any computer, ordinary or quantum, comes as close as one wishes and no stage of it contains them concerns how far a simulation can reach. It says nothing about speed or cost: it does not say that any computer is faster or slower, or put any problem into or out of any class of difficulty. How many steps a given closeness takes is a question about particular programs.

How near present measurement comes

Each proposed discrete rule beneath the smooth quantum wave fixes a precision beyond which its statistics would differ from the world's, and for any one fixed program, beyond its smallest step, a sensitive enough measurement eventually tells the program's statistics from the world's. Precision measurement and quantum information science (the study of how quantum systems store and process information) are approaching the range where such statistics can be tested very finely.

For Chapter 22: Computation and the Mind

Three points that Chapter 22 states in words or sets beside its argument are given here for the specialist.

One removal would be enough

Chapter 22 states that computation removes all three members of the bound triple, hard randomness, directed time and subjectivity, as a single act. The argument would go through if computation removed only one of the three, since by the binding of the triple the other two would go with it. That computation removes all three directly is more than the conclusion needs. Computation does in fact remove all three members of the bound triple directly, and that fact shows how far apart the two ways of describing stand.

Orchestrated objective reduction

The argument of Chapter 22 is a different one from the hypothesis of orchestrated objective reduction, put forward by Stuart Hameroff and Roger Penrose in 1996. That hypothesis places consciousness in quantum coherence (a quantum state kept whole, undisturbed by its surroundings) sustained in one structure of the brain's cells, the microtubules. It meets the standard objection, made by Max Tegmark in 2000, that such coherence is lost far too quickly in warm brain tissue to bear on what nerve cells do. The argument of Chapter 22 depends on no particular biological mechanism and needs no sustained entanglement in the brain, so that objection does not reach it.

Proofs and programs

Chapter 22 states that computation does not reduce to logic: a proof is valid whatever the order in which its steps are checked, while a computation must be carried out in order, each step depending on the last. The correspondence of proofs with programs, which William Howard set out in 1980, is a correspondence of structure, and it leaves that difference standing.

For Chapter 22: The Information Bound

Chapter 22 states in words a bound on what a computed world can contain. Here it is written as a formula, with the same conditions. The four names are those of the chapter: the rule, a computational recipe for rewriting a pattern; the starting condition, the pattern the rule begins from; the computed result, the pattern the rule produces from the starting condition; and an observation, any complex thing found in the world.

Write information(thing), read aloud as "the information of the thing", for the algorithmic information content of a thing: the length of the shortest description of it, in a fixed description language (the measure of Kolmogorov 1965, reached independently by Solomonoff and by Chaitin in the same years). The number of steps the rule is run is counted as part of the starting condition. Then, insofar as the computed result is the deterministic outcome of the rule and the starting condition alone, with nothing from outside added, and insofar as the notion of a theory is well defined,

$$\text{information}(\text{computed result}) \leq \text{information}(\text{rule}) + \text{information}(\text{starting condition})$$

read aloud as: the information of the computed result is at most the information of the rule plus the information of the starting condition. The inequality holds up to a constant fixed by the choice of description language.

The bound is the conservation of information under deduction, which is inherent in the idea of mathematics itself: were it otherwise, a valid inference would produce from its premises something its premises did not contain, and the notion of a valid inference would be inconsistent. Where an observation is close enough to some computed result, the observation's apparent information is compressed into the rule and the starting condition. A rule that is the shared basis for many such computed results is called a law of nature, the best available compression of everything observed. A theory of everything is the claim that a single rule, with varying starting conditions, compresses every observation there is.

The bound fixes where the information in any purely omniscient theory comes from: a theory can move information from the observations into the rule, and it cannot produce information from nothing. So every such theory takes some information as given and wholly incompressible, and however many rounds of theory building are attempted, the remainder is an incompressible residue of hard randomness, something not predictable on the basis of any theory. Compression can never reduce all observable information to zero bits. Chapter 22 states each of these results in words, and the Incommensuration Theorem gives the same exclusion from the side of comparison: it forbids perfected determinism (symmetry perfected with continuity) and pure randomness (asymmetry perfected with discontinuity), and what every actual world has between the two poles is causation together with that residual.

For the Afterword: The Books of Derivation

The Afterword says that this book leans on two things argued in full in the books of derivation. The Afterword names the first of the two things, the match between the six parts of comparison and the six measures, in the Afterword's own text. The second of the two things is gravity as the context of physics. The books of derivation reach gravity as the context of physics by three routes that do not pass through the grid of the six measures at all. The first two routes are the four marks that any context must carry, in *Measurement; The Steps from Comparison to Physics*, and the law of continuity for everything that is of context, in *Incommensuration; The Theorem and the Three Axioms*. The third route is the part of the combined geometry that belongs to the arena, in *Non-Commutative Geometry and the Immanent Metaphysics*. The form of the physical world, its algebra (in non-commutative geometry, the structure of quantities from which the geometry and its forces are built), its symmetries and its charges, is derived in that last volume; the values remain values of this universe.

The steps from comparison to the measures of physics, taken in *Measurement*, rest on the ground of comparison, on five premises, two taken from established physics and three identifications stated as premises (numbered P15 to P19 in that book), on four named principles stated with their reasons, and on further identifications, made in the first person beside the observations that bear on them.

Terms

Each term is given in the sense this book uses, in the words of the chapter where it is first explained; the chapter number follows in brackets.

Absolute elsewhere. Every event outside both light cones of a given event: the events that no cause joins to it, in either direction. [9, 12]

Absolutely applied. A character is applied absolutely when it holds exactly, at the foundation of a description, with no exception at any scale. [3]

Action. A quantity made by multiplying an energy by a duration, or a momentum by a distance; Planck's constant is an amount of action. [1, 18]

Arena. Space and time together: the setting of every measurement, in which every length and every duration is measured. In general relativity, the shape of the arena is gravity. [14, 17]

Arrow of time. The one-way running of time from past to future. The structural arrow is the arrow found in every single comparison; the entropic arrow is the direction in which entropy increases. [6, 13]

Asymmetry. Difference of content across difference of context: change the setting, and the content changes with it. [2]

Axiom. A basic statement from which other statements are reasoned; the three axioms of this book are each drawn from the Incommensuration Theorem. [4]

Bell's theorem. No account in which every outcome is fixed by hidden local causes can match all the predictions of quantum theory, and the measurements confirm those predictions. [10]

Binary hypothesis. The proposal that all the information in the universe can be written as ones and zeros; the strong form of the simulation hypothesis. [21]

Bit. The smallest amount of information, the answer to one yes-or-no question; a one or a zero. [21, 22]

Black hole. A region of space and time from inside which nothing, not even light, can get out; its outer surface is the event horizon. [14, 15]

Block universe. The view that past, present and future are equally real, that time is a fixed dimension rather than a directed process, and that the felt direction of time belongs to us and not to the world. [13]

Blue-shift. The squeezing of light and other radiation into shorter and shorter waves, which moves its color toward blue and raises its energy; for a body falling toward the inner horizon of a black hole it grows without limit. [15]

Bound triple. Subjectivity, directed time and hard randomness: three aspects of one act of comparison proceeding in a direction, each requiring the next, so that none is present without the other two. [8]

Catastrophic theorem of determinism. If any amount of indeterminism is evident or available anywhere, to any degree, then no scale is completely deterministic. It follows from the definition of determinism as the complete specification of every content for every context, at every scale. [9, 11]

C constant and H constant. The two constants that bound any domain, named after the letters physicists write for the speed of light and Planck's constant: the C constant at the domain's reach, the H constant at its grain. Either is a domain constant. [18, 19]

Category error. A question or a claim that applies a concept where the concept does not belong: an error of kind, as distinct from an error of fact, so that the request is malformed rather than hard. [6, 11]

Cauchy horizon. The edge of the predicted world: the boundary beyond which the starting data no longer fix the shape of space and time. [14]

Causality, cause. A consistency that an observer finds between two things that happen in the world, one following the other, again and again. Subjective causality is the foundation; objective causality is derived from it. [6]

Causation. The one-way, unbroken order at each event that physics finds at every scale: continuity with asymmetry. [6]

Choice. A process with three parts: what is possible, a selection, and a consequence. A choice brings about something that was open but not yet laid out; a decision selects from a set of possibilities already defined. [7]

Chronology horizon. The Cauchy horizon of a time machine: the edge of the region still settled by the world as it was before the machine was made. [14, 15]

Class and instance. A class is a kind, taken as a whole; an instance is one particular case of a kind. [4, 18]

Closed timelike curve, time loop. A path through space and time that a body with mass could follow and arrive again at the very event from which it set out. [15]

Comparer. Whatever makes a comparison; the word names that role, and in this book it never means a mind. [1]

Comparison. The act of finding whether two things are the same or different, in some respect. [1]

Compounds. The four patterns continuity, discontinuity, symmetry and asymmetry, each made of two findings at once, one about the content and one about the context. [2]

Computation. A procedure that goes, step by step and by fixed rules, from one state to the next, where each state is a finite string of symbols. A program is the written set of those rules. [21]

Computationalism. The thesis that all of physical reality is, in principle, computational: that every process can be represented as discrete operations on discrete states. [22]

Conjugate. Two things are conjugate when each can be perfected only at the cost of the other, and neither comes before the other. [3]

Consciousness. The bound triple taken all at once, in a context, as one act happening. [22]

Conservation. Staying the same in amount through every change; every conservation law is a law of symmetry. [2]

Content. What is being compared, what the comparison is about. [1]

Context. The respect in which the two things compared are two, such as two places or two moments: the setting that makes one finding distinct from another. [1]

Continuity. Sameness of content within sameness of context: hold the setting fixed, and the content does not jump. [2]

Continuum. The smooth whole of every possible value, with no gaps anywhere in it, in which between any two values there is always another. No list reaches it. [2, 21]

Correlation. A pattern of agreement between two records, found when the two records are set side by side. [8, 10]

Cosmological event horizon. The edge beyond which light sent out now will never reach us. [8, 18, 19]

Creation. The side of every event by which something new comes to be: choice's counterpart at the infinite scale, generating the range of what is open. It is real, and it does not exist. [7, 18]

Cross-modal reach. Each great theory applies the tool proper to one modality to the domain proper to the other; its reach into that domain is dense and never exhaustive. [2, 17]

Deep why and shallow why. The shallow why asks for the pattern of which this is a case, and physics answers it with a law. The deep why asks for the possibility from which a whole world is chosen, which no measurement can return. [20]

Defining constant. A constant whose value is fixed exactly, by definition. Since 2019 the International System of Units has fixed each of its units by the exact value of one of seven defining constants. [19]

Determinism. The claim that everything is fully and exactly fixed: every value of every content for every context, with no observer and no time needed. [5]

Discontinuity. Difference of content within sameness of context: hold the setting fixed, and the content changes anyway, a jump, a break, a threshold crossed. [2]

Domain, domain of comparison. Everything that a given kind of comparison can compare: the range within which comparisons can be made. Its inner edge, the microscopic boundary, is marked by Planck's constant; its outer edge, the macroscopic boundary, by the speed of light. [1, 12, 19]

Domain binding. Identifying a domain's own cases of symmetry and continuity and reading the theorem's limit in that domain's language; a different thing from matching ideas to the three modal roles. [7]

Domain curvature. The property by which no domain can be fully described place by place: features of the relations among points, and of no point. [18]

Entangled. Two particles whose outcomes, measured far apart, agree in a pattern that no separate instructions carried by each particle could give them. [9, 10]

Entropy. A count of how many ways a thing's small parts could be arranged while the thing looks just the same from outside. [13]

Epistemic. Concerning knowing: how anything comes to be known, checked and held. Epistemology is the study of knowing. [1, 3]

Epistemic Sandwich. An argument built on the free will theorem of Conway and Kochen, running across three layers of physical description at three scales, with our own scale as the filling. The choice it leaves is between a real arrow of time at every layer of description and no knowledge at all. [13]

Errors of extrema, the four errors. The four errors that arise when the relationship between self and world is taken to an extreme: perception taken as absolutely independent (realism) or absolutely dependent (idealism), and expression taken as absolutely independent (prejudice) or absolutely dependent (reaction). All four fall for one reason: every action has at least some reaction. [6]

Event. A single happening at one place and one moment. [1, 12]

Exist, real, objective. To exist is to have a place in a pattern, to stand as structure, the same for every comparer. To be real is to be encountered, here and now, in an actual act of comparison. To be objective is to hold across perspectives; objectivity comes in degrees, as the agreement of many findings across many comparers. [4]

Faithful reversal and lawful reversal. A lawful reversal runs a process back as the laws allow. A faithful reversal would pass back through exactly the states the original passed through and arrive at exactly the condition it started from. [13]

Falsifiable. A claim is falsifiable when it forbids something: when it says that a specified observation will not occur, so that the observation, if it occurred, would show the claim false. The test belongs to claims of why and how. [20]

Finding. What one act of comparison yields: one outcome, laid down as a record. [1]

First person, third person. To act in the first person is to act at one's own scale; to observe in the third person is to observe a domain from a scale other than its own. The two are also the two positions of description, practice and theory, and the uses meet, since observation from another scale is description. [9]

Focus of time. The frame of argument in which time is the thing being considered, with all three axioms of Chapter 4 held together. [14, 17]

Forbidden pairing. Perfect symmetry with perfect continuity, or perfect asymmetry with perfect discontinuity, together in one thing; the books of derivation call each a forbidden conjunction. [2, 3]

Formalism. The mathematics in which a theory is written, with the rules for working it. [1, 16]

Four theorems of being. Four propositions that follow from the universal law and the Incommensuration Theorem: two about the self, the subjective side, and two about the world, the objective side; two about content and two about context. The first is that subjective context is continuous. [2, 7]

Free will theorem. The theorem of John Conway and Simon Kochen (2006, strengthened 2009): if experimenters are free, then, given that quantum theory and relativity apply and there is no backward causation, the responses of particles are not fixed by anything that came before them. [9]

Gauge. Structure in a model that refers to nothing in the world, such as the free choice of potential and of overall phase. [20]

General relativity. The theory of gravity as the shape of space and time. [1, 2]

Geometric time, event time, evolution time. Geometric time is time as part of the shape of space and time, the time clocks measure along their paths. Event time is the order of findings for a comparer. Evolution time is the smooth dial along which the equation of quantum theory carries a state forward. [13, 14]

Great Mystery. The place where the unknown and the unknowable converge. [9, 18]

Hard randomness. An outcome that no description of what came before fixes, however complete that description is, for any comparer whatsoever, even one with unlimited access to the state of things: the residual the theorem guarantees, and one event with choice, described from outside. It is real, and kept apart from pure randomness. [6, 7, 8]

Hidden variables. Values held to fix every outcome in advance, taken to be out of reach of any measurement. [4, 10]

Horizon. An edge of what can be known from here, which moves back as it is approached, so that any one horizon can be passed while the kind of edge is always there. [18]

Immanent, omniscient, transcendent. The three modalities: the names of the three basic roles that every domain needs, each assigned to a term only relative to the triple in which it is considered. The immanent is the mode of process: the relational, the here and now, what is happening between things. The omniscient is the mode of structure: what holds the same for everyone who looks. The transcendent is the mode of what makes things possible: the evaluative, the creative, the potential. [2, 4]

Incommensuration Theorem. Symmetry and continuity cannot both be applied, at once and absolutely, to one comparison, and neither can asymmetry and discontinuity. It is a structural fact about the nature of knowing, inherent in every act of knowing, and the deepest and most powerful epistemic basis possible. It was first stated in the Technical Appendix of *An Immanent Metaphysics*. [3]

Incompressible residue. What is left after every regularity has been taken out, and what no further taking out can shrink. [8, 22]

Information deficit. What the share of the world that an observer can reach leaves out; The Gravitational Passage calls it "the price of being inside". [14, 16]

Infra-physics. The layer beneath every physical theory: the layer of structural necessity on which the facts of physics rest. [20]

Invariant. A content that stays unchanged when its setting changes; an invariant quantity is one that every comparer finds the same. [2, 12]

Light cone. At each event, the set of every event that can be joined to it by cause: the forward cone holds every event it can affect, the backward cone every event that can affect it. [12]

Listable. A collection that can be set out in order, first, second, third, without leaving anything out. Everything a computation could ever produce is listable. [21]

Locality. In the broad sense, every agreement between what is found in two places is carried by causes that travel from place to neighboring place. In the narrow sense physics keeps, no signal crosses a distance in no time. [10]

Macroscopic, mesoscopic, microscopic. The scale of the very large, the cosmos; the middle scale, the scale of the comparer and its instruments; and the scale of the very small, particles and fields. In this book "mesoscopic" always names the observer's own scale. [1, 13, 17]

Mass inflation. The effect, found by Poisson and Israel (1990) and by Ori (1991), by which radiation falling into a charged or spinning black hole acts back on the shape of its interior, so that the enclosed mass grows without limit at the Cauchy horizon. [15]

Meager. Made of countably many nowhere dense pieces: thin in the sense of position. [15]

Measurement. A comparison made from inside the world: a count, a reading or a frequency set against a standard. [1]

Metaphysics. The study of the relations between domains, as physics is the study of the interactions within one. [20, 21]

Microstate amplification. The process by which small differences at a fine-grained level of a system are carried up and magnified into differences at the level that matters for the system's behavior. [9]

Modal boundary. In The Gravitational Passage, the Planck scale, where the continuously asymmetric and the discontinuously symmetric descriptions exchange dominance. [14]

Modal rotation. The passage from one position's description of an event to another's: the event stays as it is, and the position from which it is described changes. A role belongs to a position of description, and never to a thing. [7, 8]

Model, world, correspondence. The three parts of every theory: the object domain, the world the theory is about; the model, the symbols and relations in which it is described; and the correspondence, the links made between the two by acts of comparison. [20]

Momentum. The quantity of motion a moving body carries, larger for a heavier body and larger for a faster one. [2, 3]

Names no concept. A phrase names no concept when every account of what would satisfy it either describes something other than what the phrase is used to mean, or asks for a set of conditions, all at once, that nothing satisfies. [16]

Non-polarization principle. All worlds are causal to some degree above zero. [6]

Nowhere dense. Thick nowhere: however small a patch is chosen, there is a smaller patch inside it that the set misses. [15]

Object domain. What a theory is about, the first of the three parts of every theory; in physics, the physical world itself, always richer than any description of it. [20]

Past Hypothesis. The assumption that the universe began in a state of extraordinarily low entropy, which the account of the arrow of time from entropy needs. [13, 16]

Phonon. One step of energy in a vibration of a crystal lattice: the smallest packet of vibration in a crystal, a quantum of the description of sound, and not a separate thing in the crystal. [17]

Planck's constant. The quantum of action: the smallest package of process, an energy multiplied by a time. By itself it sets no smallest length and no shortest time; how fine a comparison can be made depends on the energy brought to it. [12, 18, 19]

Planck length, Planck scale. The scale at which a body's gravitational size and its quantum size are equal, where the description of the setting and the description of what is found change places as to which one dominates. It is a bound on findings, not a least step of space and time. [14, 18]

Predicted world. The largest region that the starting conditions of a world settle, with nothing left to choose (physicists say the maximal globally hyperbolic development). [14, 15]

Principle of Identity. What cannot be told apart, by any test that could even be imagined, must be counted as the same. [21]

Pseudo-randomness. Numbers made by a fixed recipe that only look random; a pseudo-random generator is fully fixed by its seed. [21, 22]

Pure randomness. A content found different whatever its context: complete indeterminism, with nothing at all fixed. It is an idealization, and kept apart from hard randomness. [5, 6]

Quantization. The coming of things in whole units that cannot be split. [1, 2]

Quantum theory. The theory of matter and light, whose findings come in countable steps. [1, 2]

Residual. What no lawful description fixes; the theorem guarantees a residual in every description. [6, 7, 8]

Root Tautology. Interaction, comparison and relation are one act, named three ways. [1]

Signal. Anything that passes from one event to another and lets a record at the second depend on an act at the first. [1, 12]

Simulation, simulation hypothesis. A simulation is a computation built to copy the behavior of some physical thing. The simulation hypothesis is the proposal that our world is such a computation, running in a more basic world. [21]

Singularity. A place where the description of space and time breaks down, past which no continuation of a stated smoothness exists. [14, 15]

Six intrinsics. Sameness and difference, content and context, the subjective and the objective: the six parts that belong to every comparison by its nature. [1]

Six measures. Mass, space, force, time, probability and possibility: what a physical measurement measures, three contents each measured in its own context. [1, 19]

Soft randomness. Unpredictability that comes from a limit of what we know, and that better knowledge could shrink: a statement about our knowledge. Compare hard randomness. [8]

Spacelike separated. Two events so far apart in space, and so close in time, that no signal traveling at the speed of light could pass between them. [9, 12]

Speed of light. At root, the outer edge of the domain of comparison itself, and not the speed of any one thing: the fastest that any influence can pass from one event to another. Light, which travels in empty space at exactly that speed, is its perfect example. [12]

Stance. The position from which a comparison is made: the subjective part of every comparison. For a kitchen scale, the scale itself, with its place, its setting and the way it stands. [1, 4, 8]

Strong cosmic censorship. Penrose's conjecture that the predicted world of typical starting conditions cannot be continued past its edge, at a stated smoothness. [14, 15]

Structural arrow. The direction of time that follows from the parts of comparison alone. It is present in a single comparison, and it needs no large number of particles, no coarse-graining, no concept of entropy and no Past Hypothesis. [13, 16]

Subjective, objective. The subjective is the stance from which a comparison is made; the objective is the side about which it is made, the properties of the things compared, which do not depend on who is comparing them. [1]

Superdeterminism. Determinism extended to the experimenters' choice of settings, so that the setting and the particle share one causal history. [6, 11]

Superposition. A combination of alternatives held together in one description, each with its chance. [4, 13]

Survivors. The two pairings the theorem allows when applied absolutely: continuity with asymmetry, smooth and going somewhere, and symmetry with discontinuity, the same everywhere and countable. [3]

Symmetry. Sameness of content across difference of context: change the setting, and the content is unchanged. [2]

Teleology. The coherence of a whole, which descriptions in terms of purposes and ends try to capture. [5, 7]

Test in principle. A test whose logic works whether or not any instrument can perform it. [21]

Time-binding. The arrow of time as lived, binding past to present to future in one unbroken, one-way thread. The word is Alfred Korzybski's (1921), where it means the passing of knowledge across generations; the use here is narrower. [22]

Time machine. In the physicists' sense, a region where paths that loop back in time appear in a world that had none before them; its edge is the chronology horizon. [14, 15]

Triple of the real. Change, causality and choice: the three members that everything real has. [6, 7]

Uncertainty relation. The rule that the two members of a conjugate pair, such as position and momentum, cannot both be fixed sharply at once. [1, 18]

Universal law. The principle of structure that says which kind of law governs each side of reality: symmetry law for content, continuity law for context. [2, 19]

Unknowable. What cannot be known, even in principle, by any observer, at any time, by any method, because comparison itself does not reach it. [18]

Valid, sound. An argument is valid when its conclusion follows from its starting points; it is sound when it is valid and its starting points are also true of what it is about. [5]

Zero-positive. A degree with a floor at zero, nothing below it, and no ceiling above. [1, 19]

Works Cited

Abdo, A. A., and others (Fermi LAT and Fermi GBM Collaborations) (2009). A Limit on the Variation of the Speed of Light Arising from Quantum Gravity Effects. *Nature* 462, 331–334.

Albert, D. Z. (2000). *Time and Chance*. Cambridge, Massachusetts: Harvard University Press.

Angelopoulos, A., and others (CPLEAR Collaboration) (1998). First Direct Observation of Time-Reversal Non-Invariance in the Neutral-Kaon System. *Physics Letters B* 444, 43–51.

Anscombe, G. E. M. (1971). *Causality and Determination*. Cambridge: Cambridge University Press.

Aspect, A., Dalibard, J., and Roger, G. (1982). Experimental Test of Bell's Inequalities Using Time-Varying Analyzers. *Physical Review Letters* 49, 1804–1807.

Aspect, A., Grangier, P., and Roger, G. (1982). Experimental Realization of Einstein-Podolsky-Rosen-Bohm Gedankenexperiment: A New Violation of Bell's Inequalities. *Physical Review Letters* 49, 91–94.

ATLAS Collaboration (2017). Evidence for Light-by-Light Scattering in Heavy-Ion Collisions with the ATLAS Detector at the LHC. *Nature Physics* 13, 852–858.

ATLAS Collaboration (2019). Observation of Light-by-Light Scattering in Ultraperipheral Pb+Pb Collisions with the ATLAS Detector. *Physical Review Letters* 123, 052001.

Austin, J. L. (1962). *Sense and Sensibilia*. Oxford: Oxford University Press.

Beggs, J. M., and Plenz, D. (2003). Neuronal Avalanches in Neocortical Circuits. *Journal of Neuroscience* 23, 11167–11177.

Bekenstein, J. D. (1973). Black Holes and Entropy. *Physical Review D* 7, 2333–2346.

Bekenstein, J. D. (1981). Universal Upper Bound on the Entropy-to-Energy Ratio for Bounded Systems. *Physical Review D* 23, 287–298.

Bell, J. S. (1964). On the Einstein Podolsky Rosen Paradox. *Physics* 1, 195–200.

Bell, J. S. (2004). *Speakable and Unspeakable in Quantum Mechanics*. Second edition. Cambridge: Cambridge University Press. (The paper "Free Variables and Local Causality".)

Bennett, C. H. (1973). Logical Reversibility of Computation. *IBM Journal of Research and Development* 17, 525–532.

BIG Bell Test Collaboration (2018). Challenging Local Realism with Human Choices. *Nature* 557, 212–216.

Bohm, D. (1952). A Suggested Interpretation of the Quantum Theory in Terms of "Hidden" Variables, I and II. *Physical Review* 85, 166–193.

Bohr, N. (1928). The Quantum Postulate and the Recent Development of Atomic Theory. *Nature* 121, 580–590.

Bohr, N. (1935). Can Quantum-Mechanical Description of Physical Reality Be Considered Complete? *Physical Review* 48, 696–702.

Born, M. (1926). Zur Quantenmechanik der Stoßvorgänge. *Zeitschrift für Physik* 37, 863–867.

Bostrom, N. (2003). Are You Living in a Computer Simulation? *Philosophical Quarterly* 53, 243–255.

Brewer, E. A. (2000). Towards Robust Distributed Systems. Keynote address, ACM Symposium on Principles of Distributed Computing, Portland, Oregon.

Bureau International des Poids et Mesures (2019). *The International System of Units* (*SI*), the SI Brochure. Sèvres: BIPM.

Butterfield, J., and Isham, C. J. (1999). On the Emergence of Time in Quantum Gravity. arXiv:gr-qc/9901024.

Callender, C. (2000). Is Time "Handed" in a Quantum World? *Proceedings of the Aristotelian Society* 100, 247–269.

Callender, C. (2021). Thermodynamic Asymmetry in Time. In E. N. Zalta (editor), *The Stanford Encyclopedia of Philosophy*, Fall 2021 edition.

Cantor, G. (1874). Über eine Eigenschaft des Inbegriffes aller reellen algebraischen Zahlen. *Journal für die reine und angewandte Mathematik* 77, 258–262.

Cantor, G. (1891). Über eine elementare Frage der Mannigfaltigkeitslehre. *Jahresbericht der Deutschen Mathematiker-Vereinigung* 1, 75–78.

Cardoso, V., Costa, J. L., Destounis, K., Hintz, P., and Jansen, A. (2018). Quasinormal Modes and Strong Cosmic Censorship. *Physical Review Letters* 120, 031103. arXiv:1711.10502.

Cardoso, V., Costa, J. L., Destounis, K., Hintz, P., and Jansen, A. (2018). Strong Cosmic Censorship in Charged Black-Hole Spacetimes: Still Subtle. *Physical Review D* 98, 104007.

Carr, H. Y., and Purcell, E. M. (1954). Effects of Diffusion on Free Precession in Nuclear Magnetic Resonance Experiments. *Physical Review* 94, 630.

Carter, N. (2009). *Visual Group Theory*. Washington, DC: Mathematical Association of America.

Chaitin, G. J. (1966). On the Length of Programs for Computing Finite Binary Sequences. *Journal of the ACM* 13, 547–569.

Chamseddine, A. H., and Connes, A. (1997). The Spectral Action Principle. *Communications in Mathematical Physics* 186, 731–750.

Choquet-Bruhat, Y., and Geroch, R. (1969). Global Aspects of the Cauchy Problem in General Relativity. *Communications in Mathematical Physics* 14, 329–335.

Christenson, J. H., Cronin, J. W., Fitch, V. L., and Turlay, R. (1964). Evidence for the 2π Decay of the K_2^0 Meson. *Physical Review Letters* 13, 138–140.

Christodoulou, D. (2009). *The Formation of Black Holes in General Relativity*. EMS Monographs in Mathematics. Zürich: European Mathematical Society. arXiv:0805.3880.

Chruściel, P. T. (1991). On Uniqueness in the Large of Solutions of Einstein's Equations ("Strong Cosmic Censorship"). *Proceedings of the Centre for Mathematical Analysis, Australian National University* 27.

Clauser, J. F., Horne, M. A., Shimony, A., and Holt, R. A. (1969). Proposed Experiment to Test Local Hidden-Variable Theories. *Physical Review Letters* 23, 880–884.

Conway, J. H., and Kochen, S. (2009). The Strong Free Will Theorem. *Notices of the American Mathematical Society* 56, 226–232.

Conway, J., and Kochen, S. (2006). The Free Will Theorem. *Foundations of Physics* 36, 1441–1473.

Dafermos, M. (2003). Stability and Instability of the Cauchy Horizon for the Spherically Symmetric Einstein–Maxwell–Scalar Field Equations. *Annals of Mathematics* 158, 875–928.

Dafermos, M. (2005). The Interior of Charged Black Holes and the Problem of Uniqueness in General Relativity. *Communications on Pure and Applied Mathematics* 58, 445–504. arXiv:gr-qc/0307013.

Dafermos, M., and Luk, J. (2025). The Interior of Dynamical Vacuum Black Holes I: The C^0-Stability of the Kerr Cauchy Horizon. *Annals of Mathematics* 202 (2), 309–630. arXiv:1710.01722.

Dafermos, M., and Shlapentokh-Rothman, Y. (2018). Rough Initial Data and the Strength of the Blue-Shift Instability on Cosmological Black Holes with $\Lambda > 0$. *Classical and Quantum Gravity* 35, 195010. arXiv: 1805.08764.

Davis, T. M., and Lineweaver, C. H. (2004). Expanding Confusion: Common Misconceptions of Cosmological Horizons and the Superluminal Expansion of the Universe. *Publications of the Astronomical Society of Australia* 21, 97–109.

Descartes, R. (1637). *Discours de la méthode*. Leiden: Ian Maire.

Descartes, R. (1641). *Meditationes de prima philosophia*. Paris: Michel Soly.

Dias, Ó. J. C., Eperon, F. C., Reall, H. S., and Santos, J. E. (2018). Strong Cosmic Censorship in de Sitter Space. *Physical Review D* 97, 104060. arXiv: 1801.09694.

Dias, Ó. J. C., Reall, H. S., and Santos, J. E. (2018). Strong Cosmic Censorship: Taking the Rough with the Smooth. *Journal of High Energy Physics* 2018 (10), 001.

Dias, Ó. J. C., Reall, H. S., and Santos, J. E. (2019). Strong Cosmic Censorship for Charged de Sitter Black Holes with a Charged Scalar Field. *Classical and Quantum Gravity* 36, 045005. arXiv:1808.04832.

Earman, J. (1986). *A Primer on Determinism*. Dordrecht: D. Reidel.

Earman, J. (2002). What Time Reversal Invariance Is and Why It Matters. *International Studies in the Philosophy of Science* 16, 245–264.

Earman, J., and Norton, J. (1987). What Price Spacetime Substantivalism? The Hole Story. *British Journal for the Philosophy of Science* 38, 515–525.

Einstein, A. (1905). Über einen die Erzeugung und Verwandlung des Lichtes betreffenden heuristischen Gesichtspunkt. *Annalen der Physik* 17, 132–148.

Einstein, A. (1905). Zur Elektrodynamik bewegter Körper. *Annalen der Physik* 17, 891–921.

Einstein, A. (1907). Über das Relativitätsprinzip und die aus demselben gezogenen Folgerungen. *Jahrbuch der Radioaktivität und Elektronik* 4, 411–462.

Einstein, A. (1916). Die Grundlage der allgemeinen Relativitätstheorie. *Annalen der Physik* 49, 769–822.

Einstein, A., Podolsky, B., and Rosen, N. (1935). Can Quantum-Mechanical Description of Physical Reality Be Considered Complete? *Physical Review* 47, 777–780.

Everett, H., III (1957). "Relative State" Formulation of Quantum Mechanics. *Reviews of Modern Physics* 29, 454.

Fan, X., Myers, T. G., Sukra, B. A. D., and Gabrielse, G. (2023). Measurement of the Electron Magnetic Moment. *Physical Review Letters* 130, 071801.

Friedman, J., Morris, M. S., Novikov, I. D., Echeverria, F., Klinkhammer, G., Thorne, K. S., and Yurtsever, U. (1990). Cauchy Problem in Spacetimes with Closed Timelike Curves. *Physical Review D* 42, 1915.

Ghirardi, G. C., Rimini, A., and Weber, T. (1980). A General Argument against Superluminal Transmission through the Quantum Mechanical Measurement Process. *Lettere al Nuovo Cimento* 27, 293–298.

Gilbert, S., and Lynch, N. (2002). Brewer's Conjecture and the Feasibility of Consistent, Available, Partition-Tolerant Web Services. *ACM SIGACT News* 33 (2), 51–59.

Giorgi, E., Klainerman, S., and Szeftel, J. (2024). Wave Equations Estimates and the Nonlinear Stability of Slowly Rotating Kerr Black Holes. *Pure and Applied Mathematics Quarterly* 20 (7), 2865–3849. arXiv:2205.14808.

Giustina, M., and others (2015). Significant-Loophole-Free Test of Bell's Theorem with Entangled Photons. *Physical Review Letters* 115, 250401.

Gödel, K. (1931). Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I. *Monatshefte für Mathematik und Physik* 38, 173–198.

Gödel, K. (1949). A Remark about the Relationship between Relativity Theory and Idealistic Philosophy. In P. A. Schilpp (editor), *Albert Einstein: Philosopher-Scientist*, Library of Living Philosophers, 557–562.

Gödel, K. (1949). An Example of a New Type of Cosmological Solutions of Einstein's Field Equations of Gravitation. *Reviews of Modern Physics* 21, 447.

Gosse, P. H. (1857). *Omphalos: An Attempt to Untie the Geological Knot*. London: John Van Voorst.

Gurriaran, S. (2026). Non-Linear Instability of the Kerr Cauchy Horizon near i^+. Preprint, arXiv:2603.17911.

Häfner, D., Hintz, P., and Vasy, A. (2025). Linear Stability of Kerr Black Holes in the Full Subextremal Range. Preprint, arXiv:2506.21183.

Hahn, E. L. (1950). Spin Echoes. *Physical Review* 80, 580.

Hameroff, S., and Penrose, R. (1996). Orchestrated Reduction of Quantum Coherence in Brain Microtubules: A Model for Consciousness. *Mathematics and Computers in Simulation* 40, 453–480.

Handsteiner, J., and others (2017). Cosmic Bell Test: Measurement Settings from Milky Way Stars. *Physical Review Letters* 118, 060401.

Harrington, L. (1978). Analytic Determinacy and $0^{\#}$. *Journal of Symbolic Logic* 43, 685–693.

Hawking, S. W. (1975). Particle Creation by Black Holes. *Communications in Mathematical Physics* 43, 199–220; erratum 46, 206 (1976).

Hawking, S. W. (1992). Chronology Protection Conjecture. *Physical Review D* 46, 603–611.

Hawking, S. W., and Ellis, G. F. R. (1973). *The Large Scale Structure of Space-Time*. Cambridge: Cambridge University Press.

Hawking, S., and Mlodinow, L. (2010). *The Grand Design*. New York: Bantam.

Heisenberg, W. (1927). Über den anschaulichen Inhalt der quantentheoretischen Kinematik und Mechanik. *Zeitschrift für Physik* 43, 172–198.

Hensen, B., and others (2015). Loophole-Free Bell Inequality Violation Using Electron Spins Separated by 1.3 Kilometres. *Nature* 526, 682–686.

Hintz, P., and Vasy, A. (2017). Analysis of Linear Waves near the Cauchy Horizon of Cosmological Black Holes. *Journal of Mathematical Physics* 58, 081509. arXiv:1512.08004.

Hod, S. (2019). Strong Cosmic Censorship in Charged Black-Hole Spacetimes: As Strong as Ever. *Nuclear Physics B* 941, 636. arXiv: 1801.07261.

Hod, S. (2020). A Proof of the Strong Cosmic Censorship Conjecture. *International Journal of Modern Physics D* 29, 2042003. arXiv:2012.01449.

Hoefer, C. (2016). Causal Determinism. In E. N. Zalta (editor), *The Stanford Encyclopedia of Philosophy*, Spring 2016 edition.

Hoffman, D. D. (2019). *The Case Against Reality: Why Evolution Hid the Truth from Our Eyes*. New York: W. W. Norton.

Hollands, S., Wald, R. M., and Zahn, J. (2020). Quantum Instability of the Cauchy Horizon in Reissner–Nordström–de Sitter Spacetime. *Classical and Quantum Gravity* 37, 115009.

Hossenfelder, S. (2020). Superdeterminism: A Guide for the Perplexed. Preprint.

Hossenfelder, S., and Palmer, T. (2020). Rethinking Superdeterminism. *Frontiers in Physics* 8, 139.

Howard, W. A. (1980). The Formulae-as-Types Notion of Construction. In J. P. Seldin and J. R. Hindley (editors), *To H. B. Curry: Essays on Combinatory Logic, Lambda Calculus and Formalism*, 479. London: Academic Press.

Hume, D. (1739). *A Treatise of Human Nature*. London: John Noon.

Hunt, B. R., Sauer, T., and Yorke, J. A. (1992). Prevalence: A Translation-Invariant "Almost Every" on Infinite-Dimensional Spaces. *Bulletin of the American Mathematical Society* 27, 217–238.

Isham, C. J. (1995). Structural Issues in Quantum Gravity. arXiv:gr-qc/9510063.

Jacobson, T. (1995). Thermodynamics of Spacetime: The Einstein Equation of State. *Physical Review Letters* 75, 1260–1263.

Janet, P. (1877). Une illusion d'optique interne. *Revue philosophique de la France et de l'étranger* 3, 497–502.

Kay, B. S., Radzikowski, M. J., and Wald, R. M. (1997). Quantum Field Theory on Spacetimes with a Compactly Generated Cauchy Horizon. *Communications in Mathematical Physics* 183, 533–556. arXiv:gr-qc/9603012.

Kennard, E. H. (1927). Zur Quantenmechanik einfacher Bewegungstypen. *Zeitschrift für Physik* 44, 326–352.

Klainerman, S., and Szeftel, J. (2023). Kerr Stability for Small Angular Momentum. *Pure and Applied Mathematics Quarterly* 19 (3), 791–1678. arXiv:2104.11857.

Kochen, S., and Specker, E. P. (1967). The Problem of Hidden Variables in Quantum Mechanics. *Journal of Mathematics and Mechanics* 17, 59–87.

Kolmogorov, A. N. (1965). Three Approaches to the Quantitative Definition of Information. *Problems of Information Transmission* 1, 1–7.

Korzybski, A. (1921). *Manhood of Humanity*. New York: E. P. Dutton.

Krauss, L. M. (2012). *A Universe from Nothing: Why There Is Something Rather than Nothing*. New York: Free Press.

Landauer, R. (1961). Irreversibility and Heat Generation in the Computing Process. *IBM Journal of Research and Development* 5, 183–191.

Landry, Forrest. *An Immanent Metaphysics*. First published 2001. Wytheville, Virginia: UVSM Publications. ISBN 979-8-9896628-4-5. Cited by aphorism number.

Landry, F. *Accessing the Real*. Wytheville, Virginia: UVSM Publications.

Landry, F. (2026). *The Basis of Physics; Nineteen Essays on the Structural Limits of Physical Law*. Wytheville, Virginia: UVSM Publications. ISBN 979-8-90710-072-5. DOI: https://doi.org/10.67291/WHOM6253

Landry, F. (2026). *Category Theory and the Immanent Metaphysics; A Formal Correspondence*. Wytheville, Virginia: UVSM Publications. ISBN 979-8-90710-071-8.

Landry, F. (2026). *Causality Without Determinism; Free Will, Measurement and Mind*. Wytheville, Virginia: UVSM Publications. ISBN 979-8-90710-075-6.

Landry, F. (2026). *Comparison; The Necessary Basis*. Wytheville, Virginia: UVSM Publications. ISBN 979-8-90710-062-6.

Landry, F. (2026). *The Gravitational Passage; Deriving General Relativity from the Immanent Metaphysics*. Wytheville, Virginia: UVSM Publications. ISBN 979-8-90710-056-5.

Landry, F. (2026). *Incommensuration; The Theorem and the Three Axioms*. Wytheville, Virginia: UVSM Publications. ISBN 979-8-90710-063-3.

Landry, F. (2026). *The Kinds of Particle*. Wytheville, Virginia: UVSM Publications. ISBN 979-8-90710-060-2.

Landry, F. (2026). *The Last Turtle; Essays from Comparison and the Incommensuration Theorem to Causation, the Arena and the Limits of Computation*. Wytheville, Virginia: UVSM Publications. ISBN 979-8-90710-215-6. DOI: https://doi.org/10.67291/ONZV8400

Landry, F. (2026). *Measurement; The Steps from Comparison to Physics*. Wytheville, Virginia: UVSM Publications. ISBN 979-8-90710-064-0. DOI: https://doi.org/10.67291/VZSP1393

Landry, F. (2026). *The Measures of the Kinds*. Wytheville, Virginia: UVSM Publications. ISBN 979-8-90710-116-6.

Landry, F. (2026). *Non-Commutative Geometry and the Immanent Metaphysics*. Wytheville, Virginia: UVSM Publications. ISBN 979-8-90710-100-5.

Landry, F. (2026). *Physics Is Not Math; Metaphysics, Physics, Computation and Mathematics, and Why Their Order Cannot Be Otherwise*. Wytheville, Virginia: UVSM Publications. ISBN 979-8-90710-070-1.

Landry, F. (2026). *The Quantum Bridge; Entanglement, Error Correction, and the Limits of Simulation*. Wytheville, Virginia: UVSM Publications. ISBN 979-8-90710-107-4. DOI: https://doi.org/10.67291/NHFX3629

Landry, F. (2026). *The Quantum Passage; Quantum Mechanics from the Structure of Comparison*. Wytheville, Virginia: UVSM Publications. ISBN 979-8-90710-073-2.

Landry, F. (2026). *The Six Cells of Mathematics; How the Incommensuration Theorem Maps Mathematical Practice*. Wytheville, Virginia: UVSM Publications.

Landry, F. (2026). *The Story of the Particles*. Wytheville, Virginia: UVSM Publications. ISBN 979-8-90710-053-4.

Landry, F. (2026). *The Strengths of the Kinds*. Wytheville, Virginia: UVSM Publications. ISBN 979-8-90710-115-9.

Landry, F. *The Structure of Effective Choice*. Wytheville, Virginia: UVSM Publications.

Landry, F. (2026). *Third Resolutions in Western Philosophy*. Wytheville, Virginia: UVSM Publications. (The essay "The Simulation Hypothesis Dissolved".)

Landry, F. (2026). *Two Theories, One World; Gravity, the Quantum and the Mistake Behind the Search for One Theory*. Wytheville, Virginia: UVSM Publications. ISBN 979-8-90710-068-8.

Landry, F. *Unlocking Reality*. Wytheville, Virginia: UVSM Publications.

Landry, F. (2026). *The Unquantized Arena; A Structural Diagnosis of Quantum Gravity and the Measurement Problem*. Wytheville, Virginia: UVSM Publications. ISBN 979-8-90710-067-1.

Landry, F. (2026). *Why Time Travel Is Not a Concept; A Meditation on the Three Times and the Cauchy Horizon*. Wytheville, Virginia: UVSM Publications. ISBN 979-8-90710-121-0. DOI: https://doi.org/10.67291/SNQF5135

Laplace, P.-S. (1814). *Essai philosophique sur les probabilités* (A Philosophical Essay on Probabilities). Paris: Veuve Courcier.

Lees, J. P., and others (BaBar Collaboration) (2012). Observation of Time-Reversal Violation in the B^0 Meson System. *Physical Review Letters* 109, 211801.

Leibniz, G. W. (1714). Principles of Nature and Grace, Based on Reason. In *Philosophical Essays*, translated by R. Ariew and D. Garber. Indianapolis: Hackett, 1989.

Leibniz, G. W., and Clarke, S. (1717). *A Collection of Papers, Which passed between the late Learned Mr. Leibnitz, and Dr. Clarke, In the Years 1715 and 1716*. London.

Lorenz, E. N. (1963). Deterministic Nonperiodic Flow. *Journal of the Atmospheric Sciences* 20, 130.

Loschmidt, J. (1876). Über den Zustand des Wärmegleichgewichtes eines Systems von Körpern mit Rücksicht auf die Schwerkraft. *Sitzungsberichte der Kaiserlichen Akademie der Wissenschaften, Wien, Mathematisch-Naturwissenschaftliche Classe* 73, 128–142.

Luk, J., and Oh, S.-J. (2019). Strong Cosmic Censorship in Spherical Symmetry for Two-Ended Asymptotically Flat Initial Data I. The Interior of the Black Hole Region. *Annals of Mathematics* 190 (1), 1–111. arXiv: 1702.05715.

Luk, J., and Oh, S.-J. (2019). Strong Cosmic Censorship in Spherical Symmetry for Two-Ended Asymptotically Flat Initial Data II. The Exterior of the Black Hole Region. *Annals of PDE* 5, 6. arXiv:1702.05716.

Luk, J., and Sbierski, J. (2016). Instability Results for the Wave Equation in the Interior of Kerr Black Holes. *Journal of Functional Analysis* 271 (7), 1948–1995. arXiv:1512.08259.

Luk, J., and Sbierski, J. (2026). The Formation of a Weak Null Singularity in the Interior of Generic Rotating Black Holes. Preprint, arXiv: 2604.04877.

Mach, E. (1883). *Die Mechanik in ihrer Entwickelung, historisch-kritisch dargestellt*. Leipzig: Brockhaus.

Mach, E. (1893). *The Science of Mechanics: A Critical and Historical Account of Its Development*. Translated by T. J. McCormack. Chicago: Open Court.

Malament, D. B. (1977). The Class of Continuous Timelike Curves Determines the Topology of Spacetime. *Journal of Mathematical Physics* 18, 1399–1404.

Martin, D. A. (1970). Measurable Cardinals and Analytic Games. *Fundamenta Mathematicae* 66, 287–291.

Maudlin, T. (2007). *The Metaphysics Within Physics*. Oxford: Oxford University Press.

Maxwell, J. C. (1865). A Dynamical Theory of the Electromagnetic Field. *Philosophical Transactions of the Royal Society of London* 155, 459–512.

Minkowski, H. (1909). Raum und Zeit. *Physikalische Zeitschrift* 10, 104–111.

Morris, M. S., and Thorne, K. S. (1988). Wormholes in Spacetime and Their Use for Interstellar Travel: A Tool for Teaching General Relativity. *American Journal of Physics* 56, 395.

Morris, M. S., Thorne, K. S., and Yurtsever, U. (1988). Wormholes, Time Machines, and the Weak Energy Condition. *Physical Review Letters* 61, 1446.

Nagel, T. (1974). What Is It Like to Be a Bat? *The Philosophical Review* 83, 435–450.

Nagel, T. (1986). *The View from Nowhere*. New York: Oxford University Press.

Newton, I. (1687). *Philosophiae Naturalis Principia Mathematica*. London.

Noether, E. (1918). Invariante Variationsprobleme. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse*, 235–257.

Norton, J. D. (2008). The Dome: An Unexpectedly Simple Failure of Determinism. *Philosophy of Science* 75, 786–798.

Ori, A. (1991). Inner Structure of a Charged Black Hole: An Exact Mass-Inflation Solution. *Physical Review Letters* 67, 789.

Penrose, R. (2004). *The Road to Reality: A Complete Guide to the Laws of the Universe*. London: Jonathan Cape.

Peres, A. (1991). Two Simple Proofs of the Kochen-Specker Theorem. *Journal of Physics A* 24, L175–L178.

Perlmutter, S., and others (1999). Measurements of Ω and Λ from 42 High-Redshift Supernovae. *The Astrophysical Journal* 517, 565–586.

Planck Collaboration (2020). Planck 2018 Results. VI. Cosmological Parameters. *Astronomy and Astrophysics* 641, A6.

Poisson, E., and Israel, W. (1990). Internal Structure of Black Holes. *Physical Review D* 41, 1796–1809.

Popper, K. R. (1959). *The Logic of Scientific Discovery*. London: Hutchinson.

Price, H. (1996). *Time's Arrow and Archimedes' Point*. Oxford: Oxford University Press.

Rauch, D., and others (2018). Cosmic Bell Test Using Random Measurement Settings from High-Redshift Quasars. *Physical Review Letters* 121, 080403.

Raychaudhuri, A. K. (1955). Relativistic Cosmology. I. *Physical Review* 98, 1123–1126.

Rhim, W.-K., Pines, A., and Waugh, J. S. (1971). Time-Reversal Experiments in Dipolar-Coupled Spin Systems. *Physical Review B* 3, 684.

Riess, A. G., and others (1998). Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant. *The Astronomical Journal* 116, 1009–1038.

Rovelli, C. (2004). *Quantum Gravity*. Cambridge: Cambridge University Press.

Russell, B. (1913). On the Notion of Cause. *Proceedings of the Aristotelian Society* 13, 1.

Sbierski, J. (2018). The C^0-Inextendibility of the Schwarzschild Spacetime and the Spacelike Diameter in Lorentzian Geometry. *Journal of Differential Geometry* 108 (2), 319–378. arXiv:1507.00601.

Sbierski, J. (2023). Instability of the Kerr Cauchy Horizon under Linearised Gravitational Perturbations. *Annals of PDE* 9, 7. arXiv:2201.12295.

Sbierski, J. (2026). Lipschitz Inextendibility of Weak Null Singularities from Curvature Blow-Up. *Inventiones Mathematicae* 243 (3), 961–991. arXiv: 2409.18838.

Schrödinger, E. (1926). Quantisierung als Eigenwertproblem. *Annalen der Physik* 79, 361–376.

Schrödinger, E. (1935). Die gegenwärtige Situation in der Quantenmechanik. *Naturwissenschaften* 23, 807–812, 823–828, 844–849.

Shalm, L. K., and others (2015). Strong Loophole-Free Test of Local Realism. *Physical Review Letters* 115, 250402.

Shen, D. (2023). Construction of GCM Hypersurfaces in Perturbations of Kerr. *Annals of PDE* 9, 11. arXiv:2205.12336.

Smolin, L. (2001). *Three Roads to Quantum Gravity*. London: Weidenfeld & Nicolson.

Solomonoff, R. J. (1964). A Formal Theory of Inductive Inference, Part I. *Information and Control* 7, 1–22.

Susskind, L. (1995). The World as a Hologram. *Journal of Mathematical Physics* 36, 6377–6396.

Szeftel, J. (2026). Kerr Stability in the Full Subextremal Range. Preprint, arXiv:2609.40068.

't Hooft, G. (1993). Dimensional Reduction in Quantum Gravity. In *Salamfestschrift*, Singapore: World Scientific. arXiv:gr-qc/9310026.

't Hooft, G. (2016). *The Cellular Automaton Interpretation of Quantum Mechanics*. Cham: Springer.

Tegmark, M. (2000). Importance of Quantum Decoherence in Brain Processes. *Physical Review E* 61, 4194–4206.

Turing, A. M. (1936). On Computable Numbers, with an Application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society* s2-42, 230–265.

Van de Moortel, M. (2025). The Strong Cosmic Censorship Conjecture. *Comptes Rendus Mécanique* 353, 415–454. arXiv:2501.13180.

Verlinde, E. (2011). On the Origin of Gravity and the Laws of Newton. *Journal of High Energy Physics* 2011 (4), 29.

Wheeler, J. A. (1989). Information, Physics, Quantum: The Search for Links. In *Proceedings of the 3rd International Symposium on Foundations of Quantum Mechanics*, Tokyo.

Wigner, E. P. (1960). The Unreasonable Effectiveness of Mathematics in the Natural Sciences. *Communications on Pure and Applied Mathematics* 13, 1.

Wolfram, S. (2020). *A Project to Find the Fundamental Theory of Physics*. Champaign, Illinois: Wolfram Media.

Zeh, H. D. (2007). *The Physical Basis of the Direction of Time*. Fifth edition. Berlin: Springer.

Index

Page numbers in bold mark where a term is defined.

absolute elsewhere, 19, 199, 402
action, 18, 170, 258, 403, 432
arena, 259, **298**, 367, 405, 446, 511, 568, 597
arrow of time, 22, **115**, 150, 182, 244, 274, 358, 385, 525, 538, 597
  entropic arrow, 274, 358, 580
  structural arrow, 160, 181, 280, 358, 525, 567, 597
Aspect, Alain, 124, 149, 210, 277, 418
asymmetry, **28**
  in the plan of the book, vi
  in measurement as comparison, 22
  in the statement of the theorem, 48
  in the three axioms, 72
  in the exclusion of determinism, 98
  in what a cause is, 118
  in choice and the self, 142
  in hard randomness, 166
  in Bell's result, 214
  in superdeterminism, 245
  in signals and the light cone, 260
  in the direction of time, 275
  in the three times, 302
  in why time travel names nothing, 342
  in the two theories, 368

in the three horizons, 412
in the two constants, 439
in what physics cannot reach, 483
in simulation, 521
in computation and mind, 534
in the exclusions gathered, 563
in the technical forms, 579
axioms, the three, **73**
in the plan of the book, viii
in measurement as comparison, 4
in the statement of the theorem, 63
in what a cause is, 119
in choice and the self, 143
in hard randomness, 183
in the free will theorem, 201
in signals and the light cone, 257
in the three times, 299
in why time travel names nothing, 343
in the two constants, 455
in what physics cannot reach, 477
in simulation, 515
in computation and mind, 543
in the exclusions gathered, 563
in the technical forms, 606
first axiom, 73, 216, 244, 372, 423, 447, 479, 600
second axiom, 73, 145, 265, 282, 300, 421, 445, 579
third axiom, 63, 72, 109, 144, 237, 283, 302, 389, 410, 445, 482, 563, 600
Bekenstein, Jacob, 581
Bell's theorem, viii, 22, 169, 202, **209**, 231, 376, 414, 542, 567

Bell, John, viii, 22, 67, 87, 168, 188, 208, 230, 266, 330, 376, 404, 542, 567
binary hypothesis, 513
black hole, 19, **308**, 323, 356, 383, 405, 482, 587
block universe, 114, 284, 357
blue-shift, 327, 586
Bohm, David, 65, 199, 224, 241
Bohr, Niels, 60, 85, 375, 566
Bostrom, Nick, 526, 526, 551, 551
Bostrom, Nick, 526, 526, 551, 551
bound triple, **182**, 199, 244, 282, 525, 537, 565, 604
C constant and H constant, 411, 449
Cantor, Georg, 514, 602
catastrophic theorem of determinism, 196, 231, 566
category error, vii, **123**, 236, 263, 305, 354, 487, 536, 570
Cauchy horizon, 305, 320, 342, 580
causality, **113**
  in the three axioms, 72
  in the exclusion of determinism, 101
  in choice and the self, 142
  in hard randomness, 167
  in the free will theorem, 200
  in Bell's result, 213
  in superdeterminism, 232
  in signals and the light cone, 255
  in the direction of time, 282
  in the two theories, 390
  in the three horizons, 415
  in the two constants, 441
  in computation and mind, 540

in the exclusions gathered, 565
in the technical forms, 578
causation, 19, 109, **118**, 156, 167, 191, 220, 240, 256, 289, 554, 564, 577
choice, **141**
in measurement as comparison, 11
in the three axioms, 72
in the exclusion of determinism, 101
in what a cause is, 114
in hard randomness, 167
in the free will theorem, 191
in Bell's result, 213
in superdeterminism, 231
in signals and the light cone, 264
in the direction of time, 282
in why time travel names nothing, 357
in the two theories, 390
in the three horizons, 428
in the two constants, 437
in what physics cannot reach, 471
in simulation, 513
in computation and mind, 543
in the exclusions gathered, 564
in the technical forms, 577
chronology horizon, 311, 320, 583
closed timelike curve, 317, 319, 348, 582
comparer, **2**
in the plan of the book, vii
in the three axioms, 78
in the exclusion of determinism, 108

in what a cause is, 114
in choice and the self, 141
in hard randomness, 164
in Bell's result, 215
in superdeterminism, 238
in signals and the light cone, 255
in the direction of time, 282
in the three times, 300
in time machines, 322
in why time travel names nothing, 342
in the two theories, 374
in the three horizons, 411
in the two constants, 456
in simulation, 521
in computation and mind, 537
in the exclusions gathered, 565
comparison, **2**
in the plan of the book, vi
in the two kinds of law, 26
in the statement of the theorem, 47
in the three axioms, 70
in the exclusion of determinism, 95
in what a cause is, 112
in choice and the self, 142
in hard randomness, 164
in the free will theorem, 191
in Bell's result, 209
in superdeterminism, 231
in signals and the light cone, 255

in the direction of time, 274
in the three times, 299
in time machines, 322
in why time travel names nothing, 342
in the two theories, 368
in the three horizons, 400
in the two constants, 432
in what physics cannot reach, 470
in simulation, 508
in computation and mind, 534
in the exclusions gathered, 563
in the technical forms, 584
compounds, the four, **27**, 73, 214, 275, 368, 442, 500, 539
computation, xi, 24, 52, 129, 164, 499, **506**, 533, 570, 591
computationalism, 533
conjugate, **54**
in measurement as comparison, 18
in the three axioms, 73
in the exclusion of determinism, 100
in what a cause is, 122
in choice and the self, 147
in hard randomness, 170
in the free will theorem, 200
in Bell's result, 216
in superdeterminism, 246
in why time travel names nothing, 344
in the two theories, 372
in the three horizons, 406
in the two constants, 447

in what physics cannot reach, 479
in computation and mind, 545
in the exclusions gathered, 563
consciousness, 41, 88, 134, 142, **533**, 604
conservation, viii, **28**, 59, 114, 150, 212, 371, 418, 442, 483, 553, 566, 575
content, **5**
in the plan of the book, vi
in the two kinds of law, 26
in the statement of the theorem, 47
in the three axioms, 72
in the exclusion of determinism, 95
in what a cause is, 112
in choice and the self, 142
in hard randomness, 166
in the free will theorem, 189
in Bell's result, 212
in superdeterminism, 234
in signals and the light cone, 257
in the direction of time, 275
in time machines, 332
in why time travel names nothing, 342
in the two theories, 368
in the three horizons, 405
in the two constants, 435
in what physics cannot reach, 472
in simulation, 508
in computation and mind, 536
in the exclusions gathered, 565
in the technical forms, 575

context, **5**
in the plan of the book, vi
in the two kinds of law, 26
in the statement of the theorem, 48
in the three axioms, 73
in the exclusion of determinism, 95
in what a cause is, 112
in choice and the self, 142
in hard randomness, 166
in the free will theorem, 189
in Bell's result, 212
in superdeterminism, 234
in signals and the light cone, 257
in the direction of time, 275
in the three times, 306
in why time travel names nothing, 342
in the two theories, 368
in the three horizons, 405
in the two constants, 435
in what physics cannot reach, 472
in simulation, 508
in computation and mind, 536
in the exclusions gathered, 565
in the technical forms, 575
continuity, **27**
in the plan of the book, vi
in measurement as comparison, 12
in the statement of the theorem, 48
in the three axioms, 72

in the exclusion of determinism, 98
in what a cause is, 112
in choice and the self, 142
in hard randomness, 166
in the free will theorem, 193
in Bell's result, 213
in superdeterminism, 231
in signals and the light cone, 258
in the direction of time, 275
in the three times, 302
in why time travel names nothing, 342
in the two theories, 368
in the three horizons, 407
in the two constants, 439
in what physics cannot reach, 483
in simulation, 507
in computation and mind, 534
in the exclusions gathered, 563
in the technical forms, 579
continuum, xi, 43, 514, 533
Conway, John, x, 16, 147, 185, 189, 290, 566, 577
Copenhagen interpretation, 183, 541
cosmological constant, 66, 314, 327, 494, 581
cosmological event horizon, 176, 402, 454
CPT theorem, 292, 579
creation, 32, 72, 130, **142**, 243, 400, 455, 477, 560, 565, 599
cross-modal reach, 44, 371, 443, 500, 545
Dafermos, Mihalis, 328, 587
de Sitter space, 590

decoherence, 85, 422
deep why, 475
defining constant, 450
determinism, **95**, 231
  in the plan of the book, vii
  in measurement as comparison, 13
  in the two kinds of law, 34
  in the statement of the theorem, 54
  in the three axioms, 91
  in what a cause is, 112
  in choice and the self, 141
  in hard randomness, 166
  in the free will theorem, 189
  in Bell's result, 213
  in signals and the light cone, 255
  in the direction of time, 284
  in the three times, 307
  in time machines, 323
  in the two theories, 376
  in computation and mind, 534
  in the exclusions gathered, 564
  in the technical forms, 579
diagonal argument, 514
discontinuity, **28**
  in the plan of the book, vi
  in measurement as comparison, 22
  in the statement of the theorem, 48
  in the three axioms, 72
  in the exclusion of determinism, 98

in what a cause is, 118
in choice and the self, 151
in hard randomness, 166
in Bell's result, 214
in superdeterminism, 246
in signals and the light cone, 258
in the direction of time, 275
in the three times, 302
in why time travel names nothing, 342
in the two theories, 368
in the three horizons, 412
in the two constants, 442
in what physics cannot reach, 483
in simulation, 513
in computation and mind, 534
in the exclusions gathered, 563
in the technical forms, 579
domain curvature, 414, 599
domain of comparison, **19**
in the exclusion of determinism, 100
in hard randomness, 175
in the free will theorem, 191
in Bell's result, 223
in superdeterminism, 246
in signals and the light cone, 258
in the direction of time, 293
in the two theories, 379
in the three horizons, 402
in what physics cannot reach, 470

in simulation, 512
in computation and mind, 537
in the exclusions gathered, 567
in the technical forms, 599
Einstein, Albert, 29, 60, 106, 192, 218, 254, 274, 298, 318, 350, 366, 405, 519, 582
entanglement, 16, **188**, 209, 266, 414, 520, 547, 599
entropy, 37, **274**, 358, 384, 451, 567, 580
epistemic, vii, 56, 193, 290, 573, 578
Epistemic Sandwich, 290, 578
errors of extrema, 132
event, xii, **19**, 124, 254, 284, 300, 345, 385, 538
Everett, Hugh, 183, 199
exclusions, structural, 110, 139, 204, 228, 249, 270, 294, 317, 363, 376, 417, 599
exist, to, 21, **78**, 108, 143, 192, 215, 370, 486, 521
faithful reversal, **276**, 309, 359, 488, 571, 579
falsifiability, 486, 535, 601
focus of time, 299, 369
forbidden pairing, **36**, 53, 216, 346, 381, 417, 599
four theorems of being, 37, 150
Fourier transform, 201, 423
free will theorem, x, 16, 131, 185, **189**, 232, 290, 520, 566
gauge, 119, 440, 481, 601
general relativity, **19**
in the plan of the book, viii
in the two kinds of law, 42
in the statement of the theorem, 66
in the three axioms, 81

in the exclusion of determinism, 106
in what a cause is, 120
in superdeterminism, 246
in signals and the light cone, 257
in the direction of time, 290
in the three times, 298
in time machines, 319
in why time travel names nothing, 356
in the two theories, 366
in the three horizons, 416
in the two constants, 446
in what physics cannot reach, 482
in simulation, 506
in the exclusions gathered, 567
in the technical forms, 578
geometric time, xii, **284**, 299, 324, 342, 385, 538, 580
gravitation, 312, 377, 447, 569
Great Mystery, 201, 400
Gödel, Kurt, 320, 357, 376, 490, 582
hard randomness, x, 101, **120**, 164, 199, 222, 241, 282, 419, 510, 534, 565, 604
Hawking, Stephen, 257, 321, 356, 417, 476, 581
Heisenberg, Werner, 169, 200, 403, 460
hidden variables, 22, 193
horizon, **400**
in the plan of the book, xi
in measurement as comparison, 19
in the statement of the theorem, 66
in what a cause is, 126
in hard randomness, 176

in the free will theorem, 199
in signals and the light cone, 258
in the three times, 305
in time machines, 320
in why time travel names nothing, 342
in the two theories, 383
in the two constants, 449
in what physics cannot reach, 482
in the exclusions gathered, 565
in the technical forms, 580
identity, principle of, 133, 175, 508, 602
immanent, **30**
in measurement as comparison, 7
in the statement of the theorem, 53
in the three axioms, 72
in what a cause is, 113
in choice and the self, 142
in hard randomness, 178
in the free will theorem, 201
in Bell's result, 215
in superdeterminism, 236
in signals and the light cone, 255
in the direction of time, 282
in the three times, 299
in why time travel names nothing, 343
in the two theories, 369
in the three horizons, 400
in the two constants, 436
in what physics cannot reach, 472

in simulation, 524
in computation and mind, 534
in the exclusions gathered, 563
in the technical forms, 578
Incommensuration Theorem, **47**
in the plan of the book, vi
in measurement as comparison, 22
in the two kinds of law, 36
in the three axioms, 72
in the exclusion of determinism, 98
in what a cause is, 112
in choice and the self, 141
in hard randomness, 165
in the free will theorem, 193
in Bell's result, 214
in superdeterminism, 231
in signals and the light cone, 260
in the direction of time, 276
in the three times, 302
in why time travel names nothing, 342
in the two theories, 369
in the three horizons, 403
in the two constants, 442
in what physics cannot reach, 477
in simulation, 521
in computation and mind, 534
in the exclusions gathered, 563
in the technical forms, 579
incompressible residue, 165, 551, 606

information deficit, 314, 356
infra-physics, 485
just-so story, 554
Kerr black hole, 308, 323, 584
Kochen and Specker, 192, 577
Kochen, Simon, x, 16, 147, 185, 189, 290, 566, 577
Laplace, Pierre-Simon, 94, 234, 323, 584
Leibniz, Gottfried Wilhelm, 475, 508
light cone, x, 19, 81, 124, 199, **259**, 348, 412, 449, 491, 565
locality, 87, 191, **210**, 232, 268, 376, 415
Lorenz, Edward, 197, 517
Loschmidt, Josef, 279, 361
Mach, Ernst, 392, 460
many-worlds interpretation, 78, 199, 541
mass inflation, 328, 587
meager set, 333, 592
measurement, **2**
  in the plan of the book, v
  in the two kinds of law, 26
  in the statement of the theorem, 48
  in the three axioms, 70
  in the exclusion of determinism, 97
  in what a cause is, 116
  in choice and the self, 151
  in hard randomness, 165
  in the free will theorem, 189
  in Bell's result, 209
  in superdeterminism, 235
  in signals and the light cone, 255

in the direction of time, 287
in the three times, 298
in time machines, 323
in why time travel names nothing, 344
in the two theories, 367
in the three horizons, 403
in the two constants, 432
in what physics cannot reach, 470
in simulation, 508
in computation and mind, 535
in the exclusions gathered, 564
in the technical forms, 576
measurement problem, 20, 90
mesoscopic, **15**, 56, 103, 124, 194, 290, 370, 407, 454, 547, 578
mind uploading, 533
modal boundary, 312
modal rotation, 158, 179, 576
modalities, **30**
in the three axioms, 71
in what a cause is, 114
in choice and the self, 142
in hard randomness, 178
in Bell's result, 215
in superdeterminism, 236
in signals and the light cone, 255
in the direction of time, 282
in the three times, 299
in why time travel names nothing, 343
in the two theories, 369

in the three horizons, 400
in the two constants, 437
in what physics cannot reach, 472
in simulation, 522
in computation and mind, 534
in the technical forms, 576
modality, rule for matching, 369, 401, 436
model, 104, 121, 168, 293, **478**, 547, 583
Nagel, Thomas, 115, 542
Newton, Isaac, 29, 104, 274, 312, 368, 405, 445, 480, 511, 578
Noether, Emmy, viii, 33, 67, 212, 371, 423, 483, 566, 575
non-polarization principle, 127, 146
Norton's dome, 104
Novikov, Igor, 305, 322, 350, 584
nowhere dense set, 333, 592
object domain, 478
objective, **73**
in measurement as comparison, 6
in the two kinds of law, 27
in the statement of the theorem, 48
in what a cause is, 113
in choice and the self, 142
in Bell's result, 215
in superdeterminism, 247
in the direction of time, 275
in why time travel names nothing, 343
in the two theories, 368
in the three horizons, 409
in the two constants, 434

in what physics cannot reach, 472
in the exclusions gathered, 565
in the technical forms, 604
objectivity, **6**, 30, 75, 133, 143, 226, 439
omniscient, **30**
in the three axioms, 72
in the exclusion of determinism, 101
in what a cause is, 114
in choice and the self, 142
in hard randomness, 178
in the free will theorem, 201
in Bell's result, 215
in superdeterminism, 236
in signals and the light cone, 255
in the direction of time, 282
in the three times, 299
in why time travel names nothing, 343
in the two theories, 369
in the three horizons, 400
in the two constants, 437
in what physics cannot reach, 472
in simulation, 522
in computation and mind, 534
in the exclusions gathered, 563
in the technical forms, 578
paradox of embedding, 528
Past Hypothesis, 274, 358, 580
Penrose, Roger, 308, 323, 368, 584
phonon, 377

Planck length, 312, 405, 457, 511, 581
Planck scale, 312, 367, 405, 516
Planck's constant, xi
  in measurement as comparison, 18
  in the statement of the theorem, 54
  in the exclusion of determinism, 100
  in what a cause is, 131
  in hard randomness, 169
  in the free will theorem, 201
  in signals and the light cone, 258
  in the direction of time, 293
  in the three times, 312
  in the three horizons, 403
  in the two constants, 432
  in what physics cannot reach, 470
  in simulation, 511
  in computation and mind, 537
  in the exclusions gathered, 567
  in the technical forms, 581
Popper, Karl, 486
predicted world, 308, 323
pseudo-randomness, 525, 537
pure randomness, **101**, 119, 151, 166, 554, 566, 579
quantization, **22**, 37, 59, 312, 374, 413, 442, 499, 569, 601
quantum theory, **16**
  in the plan of the book, viii
  in the two kinds of law, 42
  in the statement of the theorem, 59
  in the three axioms, 81

in the exclusion of determinism, 99
in what a cause is, 120
in hard randomness, 165
in the free will theorem, 199
in Bell's result, 209
in superdeterminism, 232
in signals and the light cone, 266
in the direction of time, 274
in the three times, 301
in time machines, 321
in why time travel names nothing, 345
in the two theories, 367
in the three horizons, 403
in the two constants, 443
in what physics cannot reach, 481
in simulation, 506
in computation and mind, 541
in the exclusions gathered, 567
real, to be, 21, **78**, 108, 129, 143, 215, 374, 486, 521
residual, 90, 110, **122**, 157, 165, 222, 237, 284, 419, 501, 554, 565, 606
Root Tautology, xiii, **3**, 28, 57, 70, 99, 135, 248, 255, 280, 386, 436, 479
Schrödinger, Erwin, 23, 85, 274, 519
self, 7, 29, 123, 142, 201, 400, 440, 477, 507, 560, 565
signal, **6**, 254
in the plan of the book, vii
in the statement of the theorem, 67
in the three axioms, 79
in the exclusion of determinism, 100
in what a cause is, 120

in hard randomness, 169
in the free will theorem, 188
in Bell's result, 210
in superdeterminism, 233
in the direction of time, 288
in time machines, 331
in why time travel names nothing, 347
in the two theories, 395
in the three horizons, 402
in the two constants, 467
in what physics cannot reach, 488
in simulation, 530
in the exclusions gathered, 567
simulation hypothesis, 506, 534
singularity, **312**, 324, 416, 491, 567, 584
six intrinsics, **6**, 27, 48, 73, 247, 343, 409, 436, 477
six measures, viii, **12**, 32, 131, 378, 413, 434, 471, 567, 600
soft randomness, 164
spacelike separated, 191, 259, 577
speed of light, **254**
in the plan of the book, ix
in measurement as comparison, 19
in the statement of the theorem, 67
in the three axioms, 76
in the exclusion of determinism, 100
in hard randomness, 171
in the free will theorem, 188
in Bell's result, 210
in superdeterminism, 233

in the direction of time, 293
in the three times, 312
in time machines, 323
in why time travel names nothing, 349
in the three horizons, 402
in the two constants, 432
in what physics cannot reach, 470
in simulation, 511
in the exclusions gathered, 567
in the technical forms, 581
spin echo, 288, 347
stance, **6**, 48, 74, 142, 181, 199, 238, 282, 357, 523, 541, 565
Standard Model, 494, 536
strong cosmic censorship, xii, 309, 323, 342, 582
subjectivity, **181**
in measurement as comparison, 6
in the two kinds of law, 27
in the statement of the theorem, 48
in the three axioms, 73
in what a cause is, 113
in choice and the self, 142
in superdeterminism, 244
in the direction of time, 275
in why time travel names nothing, 343
in the two theories, 368
in the three horizons, 409
in the two constants, 436
in what physics cannot reach, 472
in simulation, 523

in computation and mind, 537
in the exclusions gathered, 565
in the technical forms, 604
superdeterminism, ix, 109, **123**, 172, 195, 224, 231, 264, 488, 566
superposition, 85, 287, 371
surface of unconsciousness, 201, 400
survivors, the two, **51**
in measurement as comparison, 22
in the three axioms, 76
in the exclusion of determinism, 109
in what a cause is, 118
in Bell's result, 214
in superdeterminism, 246
in the three times, 306
in why time travel names nothing, 343
in the two theories, 375
in the three horizons, 417
in the two constants, 443
in simulation, 524
in the exclusions gathered, 566
in the technical forms, 598
symmetry, **28**
in the plan of the book, vi
in measurement as comparison, 12
in the statement of the theorem, 48
in the three axioms, 72
in the exclusion of determinism, 97
in what a cause is, 112
in choice and the self, 151

in hard randomness, 166
in the free will theorem, 193
in Bell's result, 212
in superdeterminism, 231
in signals and the light cone, 258
in the direction of time, 275
in the three times, 302
in time machines, 329
in why time travel names nothing, 342
in the two theories, 368
in the three horizons, 407
in the two constants, 439
in what physics cannot reach, 483
in simulation, 507
in computation and mind, 534
in the exclusions gathered, 563
in the technical forms, 575
tachyon, 269
teleology, 102, 125, 148, 565
test in principle, 509
three horizons, xi, 293, 400, 470, 569
time machine, 295, 305, 318, 356, 583
time travel, x, 54, 91, 110, 295, 297, 318, 340, 571
triple of the real, 114, 142, 178, 441
Turing, Alan, 517, 536
uncertainty relation, **19**, 65, 99, 126, 169, 200, 403, 460, 564, 599
unknowable, 169, 201, **399**, 448, 491, 553

view from nowhere, 87, 115
Wolfram Physics Project, 551
wormhole, 297, 320, 583
zero-positive, 13, 76, 119, 461, 576