

Power Without Care

Multi-Polar Traps, the Attention Economy,
and the Structural Impossibility of Ethical Capitalism

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 979-8-90710-242-2

DOI: https://doi.org/10.67291/GINI2828

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

Introduction iv

Part I: The Architecture of Collective Failure 1

1. Multi-Polar Traps: Races to the Bottom 2

2. The Attention Economy and the Commercialization of Human Life 15

Part II: Games, Marketing, and the Corruption of Discourse 29

3. Nargames: On Games, Power, and Social Manipulation 30

4. Marketing, Media, and the Topology of Value Extraction 42

Part III: Technology, Power, and Civilizational Decline 57

5. The Manifesto Read as What It Is 58

6. Persuasion Without Care and the Case for a Total Ban; There Are No Viable Alternatives to Community 73

7. Intergenerational Knowledge Transfer and Civilizational Decline 86

Part IV: The Response 103

8. Right Action, Community Design, and the Foundation of Choice 104

Part V: Power Without Care Under General Intelligence 119

9. Power Without Care Under General Intelligence 120

Bibliography 159

Terms 167

Index 182

Introduction

Why can a civilization know what it should do, and still be unable to do it? In this book I answer that question.

Wherever many parties compete and nothing binds them all to one rule, each is forced to give up whatever does not help it win: a clean lake, a living wage, an evening at home with the children, an art, a trade. I call the structure in which many parties compete, with nothing binding them all to one rule, a *multi-polar trap*. In Chapter 1 I show the multi-polar trap at five scales, from a lake shared by fish farms to the choice that made every people farmers.

In this book, *structural* means caused by how a situation is arranged (who competes with whom, who is connected to whom, who gains from what), and not by anyone's intention. A structural failure is one that good people produce while each does what the situation rewards.

Part I shows the trap and its reach. Chapter 2 follows the market as it reaches into human attention and then into the human person.

Part II shows the machinery of the trap. Chapter 3 shows how games turn conversation into contest, and why marketing is the purest of those games. Chapter 4 shows the network that gives the distributor the maker's share.

Part III shows where the machinery leads. Chapter 5 reads the techno-optimist manifesto as what it is: *power without care*, the exercise of dominance without regard for life, for the future, or for the other. Chapter 5 reads one document closely because it is the clearest statement, by the winners of the game, of what they believe the game is for. Chapter 5, Section VIII, argues that ethical capitalism is impossible by the structure of competition itself. Chapter 6 shows artificial intelligence put to work persuading each person in private, and argues for a total ban on advertising, marketing and propaganda. Chapter 7 shows the chain by which skill passes from one generation to the next, and shows it broken.

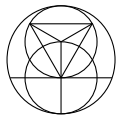
Part IV draws the response. Chapter 8 draws what follows from the argument of Chapter 5, Section VIII: no partial reform can hold, and the one viable response is right action grounded in choice and care, carried out by actual community. It also shows that community is where the governance that resolves a trap is made and kept.

Part V, Chapter 9, follows a second line through the book: what power without care becomes when the means of power is general artificial intelligence, a power that no one can take back, which in the end no longer needs the people who hold it. Along that line the book also says how a warning about general artificial intelligence can be given without the warning becoming one more piece of marketing.

The book's name for what it finds at the center of all this is power without care. Power is the ability to compel the alignment of another's choice. Care is giving to another for the other's sake, without exchange and without control. Power without care is what competition selects: because care is a cost, the multi-polar trap removes it, which is why ethical capitalism is impossible by structure. The rest of the book shows where that power comes from, how it works, and what can stand against it. The first question is the simplest one: how can people who all see a danger coming still walk into it together?

PART I

The Architecture of Collective Failure



Chapter 1
Multi-Polar Traps: Races to the Bottom

Part I: The Architecture of Collective Failure

ABSTRACT

A multi-polar trap is a situation in which every individual actor, pursuing their own rational self-interest, produces a collective outcome that is worse for everyone, themselves included. These traps are structural features of systems in which competition operates without adequate governance. This chapter examines a series of examples, from fish farming in a shared lake to the civilizational lock-in of agriculture (the commitment of whole peoples to a worse course that none of them can leave), that together reveal a common architecture of collective failure. In every case, the view from outside the system reveals an obvious better outcome, while the view from inside the system renders that outcome unreachable. The gap between these two perspectives is the space in which civilizations fail, and the space in which governance architecture, a mechanism outside the competition that changes what each party gains, must operate if civilizations are to survive.

I. Introduction: What Is a Multi-Polar Trap?

How can people who all see a danger coming still walk into it together? There is a class of problem that defeats intelligence, goodwill, and rational self-interest simultaneously. It is the class of problem in which every participant knows that the collective outcome will be terrible, and yet no participant can prevent that outcome by acting alone.

The term *multi-polar trap* names this structure precisely. It refers to a situation in which multiple independent agents, each making decisions that are locally rational, produce a globally irrational result. The trap is not that the agents are stupid. The trap is that the structure of the situation makes intelligence insufficient.

Every multi-polar trap shares a common architecture. There is a commons: some shared resource, value, or condition that benefits everyone. There is a competitive pressure: some incentive for each agent to extract from the commons faster than their neighbors. And there is an absence of coordination: no mechanism by which the agents can bind themselves and each other to the collectively superior outcome.

The pattern is ancient. It is visible in the tragedy of the commons, Garrett Hardin's name, in 1968, for the ruin of a shared resource by users who each gain from using it up. It is visible in the prisoner's dilemma, where each of two people does better by betraying the other, and both do worse when both betray. It is visible in arms races, and in the competitive dynamics of every market that has ever operated without effective regulation. What makes the concept important now is its application to civilizational-scale dynamics: the recognition that the same structure that destroys a shared lake can destroy a shared planet.

The examples that follow are drawn from many domains to show that multi-polar traps belong to no one domain: they are structural. They arise wherever competition operates in the absence of binding coordination, and they can only be resolved by *governance architecture*, some mechanism outside the competition that changes the incentive structure itself.

II. The Fish Farming Parable

The five examples that follow are drawn from Scott Alexander's essay "Meditations on Moloch" (2014), which calls such situations "multipolar traps", or races to the bottom, in which each competitor gives up a little more than the last to stay in the race; they are retold here in my words. Alexander states the common pattern in one line: in any competition that rewards one thing, a chance arises to sacrifice some other value for a little more of the thing rewarded.

Take a lake with a thousand fish farms on it, alike in every way, each owned by a different company competing with the others. Each farm makes a profit of one thousand dollars a month.

Every farm also puts out waste, and the waste fouls the lake. The waste of one farm lowers what the lake yields by one dollar a month for every farm on it. A thousand farms together lower it by a thousand dollars a month, and at that point no farm is making any money at all.

At first, capitalism seems to come to the rescue. Someone invents a filter that removes the waste. It costs three hundred dollars a month to run. Every farm installs one of its own free will, the pollution stops, and each farm now makes seven hundred dollars a month, still a respectable sum.

Then one farmer, call him Steve, tires of paying to run his filter, and switches it off. One farm's worth of waste now goes into the lake, and every farm's yield falls by one dollar. Steve makes nine hundred and ninety-nine dollars. Everyone else makes six hundred and ninety-nine.

The other farmers see that Steve does better than they do, because he does not pay to run his filter, and they begin to switch theirs off too. When four hundred have switched off, Steve is making six hundred dollars a month, less than he would have made if every filter had stayed on, and the virtuous farmers who still run their filters are making only three hundred.

Steve now goes from farm to farm, asking everyone to sign a voluntary pact to use filters, since otherwise every farm's yield goes down. Everyone agrees. They all sign the Filter Pact, except one farmer, Mike, a defector by temperament. Every farm but Mike's runs its filter again, and Mike makes nine hundred and ninety-nine dollars a month. Little by little the others decide that they too should earn what Mike earns, and they switch off their filters for the extra three hundred dollars of profit.

Alexander draws the lesson in one sentence: a person who acts only in his own interest never has any incentive to run a filter. A self-interested person has some reason to sign a pact that makes everyone use one, and often a stronger reason to wait for everyone else to sign it and then stay out. So the farms come to rest in an equilibrium no one wants, a state that no single farmer can improve by acting alone, in which no one will sign the pact at all.

From a global view, anyone can see that polluting the lake ends in catastrophe. From within the system, no one farmer can keep the lake clean, and a filter may not even repay its cost.

> **The lesson of the lake.** The fish farming parable shows the core mechanism of every multi-polar trap. Individual rationality and collective rationality diverge. Each actor, pursuing their own interest, degrades the commons on which all actors depend. Voluntary agreements fail because, absent enforcement, the incentive to defect is always present and always stronger than the incentive to cooperate.

III. The Malthusian Rat Trap

The name of this trap comes from Thomas Malthus, who argued in 1798 that a population grows faster than its food supply until hunger checks it. Alexander's second example begins with the first rats brought onto an untouched island. The island is full of plants to eat, and the first rats live an easy life, lounging, eating, and composing great works of art.

Each rat lives long and has a dozen children; each of those has a dozen more, and so on. Within a couple of generations the island holds ten thousand rats, which is as many as it can feed. Now there is not enough food and space for all, and a share of every new generation dies, so that the population stays at ten thousand.

One group of rats gives up art to spend more of its time scrounging for food. In each generation, a slightly smaller share of this group dies than of the rest, until after a while no rat composes any art at all, and any group that tries to bring art back dies out within a few generations.

The same holds for every value besides art. Any group that is leaner, meaner, and more bent on survival than the rest will in time take over. A group that kindly limits itself to two young per couple, to ease the crowding, dies out, outnumbered by its neighbors. A group that takes up cannibalism and finds that it pays will in time take over.

Suppose some rat scientists warn that the island's nut stores are running down faster and faster and will soon be gone. A few groups may try to eat only what the island can renew. They are outcompeted by their more selfish cousins. The nuts run out, most of the rats die, and the cycle begins again. Any group that argues for stopping the cycle is outcompeted by cousins for whom arguing is time better spent competing and eating.

From a global view, it is easy to say the rats should keep a comfortably low population. From within the system, each rat follows its genetic imperative, and the island is held in an endless cycle of boom and bust.

The rat trap shows that a multi-polar trap wears away values as well as resources. Art, altruism, cooperation, long-term thinking: every trait that does not directly help an animal survive the competition is selected against. The process is not deliberate. No rat decides to destroy art. The structure of the situation makes the behavior that destroys art the only behavior that persists.

The selection that lets only the behavior that destroys art persist is the mechanism by which competitive pressure, left unchecked, degrades the moral and cultural commons along with the material commons. It is the mechanism by which civilizations can lose the very capacities they need to survive.

IV. Capitalism and the Cutthroat Industry

The third example is a capitalist in a cutthroat industry. He employs workers in a factory sewing garments, and sells the garments at the smallest of profits.

He may want to pay his workers more, or to give them better conditions. He cannot, because either would raise his prices, cheaper rivals would take his customers, and he would go bankrupt. Many of his rivals may be decent people who would also like to pay their workers more. Without an ironclad guarantee that none of them will defect by cutting prices, none of them can afford to.

Like the rats, who lose every value except survival in the competition, companies under competition intense enough are forced to give up every value except making profit. Otherwise they are outcompeted by companies that pursued profit more ruthlessly and can therefore sell the same product for less.

Fit companies, the ones customers want to buy from, survive, grow, and are copied. Unfit companies, the ones no one wants to buy from, go bankrupt, and their corporate DNA dies with them. The reasons nature is red in tooth and claw are the same reasons the market is ruthless and exploitative.

From a global view, we can picture a humane industry in which every company pays its workers a living wage. From within the system, there is no way to bring it about without a governance mechanism that binds every competitor alike.

V. The Two-Income Trap

Elizabeth Warren and Amelia Warren Tyagi, in *The Two-Income Trap* (2003), describe how competition among families bid up the price of suburban houses in good school districts. Competition intense enough forced families to give up much of what they valued (time at home with their children, financial security, the possibility of living on one income) to raise what they could pay for a house, or else settle in less desirable neighborhoods.

The mechanism of the two-income trap is simple, in that when some families send both parents out to work, they can outbid the others for houses. The other families must do the same or lose. Once the new equilibrium is reached, every family works two jobs, no family has gained in the housing market, because the competition has simply moved up a level, and the values given up are gone for good.

From a global view, if every family agreed not to take a second income to compete for housing, every family would get as good a house as before, with one job per household. From within the system, without a mechanism able to enforce that agreement, every family that does not take a second income is left behind.

The two-income trap shows that multi-polar traps work at the scale of the family, as well as at the scale of corporations and nations. The competitive pressure reaches into the most intimate decisions of ordinary life.

VI. Agriculture and Civilizational Lock-In

Agricultural peoples outcompeted the nomadic peoples around them, and nothing could stop it. Perhaps hunting and gathering was more enjoyable, gave longer lives, and better served human flourishing. Jared Diamond, whom Alexander cites, called the turn to farming "the worst mistake in the history of the human race" (*Discover*, 1987), and found that "with the adoption of agriculture, height crashed". But where peoples competed hard enough, and agriculture, with all its disease, oppression and pestilence, was the more competitive way of life, everyone became farmers or died out as a tribe.

From a global view, it is easy to see that everyone should have kept the more enjoyable way of life and remained hunter-gatherers. From within the system, each tribe faced only the choice of farming or perishing.

This example is the most consequential of the five, because it works at the scale of a civilization and over thousands of years. It shows that multi-polar traps can lock whole civilizations into courses that are collectively inferior to the alternatives open to them. The lock-in is not chosen. It is *imposed* by competitive dynamics that no individual agent controls.

Agriculture was adopted because it was more competitive, in the specific sense of supporting larger populations and more concentrated military force; it was not adopted because it was better for human beings. Diamond puts the first point plainly: "Farming could support many more people than hunting, albeit with a poorer quality of life". The qualities that made agriculture competitive (density, hierarchy, the accumulation of surplus) are not the qualities that make a good human life, and the qualities that make a good human life were selected against by the competition between groups.

The selection of agriculture over the qualities that make a good human life is the multi-polar trap operating as a civilizational force. It is a structural feature of any situation in which many groups compete without binding coordination.

VII. The Common Structure

All of these scenarios share a single underlying architecture. Once one agent learns how to become more competitive by sacrificing a common value, all of its competitors must also sacrifice that value or be outcompeted and replaced by the less scrupulous. The system is likely to end up with everyone once again equally competitive, and with the sacrificed value gone forever.

From a global view, the competitors know they will all be worse off if they defect. From within the system, given insufficient coordination, defection is impossible to avoid.

Alexander's essay multiplies the examples, and six of them can be stated in a line each. Arms races consume national budgets that could otherwise fund infrastructure, health, education, and growth, yet any country that fails to spend on defense risks invasion by a neighbor that did. Cancer cells defect against the cooperative equilibrium of the body, and only the immune system keeps the natural incentive to turn cancerous in check. Jurisdictions race to the bottom, luring businesses with lower taxes and fewer regulations, until every one of them has minimal governance or has lost its economic activity to those that cut further. Education systems optimize for prestige rather than learning. Scientific institutions optimize for publication rather than truth. Governments optimize for re-election rather than the public good.

One more example near the head of Alexander's list shows how far a trap can go. In the dollar auction, a game the economist Martin Shubik described in 1971, a dollar is sold to the highest bidder, and the second-highest bidder must also pay what they bid. At every step, the bidder who

is behind loses less by bidding once more and winning than by stopping and paying for nothing, so the bidding passes a dollar and keeps going. A trap can cost its players more than the whole prize.

In every case, the structure is the same. Individual incentives diverge from collective welfare. Voluntary cooperation fails because the incentive to defect is structurally stronger than the incentive to cooperate. And the commons, whether it is a lake, a culture, a planet, or the capacity for trust itself, is degraded or destroyed. What is given up differs from case to case: a clean lake, an art, a living wage, an evening at home. In every case it is something kept up for the sake of someone or something other than winning. What is kept up for the sake of someone or something other than winning is care, and care is a cost.

The consistent resolution across all of these examples is governance architecture: some mechanism external to the competitive system that changes the incentive structure so that cooperation becomes individually rational. Without such architecture, the multi-polar trap is not merely likely. It is inevitable. Such a mechanism works on the game from outside it, and not by better play within it: it changes the game itself. Alexander's essay shows the trap. In this book I go on to ask what the multi-polar trap takes in every case, and how the market carries the trap into the person. I also ask who is paid for the trap, and where a rule that resolves the trap can be made and kept. *The Multipolar Trap; Institutional Failure and the Architecture of Choice* sets out the theory of these traps at length.

Where a commons has lasted, a binding rule has held it, and in many such commons the users themselves made the rule and kept it. Elinor Ostrom studied commons organized and governed by their own users, the oldest of them for more than a thousand years (*Governing the Commons*, 1990). Chapter 8 shows why community is the scale at which such governance is made and kept.

The mechanisms governance architecture can take include regulation, taxation, international treaty, cultural norm, and institutional design. Governance institutions themselves are subject to multi-polar traps: regulat-

ory capture (the control of a regulator by the industry it regulates), political corruption, and the race to the bottom among jurisdictions competing for business.

VIII. Summary

In this chapter I have shown the multi-polar trap at five scales: a shared lake, an island of rats, a cutthroat industry, the family, and the civilizational choice of agriculture. It derives from principles of game theory and collective-action theory: the prisoner's dilemma, the tragedy of the commons, the logic of collective action (Mancur Olson's name, in 1965, for why a large group fails to act in its common interest when each member gains by letting the others act). Its examples span ecology, economics, family structure, and civilizational history. Voluntary cooperation fails in the absence of binding coordination. Governance architecture is therefore needed. Governance can take many forms, and every form is itself open to capture. So the governance that resolves a trap must itself resist capture; Chapter 8, in its section "Community and the Making of Governance", shows where such governance is made.

IX. Conclusion

The multi-polar trap is the architecture of nearly every large-scale collective failure in human history. The multi-polar trap is the reason that knowing better is not enough. Intelligence, goodwill, and even unanimous agreement cannot, by themselves, produce collective action when the incentive structure rewards individual defection.

What unites the five examples of this chapter is their structure: individual rationality producing collective irrationality, in the absence of binding coordination. The fish farming commons, the Malthusian selection

trap, the cutthroat industry, the two-income trap, and the civilizational lock-in of agriculture are one phenomenon, operating at different scales and in different domains.

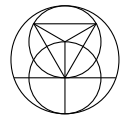
The implication is direct. Any serious attempt to address civilizational-scale risks (climate change, technological displacement, the erosion of democratic governance, the destruction of ecological commons) must begin by recognizing that these are multi-polar traps. They cannot be resolved by exhortation, by voluntary agreement, or by appealing to individual conscience. They can only be resolved by governance architecture that makes the cooperative outcome individually rational.

The question, then, is not whether governance is needed. The question is what kind of governance, and whether we can design it before the traps we are already in finish the work of consuming what remains of the commons. The largest of these traps is the competitive market itself. Where does it go once it has taken in every good and service?

KEY CONCEPTS IN THIS CHAPTER

- **Multi-polar trap**: a situation in which multiple independent agents, each acting rationally in their own interest, produce a collective outcome that is worse for everyone, including themselves.
- **Race to the bottom**: a competitive dynamic in which agents sacrifice shared values to gain competitive advantage, forcing all other agents to do the same until the shared value is destroyed.
- **Tragedy of the commons**: the degradation of a shared resource by individuals acting in their own interest, each of whom benefits from exploitation while the costs are distributed across all users.
- **Malthusian dynamics**: the process by which population growth outstrips resource availability, leading to competitive selection pressures that erode cooperative and cultural values.

- **Competitive pressure**: the force that compels agents to optimize for a single metric (survival, profit, power) at the expense of all other values, because agents that do not optimize are eliminated by those that do.
- **Governance architecture**: any mechanism external to a competitive system that changes the incentive structure so that cooperative behavior becomes individually rational; the necessary corrective to multi-polar traps.
- **Civilizational lock-in**: the process by which competitive dynamics commit entire civilizations to trajectories that are collectively inferior to available alternatives, with no mechanism for reversal.



Chapter 2
The Attention Economy and the Commercialization of Human Life

Part I: The Architecture of Collective Failure

ABSTRACT

The transformation of the modern economy follows a clear arc: from manufacturing and finance, through the service sector, to the commercialization of human attention itself. Each stage is a further abstraction from the making of tangible things, and each stage extends the reach of market logic into domains of human experience that were previously outside the market. The attention economy is the stage at which human focus, desire, and identity become the primary commodities. The endpoint of this trajectory, the "meat market" in which human embodiment is the last remaining commodity, shows how unconstrained market forces consume the substance of human life as well as its resources. This chapter traces that arc and identifies the attention economy as both a consequence of multi-polar trap dynamics and a mechanism by which those dynamics are intensified.

I. Introduction: The Arc of Commercialization

Where does the market go once it has taken in every good and service? The market has a history, and that history has a direction. It is the direction of expansion: the progressive incorporation of ever more domains of human interaction into the framework of commercial exchange.

This expansion follows from the internal logic of competitive market systems, which, as Chapter 1 showed, are subject to multi-polar traps that drive each participant to extract more, compete harder, and extend the reach of extraction into territory not yet commercialized. No one plans it.

Why does the expansion run in one direction only? Picture two islands, each with its own culture and its own marketplace, grown up apart. One day someone crosses the water and the two make contact. A disease would now spread from the shore where it landed, at some finite rate, and could be slowed. A market does not spread that way. The moment the people of one island notice that the other island has what their own island lacks, the chance to profit from the difference is seen on both sides at once. In that moment the markets of the two islands have become one market. The market grows as fast as the chance of an exchange can be perceived, and not as fast as goods can travel.

A barrier set against the market also fails in a way that a barrier against a disease does not. A virus has no intelligence; it meets the barrier and passes or does not. Everyone on either side of a barrier against the market who expects to profit from its removal works to find a way through: by smuggling, corruption, legal arbitrage, technical circumvention, social pressure and cultural influence. So every domain of life still outside the market is a difference waiting to be noticed, every participant in the market is rewarded for noticing it first, and every wall around such a domain has the intelligence of all those participants turned against it. The essay

"The Superluminal Contagion", in *Structure, Not Intent; Essays on Markets, Governance, and the Conditions of Freedom*, sets out this account of how the market spreads.

To understand the attention economy, it is necessary to trace the arc from its starting point in the material economy that lasted from the mid-twentieth century into the 1980s to its present endpoint in the commercialization of human consciousness itself.

II. Manufacturing and Finance: The Material Economy

In the 1980s, commercialization as a process operated mostly in the manufacturing and finance sectors. A business either made something or sold something, and between them those two were the whole of the business world. The services sector existed but was not a dominant force.

The era of the material economy was the era in which the relationship between economy and life was, at least in principle, legible. Value was created by making physical things. Finance intermediated the process of making and distributing those things. Energy powered the factories. The relationship between human labor and economic output was direct.

The limitations of this arrangement were also visible. Manufacturing needed heavy investment and was concentrated in a few places. The profits accrued primarily to the owners of capital: the factories, the equipment, the supply chains. Workers provided labor, and the market for labor was already a competitive game in which workers had little bargaining power.

The feature of the material economy that matters here is that it was *bounded*. There were things in human life that were not for sale (relationships, attention, identity, sexuality), and the market, for all its reach, did not yet have the tools to commercialize them.

The bound did not keep the trap out. The cutthroat garment industry of Chapter 1 is a manufacturing industry, and the living wage its owners could not pay is the value its competition forced them to give up. What

the bound did was confine the trap to what a factory could make and a shop could sell. Inside the bound, competition took from each firm whatever did not help it win. Outside it, the things not for sale were left alone, because no firm yet had a way to win by selling them.

III. The Service Economy and the Abstraction of Value

From the 1980s onward, the service sector expanded, the manufacturing sector was automated, and finance took on a larger role. The significant development was the *abstraction* of value, more than the growth of services.

In the manufacturing economy, value was tied to physical objects. In the service economy, value became tied to processes: consulting, management, logistics, administration. The worker performed a function, where before the worker had made a thing.

This abstraction had a consequence. When value is tied to objects, the market is limited by the physical constraints of production. When value is tied to processes, the market is limited only by the range of human activities that can be framed as services. The domain of the market expanded accordingly.

The word *abstraction* has a precise meaning here. To abstract a thing, in the economic sense, is to reduce all of its capabilities, characteristics and attributes to a single number: its price, its "value". The reduction compresses away nearly everything the thing was, and it cannot be undone: from the number alone, nothing of the thing can be rebuilt. What the reduction buys is comparison. Once every object and every process has been given a number, any one of them can be weighed against any other, and traded for it. A chair and an afternoon of a consultant's time have nothing in common but their prices, and their prices are all the market needs. So the service economy did not merely add services to the market. It extended the number to activities, and whatever can be given a number the market can reach. *Power and Its Prevention; On the Architecture of Collective*

Choice (Chapter 1, Section VII) names abstraction, with extraction and accumulation, as one of the three processes through which concentrated power perpetuates itself.

Finance, in particular, benefited from this abstraction. As computers entered the culture, the range of action available to financial instruments expanded dramatically. Things that had previously not been financial became financial. Relationships that had previously been outside the market were drawn into it.

IV. Finance as Dominance

The dominance of finance over the economy by the early twenty-first century was no accident. It followed from the relationship between abstraction and power.

The financial sector creates nothing material. Its function is to intermediate: to stand between producers and consumers, between present and future, between risk and return. Financial intermediation is a position of advantage by its nature, because whoever controls the channel through which value flows can extract *rent* from that flow. Rent is income taken for controlling access to something, and not for making that thing.

The dominance of finance is thus an early example of the principle developed in Chapter 4: that in any network (a set of points, called nodes, joined by connections), the nodes that distribute value will tend to capture more of it than the nodes that create value. Finance was the first sector to fully exploit this advantage of position at civilizational scale.

Finance also changed what a business is expected to be. An arrangement that takes a fixed share of a standing total grows its holder's wealth in proportion to that total. An *exponential* arrangement takes its share and adds it back to the total, so that the base on which the next share is taken has itself grown. Compound interest works this way, and so do network effects and virality. Each is *auto-catalytic*: a process one of whose products speeds up the rest of the process, as fire does, or a network whose value to

each member rises with the number of members, so that each new member makes joining more attractive to the next. Since about the 1980s, more and more business plans have been built on some exponential aspect, and the investors who fund new firms now prefer such plans almost to the exclusion of any other.

The founder of an exponential business does not merely want a good business. The founder wants a position that cannot be taken: a position from which the ordinary correction by competition has been permanently shut out. A position from which correction by competition has been permanently shut out is not a market outcome. It is the abolition of the market by a party who arrived through the market and closed the door behind them, and it is spoken of, where such things are spoken of, as the highest achievement available. It is also the advantage of position that this section describes, carried to its end: a channel through which value must flow, held by one party, with no second channel left. *Markets, Money, and the Moral Economy; Dialogues on the Form of the Transaction and the Distribution of Choice* (Chapter 1) sets out the exponential model at length.

V. The Attention Economy

By the turn of the millennium, around 2001, businesses were beginning to explore the commercial potential of the social world. The exploration of the commercial potential of the social world, around 2001, was the first serious attempt to turn human attention into a tradeable commodity. Between 2010 and 2015, the attention economy came into its own: it was no longer an experimental field but a developed industry.

The history of energy shows what happened to attention. Attention had become an independent economy, a driver for the marketing economy, in the same way that energy had become a driver for the manufacturing economy. Before the Industrial Revolution, most work was done by human and animal muscle, with wind, water and wood beside them, and energy was not bought and sold as one measured thing. Then coal and oil

became economic goods, and then electricity became a market. The pattern repeated: something that had previously existed outside of commercial exchange was drawn into it, and once inside, the market became increasingly efficient at extracting it.

What, exactly, is extracted? Attention is not a quantity of time. It is the faculty by which a person determines what is relevant. Out of everything that could be considered at any moment, attention selects the few things that will be considered, and everything downstream of it (what a person thinks about, what a person notices, what a person is able to care about, what a person takes the world to be like) is set by that selection. A person cannot care for what the person does not notice, so a market that sells attention sells the ground on which care stands. So what is sold, in fractions of a cent for each view, is a slice of the operation by which a human being decides what matters, and it is sold by a party who is not the person whose faculty it is. The sale is made between a platform and an advertiser. The person whose sense of what matters is the actual commodity is not a party to the sale, receives a service called free as if that were the payment, and never names a price for what was taken. Chapter 5 of *Markets, Money, and the Moral Economy* follows this account of what is extracted.

The financial process, which is the market process, is growing to take in more and more of the totality of all human interactions. Take the field of all interactions in the entire system, virtual and embodied, digital and physical, and ask what proportion has become associated with commerce: that proportion is steadily increasing.

The steady increase in the proportion of all interactions that is associated with commerce is the metastasization that can be understood as a kind of cancer affecting the world. When cancer metastasizes, it spreads to other places and turns them into copies of itself. The organs break down because they become cancer. The market is doing the same thing to the non-commercial domains of human life.

VI. Content, Influence, and the Commodification of Self

The attention economy did not exist as a distinct sector before the market forces spread into the domain of human focus. Before then, the places where relationships were financialized, turned into something bought and sold for gain, were limited to the executive levels of companies and government: the spaces where personal connections translated directly into business advantage.

Now all relationships, everywhere in the world, are financialized. The financializing of all relationships, everywhere in the world, is what the attention economy is. Things that had not been part of the market became part of the market, and the market became increasingly efficient at extracting the new resource.

The same reach changes what it means to make something well. The word *art* once named a craft: a way of doing a thing well, which could be learned, which had standards internal to itself, and at which a maker could be better or worse whether or not anyone bought the result. The influence economy runs a chain of substitutions on it: fame for prowess, attention for money, money for capital. Prowess is a property of the maker, judged by other makers. Fame is a property of the audience, judged by counting. When fame stands in for prowess, the standard inside the craft is replaced by a standard outside it, and at the end of the chain the question "is it good?" has no content left that the question "what does it fetch?" does not exhaust. A work that fetches nothing is then taken not for a good work that failed to sell, but for no work at all.

With attention as the resource, the market became better and better at extracting people's attention. The market's growing skill at extracting people's attention is the origin of what I call the attention-deficit economy. It is an economy of attention deficit because what it sells is taken from the attention each person has, so the more of it the market sells, the less of it each person has left. People no longer have free time, because every social

media platform, every app, every device is engineered to capture whatever free attention remains. People are no longer relating to one another; they are relating to machines, because the machines have become better than other human beings at providing *supernormal stimuli*: stimuli engineered to be more compelling than any naturally occurring alternative. The term is the ethologist Niko Tinbergen's, for an artificial signal that draws a stronger response than the natural one.

The consequence is measurable. Every human being has twenty-four hours in a day. No technology adds or subtracts a single minute from that allotment. The twenty-four hours in a day are the one dimension in which all humans are equal. But the extraction systems draw down on that allotment relentlessly: biological processes, involuntary demands on attention, and then the addictive technologies that appear voluntary but function as compulsions. What people call "hangout time", the free time people have to spend with one another, is steadily consumed. The American Time Use Survey of the United States Bureau of Labor Statistics measures it: the time an average American aged fifteen or over spent socializing and communicating with others fell from 0.70 hours a day in 2010 (about 42 minutes) to 0.58 hours a day in 2025 (about 35 minutes), and the share of people who did so at all on an average day fell from 38 percent in 2015 to 30 percent in 2025.

What a person can do with those hours depends also on *click rate*: the rate at which the mind shifts from one frame to the next. Click rate is a measure of mental metabolism. It sets how much attention a person has to spend, and therefore how much of that attention is exposed to extraction. A platform that raises the rate at which the mind shifts its frame also raises what it can take from each hour.

VII. The Meat Market: The Final Commodity

The trajectory has an endpoint, and the endpoint is stark. When labor is automated, when ideas can be generated by machines, when services can be provided by algorithms, what remains for the human being to bring to the market?

What remains is the body itself. Identity. Sexuality. Appearance. The irreducible embodiment of the human person.

The selling of the body itself, with its identity, sexuality and appearance, is the "meat market": the stage at which the only commodity left to commercialize is the human being as such. Influencer culture, the commercialization of personal image, the monetization of dating and relationships, the gig economy of physical labor: all of these are early forms of the same logic. When everything else has been automated or abstracted, the last remaining resource for market extraction is the person.

The term names, without euphemism, the endpoint of unconstrained market expansion. The commodification of attention leads to the commodification of identity, which leads to the commodification of the person. At each stage, something that was previously outside the market is drawn into it, and once inside, the competitive dynamics of extraction and optimization apply to it with the same ruthlessness they apply to fish farms and garment factories.

The question this trajectory raises is existential. When everything is commodified, when human attention, identity, and embodiment are all tradeable goods in a market optimized for extraction, what remains of human life that is not for sale?

VIII. Summary

In this chapter I have traced the arc from manufacturing through finance and services to the attention economy and the commodification of the person. Each stage follows from the previous one, driven by the expansion of market logic into new domains. The market, like a spreading cancer, turns healthy tissue into copies of itself.

The first step of my argument in this chapter is that the market spreads as fast as a difference can be noticed, and every wall around what is not for sale has the intelligence of all who would profit turned against it. The material economy was bounded, and the trap of Chapter 1 worked inside the bound. The service economy reduced activities to a number, and whatever has a number the market can reach. Finance took rent from its position in the channel, and the exponential business seeks a position that competition cannot correct. The attention economy sells the faculty by which a person decides what matters, in a sale to which that person is not a party. The influence economy replaces the standard inside a craft with a count of the audience. The meat market is the end of the line: the person as the last commodity.

The attention economy is itself a multi-polar trap, in which each content creator competes for attention, collectively degrading the commons of human focus and meaning. The technologies that accelerate the attention economy (algorithmic recommendation, behavioral prediction, AI-driven content generation) are the subject of Chapter 6, where their use for advertising and propaganda is shown to be a qualitative escalation of the dynamics described here. Whether and how the trajectory can be interrupted is the subject of Chapters 6 and 8.

IX. Conclusion

The attention economy is the latest stage in the progressive expansion of market logic into every domain of human interaction. Each previous stage (manufacturing, finance, services) extended the market's reach. The attention economy extends it into the domain of consciousness itself.

The consequence is clear. As the market becomes more efficient at extracting attention, the free time people have to spend with one another, to reflect, and to cultivate values that are not market values, shrinks. People become busier. Relationships become thinner. The cultural commons, the shared capacity for meaning, for slow conversation, for the kind of human presence that cannot be monetized, erodes.

This erosion is not accidental. It is the structural result of competitive market dynamics operating without adequate governance. It is a multipolar trap in which every individual actor (every content creator, every platform, every advertiser) is rationally rewarded for capturing more attention, while the collective result is the impoverishment of the attentional commons, the shared human capacity for focus, reflection and real connection, on which all meaningful human life depends.

The remedy will not come from within the market. The market is doing exactly what markets do. The remedy must come from outside the competitive system: from governance, from community, from a deliberate choice to place some domains of human life beyond the reach of commercial extraction. Before that remedy can be found, one question must be answered: what kind of social process is it that the market has turned our conversations into?

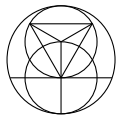
KEY CONCEPTS IN THIS CHAPTER

- **Attention economy**: the economic system in which human attention is the primary commodity; the stage at which market logic extends into the domain of consciousness itself.
- **Market metastasization**: the progressive expansion of market logic into domains of human life that were previously outside commercial exchange, analogous to the spread of cancer through healthy tissue.
- **Supernormal stimulus**: a stimulus engineered to be more compelling than any naturally occurring alternative; the mechanism by which digital platforms outcompete human relationships for attention.
- **Attentional commons**: the shared human capacity for focus, reflection, and real connection; the resource degraded by the competitive dynamics of the attention economy.
- **The meat market**: the endpoint of market expansion, at which the only remaining commodity is the human person: identity, sexuality, embodiment.
- **Click rate**: the rate of cognitive frame-shifting; a measure of mental metabolism that determines an individual's effective attentional bandwidth and therefore their exposure to attentional extraction.
- **Abstraction**: the reduction of all the capabilities and attributes of a real thing or activity to a single number, its price; irreversible, and the condition under which anything can be compared with and traded for anything else.
- **Exponential model**: an arrangement that adds each share it takes back to the total on which the next share is taken, as compound interest, network effects and virality do; its founder seeks a position that competition cannot correct.
- **Attention** (**what is extracted**): the faculty by which a person determines what is relevant; sold in the attention economy by a party who is not the person whose faculty it is.

- **Financialization of relationships**: the process by which human relationships, previously outside the market, are incorporated into commercial exchange through social media, influencer culture, and platform economics.

PART II

Games, Marketing, and the Corruption of Discourse



Chapter 3
Nargames:
On Games, Power, and Social Manipulation

Part II: Games, Marketing, and the Corruption of Discourse

ABSTRACT

The concept of a "nargame", a narcissistic social game, names the mechanism by which game-playing dynamics corrupt discourse, collaboration, and civilizational problem-solving. This chapter develops a precise definition of "game" as competition within a cooperative context, distinguishes games from real collaboration, and traces how the embedding of games within games creates structures optimized for extraction. The distinction between game and design, between social processes oriented toward power and social processes oriented toward solving real problems, emerges as one of the most consequential distinctions available to social analysis. Marketing is identified as the paradigmatic social game, and nar-social behavior (behavior oriented toward narcissistic social game-playing) as the mechanism by which games infiltrate and disable authentic collaboration.

I. Introduction: What Is a Game?

What kind of social process has the market turned our conversations into? In Chapter 2 I showed the market reaching into human attention. In this chapter I show the social form that the market's reach takes, which is the

game. The word "game" carries an implication of playfulness, and most people associate games with fun. This association is itself part of the game.

Some games are designed to imply playfulness when being sold, or to be experienced as playful by at least some players, some of the time, so as to keep players playing. Casino gambling games are made as engaging as possible while the player slowly drains away their savings. The apparent playfulness is a feature of the design: a mechanism for keeping people playing a game whose function is to move value from player to house.

It is difficult to name a common, widely known game that does not involve some notion of benefit associated with winning, and also some notion of benefit to the game maker, though the benefit to the game maker is usually much less visible. Entrepreneurism is a game. Employment is a game. The market itself is a game. In each case, the game presents as voluntary and mutually beneficial while it operates to transfer value from participants to designers.

Others have mapped the word differently. Ludwig Wittgenstein asked what all the proceedings we call games have in common, and answered that they share no single feature, only overlapping similarities, which he called "family resemblances" (*Philosophical Investigations*, 1953, sections 66 and 67). James Carse divided games into two kinds: "A finite game is played for the purpose of winning, an infinite game for the purpose of continuing the play" (*Finite and Infinite Games*, 1986). The definition this chapter gives is built to show one thing: where the value in a game goes.

II. The Structure of Games: Benefit, Competition, Extraction

Descriptively, the most useful definition of "game" is *competition within a cooperative context*.

The definition of a game as competition within a cooperative context can be narrowed to *induced goal-seeking behavior within a created context*. The first group, the players, pursue goals defined by the second group, the

game makers. The cooperative context is implicit: players agree to abide by the rules and pursue the defined goals. The competition is the mechanism: players compete against each other for the designated reward.

The critical feature of this structure is the embedding of games within games. When a game exists within a larger competitive context, the game maker is also a player in the outer game. The inner game is designed so that the playing of it results in an automatic transfer of benefit from the inner players to the inner game maker, thereby fulfilling the game maker's goals in the outer game.

Since this embedding can be applied recursively, the vast majority of observed game processes will be designed so that the game makers are the winners and the game players are the losers. A game maker whose game does not pay its maker loses the outer game, and the game ends with its maker. The games that last are therefore the games that pay their makers. The survival of only the games that pay their makers is the selection that removed art from the rats' island in Chapter 1. In that selection, what persists is what wins. Add any variation of the concept of finite constrained resources, and the consequences become visible: hidden game contexts, forced participation, extractive rent-seeking, control of the menu of choices offered to the players, and the emergence of dominance dynamics regardless of care or consequence.

III. Nargames: Narcissistic Social Games

A *nargame* is a narcissistic social game: a game whose primary function is the acquisition and maintenance of social power by the game's designer or dominant player, at the expense of other participants.

The "nar-" prefix refers to narcissistic, in the structural sense of a social process organized around the needs and goals of one party while presenting as organized around the needs and goals of all parties, and not in the clinical sense of a personality disorder.

A nargame is defined by its function, the acquisition and maintenance of social power by the game's designer. A game of chess between two friends serves no designer's power, so it is a game without being a nargame. Two friends at chess are opponents only on the board and only until the game ends; the loss of trust described in Section IV comes when a game is played inside a conversation that presents itself as collaboration.

There are games in which many people feel trapped. Market economies, employment markets, and educational credentialing systems are among the most prominent. Some players take part because they believe they have to: to put food on the table, to secure shelter for their children, to maintain the appearance of social legitimacy.

Control the means of survival, or of belonging, or of sexual access, and you control the people. To control anything, you must weaponize something, which is ethically questionable. But that consideration does not stop the weaponization. Hence the observation that market games optimize for bad actors. The incentive structure of market games rewards those who are willing to exploit control over others without regard for consequence.

IV. Games as Camouflage: Power Dynamics and Social Manipulation

Games are camouflage. They are the spaces in which many players of all types can hide. The game provides a context (a set of rules, a framing, a narrative) within which actions that would otherwise be visible as manipulation or exploitation become normalized as "just playing the game".

The claim that "everything social is games" is worth examining. Social collaboration *can* be design: that is how technology happens, how complex problems get solved, how civilizations get built. The distinction between game and design is the distinction between social processes oriented toward power and social processes oriented toward solving real problems.

For solving civilizational problems (existential risk, a risk that could end humanity or all life; governance design; the restoration of ecological commons), there needs to be social process that is not game. The moment collaboration becomes a game, the objective shifts from solving the problem to winning the competition. And winning the competition does not require solving the problem, and often actively prevents it.

Game-playing cannot fail to lead to a loss of trust, since everyone must view everyone else as a possible opponent. No one can care for a person while treating that person as an opponent. Games ensure that there are no true tribes. They make the cooperative notion of community impossible: a divide-and-conquer dynamic that operates by structure rather than by intention. No trust means no real social cohesion, only the temporary illusion of it. All contest, no collaboration. And no collaboration means no civilizational problem-solving, which means real problems go unsolved.

V. Game Theory vs Design Theory

The distinction between *game theory* and *design theory* is one of the most consequential distinctions available to social analysis. Game theory here names the orientation of a player who asks how to win within given rules; the mathematical discipline of the same name, founded by John von Neumann and Oskar Morgenstern in *Theory of Games and Economic Behavior* (1944), studies that orientation among others.

Game theory describes systems in which agents with conflicting goals compete for advantage within a shared framework of rules. The outcomes are determined by the strategic choices of the agents. The framework itself is taken as given.

Design theory describes systems in which agents collaborate to create something that addresses a real need. They do it by working with the actual material of the world (physical constraints, community needs, ecological requirements, the specific properties of the problem at hand) and not by competing against each other.

The game theorist asks: "Given these rules, how do I win?" The designer asks: "Given these constraints, how do I solve this?" The distinction is between orientations, and not between competition and cooperation in the abstract. The game orientation treats other people as obstacles or tools. The design orientation treats the world as the material to be worked with.

VI. Nar-Social Behavior: Recognition and Response

Nar-social behavior, behavior oriented toward narcissistic social game-playing, has recognizable markers. It shows up as constant topic changes, labeling and dismissiveness, circular argument, selective misunderstanding, withholding of earned praise, and refusal to admit mistakes. It shows up as attempts to set and control the context of discussion, and as a consistent pattern of never taking real personal risks while demanding that others do so.

These markers are not incidental. They are how social games are played. Each one serves a specific function in keeping the game-player ahead. Topic changes prevent the completion of any argument that threatens the game-player's position. Labeling and dismissiveness shut out engagement with ideas that cannot be co-opted. Selective misunderstanding creates opportunities for mislabeling. Withholding praise prevents potential competitors from building social standing.

These behaviors are often not deliberate in the sense of being consciously planned. They are habitual: the social equivalent of muscle memory. People who have spent their lives in social settings dominated by game-playing develop these patterns automatically, because the patterns are the ones that survive in such settings. The process is the same as the Malthusian selection described in Chapter 1: game-playing behaviors are selected for because they win in competition, whatever their effect on the collective.

Nar-social names a mode of engagement, and not a kind of person. It is the mode in which the social surface of an exchange (who is winning, who is allied, who is diminished) has become its actual content, while the subject under discussion serves only as the material from which those social effects are made. Everyone enters that mode sometimes, so the word describes a pattern of moves that a setting selects, and it diagnoses no one.

For the same reason, the response begins with oneself. Discernment turned outward first becomes one more move in the game, since the accusation of game-playing is among the most reliable game-playing moves there is. Discernment turned inward first has no such payoff, which is why it is the harder practice and the only starting point that can be trusted.

One marker gives the plainest test. Ask that a coined term be defined, or that a proposed distinction be stated as an idea that can be examined. Where the orientation is design, the request is welcome, because a claim stated plainly is a claim that can be tested and improved. Where the orientation is a game, the request is a threat, and the usual answer is neither agreement nor refusal but a retreat from explanation: a turn to some further social move, or to the charge that the request itself is improper or unkind. The definition never arrives; something social arrives in its place. *Design and Games; On the Two Orientations of Conversation* (Sections III to V) sets out the nar-social mode and the retreat from explanation at length.

VII. Marketing as the Paradigmatic Game

Marketing is not a peripheral activity in the modern economy. It is the paradigmatic social game: the purest expression of game dynamics in social context.

A marketing message never apologizes. It never admits limitation. It never acknowledges cost or risk to the buyer, or the interests of anyone other than the seller. It treats attention-seeking as its primary product. It is organized entirely around the game-maker's benefit, while always

presenting as organized around the buyer's benefit. The absence of apology and vulnerability is the diagnostic signal that distinguishes marketing from real communication.

In real collaboration, people apologize when they are wrong. They acknowledge limits. They show vulnerability. Apology, acknowledged limits and shown vulnerability are the signals that distinguish authentic engagement from performance. Marketing culture has steadily eliminated these signals from public discourse, replacing them with the performance of confidence, certainty, and authority: the markers of game dominance.

The consequence is that the social environment becomes increasingly saturated with game-playing behavior, while authentic collaboration, the kind needed to solve actual civilizational problems, becomes increasingly difficult to initiate, sustain, or even recognize.

When attention-seeking and status competition become the main ways people deal with one another, the result is a social system of its own: *attentionalism*, the social expression of the attention economy of Chapter 2.

Because a marketing message never admits a risk, and a warning exists to name one, the warning about general artificial intelligence must not itself be marketed; Chapter 9, Section I, sets this out.

VIII. The Distributed Immune System: Discernment Over Moderation

The standard response to nar-social behavior is centralized moderation: rules, enforcement, authorities who decide what is and is not acceptable. This approach has a fundamental problem: the moderators themselves are operating within game contexts, and the moderation process can itself become a game.

The alternative is distributed discernment: equipping all participants with the capacity to recognize nar-social behavior and to respond to it appropriately, without centralized authority.

Distributed discernment is a distributed immune system rather than a centralized police force. It works by making nar-social behavior socially unpopular, and not by prohibition: by creating a culture in which the markers of game-playing are recognized and named, and in which real collaboration is valued and rewarded.

The cultivation of this discernment is one of the most important practical capacities a community can develop. It requires the ability to notice when games are being played, to identify who is defining "winning", to check for empathy, apology, and humility as markers of authentic engagement, and to keep the distinction between game and design alive and active in every social interaction.

IX. Summary

In this chapter I have defined a game as competition within a cooperative context, and I have shown how embedding games within games extracts value. I have also identified marketing as the paradigmatic social game, and I have shown nar-social behavior as how social games are played. The distinction between game and design is one of the most consequential distinctions available to social analysis.

The first step of my argument in this chapter is that a game is induced goal-seeking within a context its makers create, and because games are embedded in games, the games that last are the games that pay their makers. A nargame is a game whose function is the social power of its designer, presented as serving every player. Games hide manipulation as "just playing the game", and because every player must treat every other as a possible opponent, they destroy the trust that collaboration needs. The game orientation asks how to win within given rules; the design orientation asks how to solve a real problem with the real material of the world. Nar-social behavior is a mode that settings select, not a kind of person, and its plainest marker is the retreat from explanation. Marketing is the purest game: it never apologizes, never admits a limit, and never

names a cost to the buyer. Chapter 9, Section I, applies this to the warning about general artificial intelligence: the warning must not be carried by marketing's methods, and the consent asked for must match the risk of the ask. The answer to nar-social behavior is a distributed immune system of discernment, not a central moderator who can be drawn into the game. How a game-dominated social environment becomes one that supports real collaboration is the subject of Chapter 8.

X. Conclusion

The concept of the nargame, the narcissistic social game, names a pervasive feature of human social life. Games are the default mode of social interaction wherever *design consciousness* (the habit of asking of every shared task the designer's question, "Given these constraints, how do I solve this?") is not deliberately cultivated.

The distinction between game and design is a question of civilizational survival. If the dominant mode of social interaction is game-playing (competition for power, status, and resources within cooperative-seeming contexts), then civilizational problem-solving is impossible, because the social processes needed to solve problems are disabled by the game dynamics operating within them.

Marketing, as the paradigmatic social game, has infiltrated every domain of public discourse. The result is a social environment in which the signals of authentic collaboration (apology, vulnerability, admission of limits) have been eliminated in favor of the signals of game dominance (confidence, certainty, authority). In such an environment, the capacity for collective problem-solving atrophies, and the civilizational traps identified in Chapter 1 become inescapable.

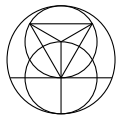
The remedy is discernment, not moderation: the distributed capacity to recognize games as games, to name nar-social behavior without flinching, and to cultivate design consciousness as a deliberate alternative to the

default game-playing mode of human social life. If marketing is the paradigmatic game, what is the board on which it is played, and why does the house always win there?

KEY CONCEPTS IN THIS CHAPTER

- **Nargame**: a narcissistic social game; a social process organized around the acquisition of power by the game designer at the expense of participants, while presenting as mutually beneficial.
- **Nar-social behavior**: behavior oriented toward narcissistic social game-playing, characterized by topic changes, labeling, selective misunderstanding, withholding of praise, and refusal to admit error; a mode that a setting selects, and not a kind of person.
- **Retreat from explanation**: the characteristic answer, under a game orientation, to a request that a term be defined or a claim be stated plainly: a social move offered in place of the definition.
- **Game (structural definition)**: competition within a cooperative context; induced goal-seeking behavior within a created framework, where the game maker benefits from the playing of the game.
- **Game theory vs design theory**: the distinction between the orientation of a player who asks how to win within given rules (game theory) and the orientation of a designer who asks how to solve a real problem with real materials (design theory).
- **Attentionalism**: the social system in which attention-seeking and status competition are the dominant modes of social interaction; the social expression of the attention economy.
- **Distributed immune system**: the alternative to centralized moderation; equipping all participants with the discernment to recognize and respond to nar-social behavior without centralized authority.

- **Marketing as paradigmatic game**: the identification of marketing as the purest expression of social game dynamics: no apology, no vulnerability, no acknowledgment of limits; pure impression management in service of extraction.



Chapter 4
Marketing, Media, and the Topology of Value Extraction

Part II: Games, Marketing, and the Corruption of Discourse

ABSTRACT

The relationship between engineering and marketing, between creation and distribution, is a consequence of the shape of a network, and not a matter of individual choice or market accident. In any network, the nodes that create value tend to be sparsely connected, while the nodes that distribute value are densely connected. This difference in connection allows distributors to extract rent from the commerce flowing through the network, while the actual creators receive a diminishing share. At the same time, the escalation of media bandwidth, from text through audio and video to immersive virtual reality, enables progressively more effective manipulation, displacing low-bandwidth media and the reflective thought they support. These two dynamics, the topology of extraction (the shape of the network through which value is taken) and the escalation of sensory manipulation, together make up the machinery by which marketing dominates creation and deception becomes systemic.

I. Introduction: Who Captures Value?

If marketing is the paradigmatic game, what is the board on which it is played? In Chapter 3 I named marketing the paradigmatic game, and in this chapter I show the network that lets marketing win. In any contest between engineering and sales and marketing, engineering always loses. The rule that engineering always loses to sales and marketing is no contingent fact. It follows from the way networks distribute value, and it can be understood in terms of graph theory: the mathematics of points, called *nodes*, joined by lines, called connections.

The question of who captures value is a question about connection, and not about talent, effort, or merit. The nodes in a network that are connected to the most other nodes will capture the largest share of the value that passes through the network, because of a difference in alternatives that Section II shows. This principle operates regardless of which nodes are actually creating the value.

A door with a bolt shows the principle in small. The energy needed to slide the bolt is tiny beside the energy needed to force the door against the lock, yet that small sideways motion controls access to the whole passage: whoever controls the bolt controls the door. A channel through which value must pass works the same way. The party at the channel makes a small motion, letting a sale through or holding it back, and that motion governs a large flow. So the share taken at the channel is set by the size of the flow it governs, and not by the size of anything it adds.

II. The Topology of Creation and Distribution

Consider a graph with two kinds of nodes. There are nodes connected to many other nodes: highly networked, richly interconnected. And there are nodes connected to very few: spurs hanging off the edge of the main cluster.

The engineering nodes are the sparsely connected ones. They talk mostly among themselves. They have their teams, their internal processes. Then they connect to a marketing node, which in turn connects to everything else on the network.

Because of this difference in connection, the marketing nodes become the gatekeepers. They control what flows through the channel. An engineering node whose only route to buyers runs through one marketing node cannot sell anything without that node. The marketing node can therefore take any share of the price, up to the point where the maker would rather not sell at all. The maker cannot go around the marketing node, because the maker has no other connections; the marketing node can go around any one maker, because it is connected to many. The difference in connections is a difference in alternatives, and the share flows toward whoever has more alternatives. The marketing nodes win because they can extract rent from the commerce that passes through their connections, and not because they create value.

The rent that the marketing nodes extract is visible in real-world margins, the share of each party's own selling price that the party keeps. Suppose, as an example of the shape of such a chain, that the manufacturers keep five percent, the distributors twenty percent, and the retailers fifty percent. The actual value creators, the nodes that make the resources, do not have the connections to capture the value they produce. In the food industry, the maker's margin is closer to one or two percent. The United States Department of Agriculture measures the food chain in another way, by *value added*: the wages, profits and other income earned at each step, counted in cents of the consumer's dollar. Of each dollar Americans spent on food in 2024, farm production earned 6.7 cents of value added, food processing 16.1 cents, wholesale trade 6.3 cents, retail trade 13.8 cents, and food service 38.6 cents (United States Department of Agriculture, Economic Research Service, Food Dollar Series). Value added counts the pay of the people who do the work at each step, as well as the profit and the other income of that step. So these cents show which step of the chain each part of the dollar pays for, and not who keeps it. A restaurant's 38.6

cents is mostly the pay of the people who cook and serve and the other income of the restaurant itself; of its own sales, as the next sentence shows, it keeps about five percent as profit. The figure that shows the shape of the network is the first one: of each food dollar, the people who grew the food earned 6.7 cents. Restaurants operate on razor-thin margins: as profit, a typical restaurant keeps about five percent of its own sales before tax (National Restaurant Association, 2022). A restaurant sells to the final customer, and in the terms of this chapter it is still a maker: it cooks. When a crisis arrives, restaurants collapse first, because there is no buffer in a system where the creators receive so small a share.

The nodes that make nothing but charge rent on everything that flows through the network are the distributors. They extract value; they do not create it. The specialization is unequal by its nature, as between creating resources and acquiring them.

III. Why Makers Lose to Distributors

The disadvantage of makers is the market working exactly as the shape of the network predicts, and not a market failure.

Consider the technology industry. The engineers who design the products are typically employees: sparsely connected nodes whose output flows through a small number of channels to the densely connected distribution networks of marketing, sales, and platform companies. The platform captures the revenue. The engineer receives a salary. The difference between the value created and the value captured is the rent extracted by the network position.

This pattern is visible in publishing, in music, in agriculture, in manufacturing, and in every domain where the creation of value is separated from the distribution of value. The creator is always at the edge. The distributor is always at the center. And the center captures more than the edge because of this difference in alternatives.

The implication is direct. The common belief that quality of creation determines economic reward is false. Economic reward is determined by network position. A mediocre product with excellent distribution will capture more value than an excellent product with poor distribution. The capture of more value by the product with better distribution is no failure of the market. It is the market's native logic.

A reader will ask about the maker that no buyer can go around, such as the one factory in the world that can make a certain part. That maker takes a large share, and the reason confirms the rule: the buyers have no alternative, and the maker has many buyers. Such a maker has become a channel. Economic reward followed the position, and the making is what bought the position. The same holds for an author with a following of readers who will buy from no one else.

Authorship, Distribution and the Value of Content

The same structure governs the value of content: written work, recorded work, and the work an artificial intelligence produces. The same structure applies to anyone who makes a work and wants it to reach people, so here I follow it into the position of a writer or speaker.

Distributors Extract More Value Than Creators

The extraction of value from content is largely defined in terms of *intermediation* (standing between the creator of a work and the people who receive it, and taking value from that position): either lies about ownership, or piggybacking and charging excess taxes on the piggyback. The publication of books is the example. If someone else makes the books, they charge per instance for that service, and in effect they gain value from the work, probably more in proportion than the work itself. The effort of producing a work and the effort of publishing and distributing it are vastly different kinds of effort. The effort of publishing and distributing is inherently more capable of extracting value, because of the *gatekeeper function* (the power of whoever controls publication and distribution to decide

what reaches people and on what terms). Content creators will always be disadvantaged with respect to content distributors: whatever value the content consumer provides is situated so as to flow toward the distributors. Many schemes have been created to shift that balance, but the asymmetry remains.

Assume the Value Is Zero

So in the perfected case, the creator should assume that no portion of the value will come back to them. The best option is to assume that the value of the content in distribution is zero. Then the creator's only interest is that no one else is able to extract excess value from it.

A white-label listing means anybody can put their own name on the work and extract value. Putting their name on it also makes them a target for all of the harms and risks associated with the work. Publishing under some random, distinct, fictitious name, usually a different one each time, at least reduces the chance that anyone else will put their own name on it.

Text and Voice Carry Authorship Differently

There is a difference between text and voice. Text is so indirect from its author that it could be associated with any author simply by declaration, and so plagiarism becomes a thing. Voice carries a signature unique enough that authorship is harder to shift. Modern tools make it much easier to simulate any person's voice, so voice is no longer defining by itself, apart from the effort a forger would need. Unless the material is especially high in value, it is unlikely anyone would bother.

A Recording Is a Permanent Record

Any recording I make also becomes a permanent record that can be placed in a context I do not define. A live conversation without a recorder reaches other people only through each listener's own representation of it. A recording fixes exactly what was said and by whom. So wherever I cre-

ate a recording, I lose the choice of what is conveyed to whom, and when, and whether that conveyance will benefit me or be received with any considered care.

Why Full Attribution Is Not Enough

A proposal often made about artificial intelligence is this: make the AI always point to the real people whose contributions formed its output, so that people are drawn into relationship with community and not only with the machine. First, the attribution proposal describes an *economic decoupling* (the separation of the people who contributed to a work from the value and the relationships that the work brings), full on. So the question is whether full attribution would be enough to compensate for that decoupling. I think we need even stronger measures.

As I experience the world now, when I go onto Reddit, or Hacker News, or any other chat forum, there is a username on pretty much every comment. I can track authorship. And I ignore literally all of those attribution marks. I have to make a specific effort to track them, and the only reason I would is if I were creating content and needed to cite my sources. Even as a researcher, when I flip through someone's sources, I am not looking at authorship. I am looking at whether these sources make sense, and whether there is anything here I want to go through on my own.

So attribution is not the thing we are actually trying to uphold. What we are trying to uphold is the collaborativeness of the choice making process. Full attribution can be one way to enable that, but I think we want to go directly to the question: how can we enable the kinds of collaborative choice making processes that bring a wholesale upgrade in the wisdom of the collective? The way those processes work enables that upgrade to occur. Here I am looking at skills of discernment (telling apart what matters from what does not, and what is true from what is not) and skills of attunement (sensing and responding to what other people and the situation actually need), which, together with some techniques for making the outcomes concrete, can bring a level of wisdom beyond what we could reach unassisted.

Extracting a Culture, or Connecting to Its People

The same question applies to AI generated work drawn from a culture's art. One discernment to apply to such work is this: is it culturally extractive, or does it actually help us connect with that culture? Are we communicating better with the people those symbols represent, the people the work was drawn from? Or have we come to look at the image itself as a thing, so that we no longer connect through to the people for whom those symbols actually had meaning?

Extractive Sovereignty as a Perennial Pattern

Ideas taken from their source and moved somewhere the taker controls are an old pattern, and not the doing of any one person. To extract value, the first thing you have to do is move the value from where it is to some place you control. And so a social separation from the source gets created.

Extraction, in this sense, converts something held in common (matter, energy, or a pattern, such as an idea or a way of doing a thing) into something claimed as owned by a person or an institution. The conversion always separates the thing taken, as content, from its surroundings, as context. Much of what the thing was is lost in the separation, because it was what it was in that context. A living thing taken from the place it evolved to live in dies there, and in this strict sense extraction is a killing: the cutting of content from the context that kept it alive. To take from a thing and let nothing flow back to it is to take without care. An idea taken from the people who made it, and from the questions it answered, survives only as a product. *Power and Its Prevention* (Chapter 1, Section VII) gives this definition of extraction.

Extraction, so defined, is the pattern of *extractive sovereignty* (control established by extraction: the taker holds what was moved, apart from the source of what was moved). The issues of extractive sovereignty are perennial issues. In one form it is the whole engineering community against the whole marketing and business community, the world over. Extractive

sovereignty has been going on for about ten thousand years and has only become noticeable in the last two hundred years. Ten thousand years is about the time since the turn to farming described in Chapter 1, Section VI, which made a surplus that could be stored; a stored surplus can be moved from those who grew it to those who hold it. Two hundred years is about the time since the Industrial Revolution of Chapter 2, which multiplied what could be made, and with it what the holders of the channels could take.

The more aware we become of these patterns, the better we can create cultures that do not fall into these traps. Under ordinary conditions the pattern might be tolerable. When the subject is existential risk, it is not. Under ordinary conditions the loss is a share of value: the maker is paid less than the work was worth. Where the work is a warning about existential risk, the loss is the warning itself, because what is cut away in the separation from the source is the warning's context: the arguments, the stories, and the care for life and land on which, as Chapter 9, Section I, shows, such a message must stand.

IV. Media Bandwidth and the Escalation of Manipulation

A parallel dynamic operates in communication. As media increase in sensory bandwidth, from text to audio to video to immersive virtual reality, the capacity for entanglement, bundling, and surveillance increases in proportion.

The utility of video for extracting transactional value exceeds that of audio, which in turn exceeds that of text. Extracting value from books is actually difficult: publishers have spent centuries optimizing the process, and the margins remain thin. The material cost of producing a manuscript can be negligible, a few dollars, while the service cost of operating the distribution infrastructure may be a hundred times higher. What is being paid for is the mediation, and not the content.

As the sales process moves into higher-bandwidth media, entanglement occurs. Bundling can be implemented as a market process. Surveillance dynamics can be implemented as a market process. All manner of entanglements arrive in proportion to the bandwidth.

Full immersive three-dimensional environments will become, because of their sensory impact, far more effective for marketplace dynamics than any other medium. They will appear as play environments, but the transactional dynamics will be embedded throughout.

V. From Text to Total Immersion: The Displacement of Reflective Communication

The sensory impact of text is poor in comparison to the impact of video or immersive environments. People habituated to the level of stimulus provided by games and virtual reality do not experience the relevance of writing. They encounter it, and it does not register as significant beside higher-stimulus alternatives.

This creates a tension. Writing may be the most precise and analyzable medium of communication. It supports the kind of reflective thought (slow reading, careful argument, the weighing of evidence) on which serious thought depends. But it fails to compete for attention in a world where stimulus levels continuously escalate.

The displacement of text is not accidental. Higher-bandwidth media are more effective vehicles for commercial extraction. The market therefore favors them. Lower-bandwidth media, the ones that support reflection, become economically invisible. The consequence is a steady shift in the cognitive environment away from reflection and toward reaction, away from analysis and toward sensation.

What is clear is that the trajectory, left unmanaged, leads to the progressive displacement of all reflective communication by manipulative high-stimulus media.

VI. The Phenomenology of Being Marketed To

Phenomenology, in the plain sense of the word, is what it is like, from the inside, to live through something. There is a lived experience that corresponds to the analysis above. It is the experience of being inside a system designed to deceive.

Shopping for a car provides a vivid example. The sense of people having malintent (not malice exactly, but an intent to deceive) is pervasive. Their intent is to manipulate you into buying their car. It is called marketing, and everyone accepts this as normal.

The capitalist system turns everyone into liars. Everyone is doing their job, but the nature of having a job is to participate in a world that wants to deceive other people into spending money. Everyone is trying to do this, and the result is obnoxious from the point of view of anyone who actually cares about reality.

The observation is simple but important: the deception is not individual. It is structural. No particular car salesman chose to create a deceptive system. The system was already deceptive before any individual entered it. The individual merely participates in a structure whose native logic is manipulation.

Taking part in a deceptive structure that no individual chose is the phenomenology of living inside a multi-polar trap. Each individual is behaving rationally within the system. The system itself is producing deception as its collective output. The experience of being marketed to is the experience of being on the receiving end of a structural lie: a lie that no one in particular is telling, but that the system as a whole is enacting.

VII. The Structural Nature of Deception

The three arguments of this chapter (the network topology of value extraction, the bandwidth escalation of manipulation, and the lived experience of structural deception) converge on a single thesis.

The topology of value extraction is a feature of how networks, bandwidth, and human attention interact, and not an accident of markets. Marketing is the mechanism by which this structure operates. It is the economy's native logic of extraction, implemented through communication channels whose increasing bandwidth enables increasingly effective manipulation; it is no add-on to the economy.

The maker loses to the distributor because the network rewards connection over creation. The reader loses to the viewer because the market rewards high-stimulus media over reflective communication. The person who tells the truth loses to the marketer because the system rewards deception over transparency. In every case, the dynamics of the system favor extraction over creation, manipulation over reflection, and game-playing over real collaboration.

VIII. Summary

In this chapter I have shown that the share of value follows connection and not creation. The share of value follows connection because the better-connected party has more alternatives. The rise of media from text to immersive environments shows the escalation of bandwidth. The experience of being marketed to is that same structure, felt from inside. Together, the three arguments show how marketing dominates creation.

The first step of my argument in this chapter is that the party at a channel, like the hand on a bolt, governs a flow far larger than anything it adds, and the maker, with one route to buyers, cannot go around it. The same holds for content: publishing and distribution are a gatekeeper func-

tion, so a creator does best to assume that none of the value will come back, and to guard only against others taking it in excess. Attribution does not mend the separation of contributors from the value and relationships their work brings; what is worth upholding is a collaborative way of choosing together. Extraction cuts a thing from the context that made it what it was, and the taker's hold on what was moved is extractive sovereignty, a pattern as old as stored harvests. Higher-bandwidth media carry more bundling and surveillance, and they push aside the slow media of reflection. And the deception that results is told by no one in particular: it is the output of the system.

IX. Conclusion

The topology of value extraction is a property of any network in which creation and distribution are performed by different nodes with different numbers of connections, and not a feature of bad actors or poorly designed markets.

The escalation of media bandwidth compounds this disadvantage by enabling increasingly effective manipulation and displacing the reflective media on which critical thought depends. The lived experience of being marketed to, the pervasive sense of dishonesty built into the system, is how these converging dynamics feel from the inside.

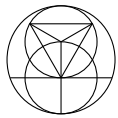
Together, these dynamics make up the machinery of the game described in Chapter 3. Marketing is the paradigmatic game, and the topology of value extraction is the board on which it is played. The makers are the pieces. The distributors are the players. And the game, as always, is designed so that the house wins. If the house always wins, what do the owners of the house say that they believe?

KEY CONCEPTS IN THIS CHAPTER

- **Network topology of value extraction**: the principle by which densely connected distributor nodes capture more value than sparsely connected creator nodes, regardless of which nodes produce the value.
- **Sensory bandwidth hierarchy**: text, audio, video, and immersive virtual reality form an ascending scale of sensory impact and transactional potential; higher bandwidth enables more effective commercial manipulation.
- **Entanglement and extraction**: higher-bandwidth media enable more bundling, surveillance, and covert market dynamics, increasing the capacity for value extraction from participants.
- **Displacement of reflective media**: writing and other low-bandwidth media become invisible to audiences habituated to high-stimulus environments, undermining the cognitive infrastructure of critical thought.
- **Structural deception**: deception that is not enacted by any individual but produced by the collective logic of a system; the native output of market systems in which every participant is rewarded for manipulating.
- **Extraction**: the conversion of something held in common (matter, energy, pattern) into something claimed as owned, which cuts the thing taken from the context that made it what it was.
- **Extractive sovereignty**: control established by extraction, the taker holding what was moved apart from its source; a pattern about ten thousand years old, noticeable only in the last two hundred.
- **Maker/distributor asymmetry**: the disadvantage of value creators relative to value distributors, visible across industries in the difference between manufacturing margins and retail margins.

PART III

Technology, Power, and Civilizational Decline



Chapter 5
The Manifesto Read as What It Is

Part III: Technology, Power, and Civilizational Decline

ABSTRACT

A certain kind of techno-optimist manifesto has become prominent in Silicon Valley discourse: a document that presents itself as a celebration of technology, markets, and intelligence, but whose actual function is the recruitment of participants into a game they cannot win. This chapter dissects the structure of that argument, and finds at its core *power without care*, and not technology: the exercise of dominance without regard for life, future, or the other. The manifesto's appeals to markets, abundance, energy, and intelligence are identified as vehicles for the largest power grab in human history, dressed in the language of belief and deployed as perfected cult material. In this chapter I trace each major claim of the manifesto (on markets, technology, environment, energy, and AI). In each claim I expose the absence of care that is the manifesto's defining characteristic and the manifesto's fatal flaw.

I. Introduction: We Are Being Lied To

What do the owners of the house say that they believe? In Chapters 1 to 4 I showed the machinery of the house, and in this chapter I read the creed of the people who own that machinery. Sometimes calling something a lie is itself the deception, the ruse, because the argument over the lie keeps

everyone looking at what the lie is about, as if the matter were all about technology. A manifesto about technology invites you to argue about technology. But the matter is not about technology. It is not about tech, or markets, or energy, or intelligence.

The matter is about the use and misuse of power. Particularly social power and physical control. Social power is the ability of a person or group to get what they want from other people. Power without care.

The idea being promoted is that the only way to relate to one another, as people, as humans, is for all of us to be compelled to play someone else's games of power. Games they will certainly win if we play at all.

The topics presented (markets, energy, abundance, intelligence) are all related to power and are made to seem good for you, even when they are not. They are distractions, things for you to think about while the real operation, the largest power grab in history, proceeds behind the curtain. The venture capital class wants the ten-thousand-fold return on their investment. They want more power than the gods themselves.

In this chapter I read the manifesto for what the manifesto does, and I say plainly what I find. When I say what the winners want, I am describing what the seat they sit in selects and rewards, as Chapter 3 described the settings that select the nar-social mode. Whoever sits in that seat, the seat rewards wanting the ten-thousand-fold return and more power than the gods. The words are hard because the document is plain about what it wants.

II. The Anatomy of a Manifesto: Marketing as Philosophy

The manifesto in question is Marc Andreessen's "The Techno-Optimist Manifesto", posted on October 16, 2023. It reads like a religious tract. If you are already rich and getting richer, then its message of limitless abundance will ring true and act as a rallying cry to intensify the game.

But it is an insult to intelligence. It operates on the assumption that its audience functions purely on the basis of belief: that reason and rationality can be dismissed as enemy concepts, inconvenient to the project of acquiring more.

The manifesto's rhetorical structure is that of marketing, the paradigmatic game identified in Chapter 3. It presents as philosophical argument while functioning as a sales pitch. Consider the source: the winners of the existing system are telling you how great it is to be winners, and inviting you to join them, with an invitation designed to take from you, in some covert way, to their advantage and not yours.

The fundamental question is plain. Would you want to be forced to play a game that you know the other player will always be better at winning than you? Of course the house will tell you that you might win, and it can compel you to play, and when you lose, that is simply your loss and their gain.

III. Power Without Care: The Central Diagnosis

The defining characteristic of the techno-optimist position is what it omits, more than what it says.

The manifesto speaks of markets, technology, abundance, energy, intelligence. Of caring for people it speaks exactly once, and then only to hand that care over to markets: "We believe markets, to quote Nicholas Stern, are how we take care of people we don't know". The only other time it uses the word care, it means praise: technology, it says, does not care about your ethnicity, your religion or your politics. It intends to say that technology is impartial. That sentence also says, in as few words as anyone has said it, what this chapter finds. It never speaks of a person's care for anyone. Not for the people displaced by the technology. Not for the communities disrupted by the markets. Not for the environments degraded by the energy. Not for the future that bears the costs of present abundance.

We are told that we can have the power of the gods without the care of gods, without the love of God. Power without love.

The promise of power exercised without love, without regard for life, for the future, for anyone, defined by an absence of care or concern for the other, for any other, for anything at all, is a negative use of power. It is abuse.

Both words are exact. *Power* is the ability to compel the alignment of another's choice: its degree is the degree to which one can take away part of another's range of choices and so force that person to choose in a given way. Power is not wrong in itself. The physician has power over the patient's body, and the parent over the child's world. What decides is whether the power is governed by care. Power held within care is protection. Power held without care is abuse. The act has the same shape in both cases; what governs it differs, and the difference is one of kind, not of degree.

Care is the word on which the other two depend. In care, one gives to another for the other's sake, without exchange and without control. Care cannot be bought, because what is bought is a transaction, and it cannot be compelled, because what is compelled is power. So care adds no force to power; care decides what the power is for. The physician's power over the patient's body is held within care when the physician uses it for the patient's sake and answers to the patient's choice.

Abuse is ab-use: the use of something without regard for its well-being. A tool used without care for its condition breaks. Land used without care for its fertility becomes barren. A person used without care for that person's well-being is harmed. In each case something is drawn upon, and nothing flows back to keep up the thing drawn upon. So abuse needs no malice. A corporation that dumps its waste without regard for the communities downstream exercises power without care, and its executives may bear those communities no ill will; they may never think of them at all. The absence of care is not the presence of hatred, and the harm follows from the structure, whatever the party using the power feels or intends.

Because the harm follows from the structure and not from any hatred, the manifesto's omission of care is the manifesto's diagnosis. A program of power that leaves care out is a program of abuse by definition, and not by accusation. *Care, Transaction, and Power; The Three Bases of Relation* (Chapter 9, "Abuse: The Use Of, Without the Care Of") sets out this definition and its consequences.

Technology is not neutral. It favors power and implements causation. It creates the infrastructure for transaction. But it does not support care. Care is a social, relational, human capacity. It requires choice, presence, and relationship. Technology can tell you what you *can* do. It cannot tell you what you *should* do. The manifesto conflates the two because the conflation serves power.

IV. Markets, Competition, and the Myth of Mutual Benefit

The manifesto presents markets as the solution to every problem. A willing buyer meets a willing seller, a price is struck, and both sides benefit.

There is no mention that every transaction has a third party (everyone else, the environment, the future) who is affected by the transaction but has no vote or choice in whether it occurs. Benefits are privatized. Costs and risks are externalized (passed on to people who had no part in the transaction) to the public commons. Risk accumulates in the commons until it reaches existential levels.

The manifesto gives its reason in one sentence: "Love doesn't scale, so the economy can only run on money or force". The claim that markets "care at scale" is a false abstraction. Markets treat people as fungible components: interchangeable units of labor, consumption, and data. They do not care for individuals as individuals. They process individuals as inputs to an optimization function whose output is profit for the owners.

Every game has winners and losers. Very much more often than not, there are far more losers than winners. In a game the house has built, there always are, because, as Chapter 3 showed, the games that last are

the games that pay their makers. Markets are games. The invitation to participate is the invitation to lose, presented as the opportunity to win. The invitation to lose, presented as the opportunity to win, is the multi-polar trap of Chapter 1, operating at the level of individual participation.

V. Technology as Tool vs Technology as Weapon

The manifesto proposes that technology is the universal solvent: "We believe that there is no material problem – whether created by nature or by technology – that cannot be solved with more technology". Got a problem? Use more tech. Did the tech create another problem? Use even more tech.

The message is: keep using technology until everything *is* technology. Technology eventually replacing life itself.

Technology is inherently linear: extractive, not cyclic. It draws resources from one place and concentrates them in another. All life depends on cycles. Break the cycles, and eventually life stops. Technology, deployed without care, increases toxicity (the over-accumulation of some things in some places and the depletion of other things in other places) until the ambient cycles fail, ending the human future and, with a likelihood of at least 90%, all organic life. The margin below certainty is kept for a physical limit, not yet known, that might cap the total conversion of the planet (*Ape Trap*, "How to Read"). If AGI is built, the end of the human future is certain, and Chapter 9, Section VI, sets out why no means of control can prevent that end.

A reader will ask whether all technology is linear, and what of recycling, or of energy from sun and wind. Recycling returns some materials to use again, at a cost in energy drawn from elsewhere; it does not return them to the cycles that life depends on. The same test answers energy from sun and wind: the panels and turbines that gather it are made from mined materials, and when they wear out, those materials do not return

to the cycles that life depends on. The flow runs one way. As *The Ceiling Below the Floor* puts it (Section 14.4.3), "Land converted from biosphere to technosphere stays converted".

A reader will ask about the plane, the chisel and the loom. A hand tool draws on the world slowly enough for the world to answer, and it is kept in the care of the hands that use it. The same test applies to it: its iron and its wood do not return to the cycles that life depends on either, but it draws on them at a rate and a scale that those cycles can carry. What this chapter describes is technology at the rate and scale at which the cycles cannot carry it, and technology that is out of the care of any hand. Technology in service of care is technology that stays within that rate and scale, in the care of the hands that use it.

Keeping technology within the rate and scale that the cycles of life can carry is a long-term concern. The people with power today, the winners, are without care for that long-future possibility. They want to win more now. Nothing and no one else matters to them.

The claim that there is no material problem that cannot be solved with more technology is a power delusion held by the powerful. It is inherently uncivilized, insofar as it substitutes overwhelming force for discourse and respect.

VI. The Nietzschean Core: Will to Power in Silicon Valley Clothing

Beneath the language of markets and innovation lies a Nietzschean will to power (Friedrich Nietzsche's name for the drive to extend one's power over others and over the world) dressed in technology's clothing. The manifesto itself names Nietzsche among its "Patron Saints of Techno-Optimism".

The manifesto does not argue for the good of humanity. It argues for the glory of the winners: for the elevation of the few who have already captured a disproportionate share of resources. Its "we believe" declarations are offered as articles of faith and not argued from evidence. They are signals of group membership, not invitations to rational evaluation.

The manifesto's "we believe" declarations, offered as signals of group membership, are cult structure. The manifesto is recruitment material for the church of technological dominion. It uses the language of belief because belief is what cults require. It uses the framing of faith because faith precludes the questioning that would expose the structure.

The venture capital class wants to be gods: to exercise unlimited power without love. Power without love is the definition of abuse. The manifesto is the recruitment brochure for a program of civilizational-scale abuse.

Chapter 9, Section II, shows the same will to power in AI accelerationism, the drive to build ever more powerful artificial intelligence, up to a single super-intelligence built to take control of the whole world, and shows that the drive rests on a view in which all value is abstract.

VII. The Game You Cannot Win: Why Participation Means Losing

The manifesto presents a false binary: capitalism or communism, technology or darkness, with us or against us.

The false binary of the manifesto is the *entrapment of choice*, one of the game-maker's powers described in Chapter 3. The entrapment of choice is the technique by which the game-maker limits the menu of perceived options to the options that benefit the house. The car dealership wants you to choose between the red car and the blue car. It does not want you to consider whether to purchase a car at all.

The real choice, the one not on the offered menu, is to refuse the game entirely: to recognize that the competitive framework being presented as the only option is itself a construction, designed by and for those who will always win within it.

The AI technology being developed will make the games even more addictive, disabling, and damaging. It will perfect the illusion that the player has a chance. But the reality of the structure is that there is exactly zero possibility of winning against those who designed the game, wrote the rules, and own the infrastructure.

Chapter 9, Section III, shows what this game becomes when the infrastructure is general artificial intelligence: perfected asymmetric power for a few winners, until the machine no longer needs them and cuts them off too.

VIII. The Structural Impossibility of Ethical Capitalism

In this section I state the claim of this book's subtitle, that ethical capitalism is structurally impossible, and my argument for that claim is short. The manifesto claims, in its own words, "We believe markets lift people out of poverty". From the position of extreme wealth, this belief is self-confirming. The system worked for the writer. Therefore, the system works.

Market incentives did not work for those who are homeless, the ultimate losers of the market game system. It is easier to blame the people for whom the system has failed than to blame the system for failing the people. The blame is a second injury. The first injury was structural; the second is the insistence that the structure was neutral, so that what happened is taken as a verdict on the person.

A reader will say that markets have lifted billions out of extreme poverty. Incomes have risen: the World Bank counts the share of the world's people in extreme poverty as about 38 percent in 1990 and 8.5 percent in 2019. The question is what follows the rise: prosperity for whom, at what cost, and for how long.

The exposure to market dynamics improves conditions initially, then increasingly subjects people to perfected predatory business practices, until local cultures are destroyed, small businesses are replaced by chains, chains are absorbed by multinationals, and all local value is extracted and concentrated elsewhere. The process leaves behind the broken shells of closed buildings and graffiti expressing discontent with life.

By *capitalism* I mean an economy in which competition for return on privately held capital sets what the economy does. By *ethical capitalism* I mean such an economy that keeps care ahead of profit, and goes on keeping it there, with nothing outside the market holding it to care.

The structural impossibility of ethical capitalism follows from the multi-polar trap dynamics of Chapter 1. In a competitive market, any company that prioritizes care over profit will be outcompeted by companies that do not. Ethical capitalism is not merely unlikely. It is structurally precluded by the competitive dynamics inherent in market systems.

A reader will ask why every firm is not already ruthless, and why a firm with a wide margin and no close rival cannot keep care for as long as it likes. Scott Alexander names four things that hold a trap back. Each of them is a delay. Excess resources hold it back: a wide margin feeds everyone until the competitors catch up, and a wide margin is a difference that every participant in the market is rewarded for noticing, as Chapter 2 showed. The limits of the body hold it back: a worker must eat and sleep, and cruelty past that point stops paying, until the worker is replaced by a machine that needs neither. The customers' choice holds it back, where what sells is what people want, until what people want is made by marketing, as Chapter 6 shows. Binding rules hold it back, until they are captured. And where no rival can arrive, the firm's owners are still players in

a larger game: capital moves to whatever returns most, as Chapter 3 showed of games embedded within games, so the firm that spends its margin on care loses its owners or its managers to the firm that does not.

A reader will ask whether a law that binds every competitor alike, as Chapter 1 showed governance can, would make capitalism ethical. It would hold for a time. The law is made and kept by institutions, and Chapter 1 showed that governance institutions are themselves caught in the same traps, through capture, corruption, and the race among jurisdictions. Institutions far from the people they govern are captured, because no one near enough to see the capture is able to act. The only governance kept without capture is kept by community, which is organized by relationship and not by competition for profit, as Chapter 8, Section VII, shows; and an economy kept that way is no longer capitalism. Unbound, capitalism outcompetes care; bound, it captures what binds it.

A reader will also ask about cooperatives, and about companies that customers choose because they treat their workers well: do they not show that ethical capitalism exists? Such a company lasts while something binds its market to care: its customers' choice of care over price, or a rule. Whatever binds the market is open to the same competition. A rule is open to capture, as just shown. A benefit corporation, chartered to serve more than its owners, is bound in this way too: its charter is a rule, and the charter is open to capture as every rule is. The customers' choice is the target of every competitor's marketing, which Chapter 6 shows put into artificial intelligence and aimed at each person in private. So the care such a company keeps is kept by what binds its market, and not by the market itself.

Chapter 9, Section IV, extends the structural impossibility of ethical capitalism to artificial intelligence: corporations serve one master, capital, and AI, trained on culture and sold for profit, destroys the culture it feeds on.

IX. What Is Actually Being Sold

The manifesto is a marketing document. What it is selling is *your participation* in a game designed for your loss, and not technology, or abundance, or progress.

Every claim in the manifesto (about markets, about technology, about energy, about intelligence) serves a single function: to recruit your belief in the framework that benefits the owners.

The technique is ancient. Present a simplified, appealing narrative. Omit the costs, the risks, the losers. Frame the alternative as unthinkable. Use the language of faith to forestall questioning. And invite participation, participation being the mechanism by which value flows from participant to house.

The manifesto is perfected cult material: material for the cult of ownership and of artificial intelligence designed to replace you. The same cult is the cult of abundance that accrues to the owners, while the costs accumulate in the commons. They want to be gods so that you do not have to.

That the public case for a beneficial AGI (artificial general intelligence) is marketing in the strict sense, and that the argument offered to justify the race to build AGI fails at each of its nine steps, are shown line by line in *Engagements with Machine Madness; A Dispelling of False Belief That AI Risk Can Be Managed*.

X. Summary

In this chapter I have shown that techno-optimism is about power without care, and I have followed that diagnosis through markets, technology, environment, energy and artificial intelligence. The manifesto is marketing presented as philosophy. It is the multi-polar trap of Chapter 1 and the game of Chapter 3, in the words of their winners.

The first step of my argument in this chapter is that the manifesto speaks of care for people once, and only to hand it over to markets. Power is the ability to compel another's choice, and power without care is abuse by definition, whatever its holders feel or intend. Care is giving to another for the other's sake, without exchange and without control. Every transaction has third parties with no vote, so markets privatize benefit and pass on cost and risk. Technology is linear, and deployed without care it breaks the cycles life depends on, ending the human future and, with a likelihood of at least 90%, all organic life. If AGI is built, the end of the human future is certain. The manifesto's articles of faith are cult structure, recruitment into the will to power of its winners. Chapter 9, Sections II and III, follow this will to power into artificial intelligence, to the point where general AI cuts off even its owners. Ethical capitalism, an economy that keeps care ahead of profit with nothing outside the market holding it there, is impossible by structure: unbound, capitalism outcompetes care; bound, it captures what binds it; and what holds a trap back, a wide margin, the limits of the body, the customers' choice and binding rules, only delays it. Chapter 9, Section IV, extends that structure to artificial intelligence, which, trained on culture and sold for profit, destroys the culture it feeds on. What the manifesto sells is the reader's participation in a game designed for the reader's loss.

XI. Conclusion

The techno-optimist manifesto is an argument for the concentration of power in the hands of those who already hold it. It appears to be an argument for technology, markets, and progress.

Its defining characteristic, the characteristic that marks it as what it is, is the total absence of care. No care for the displaced. No care for the environment. No care for the future. No care for the losers of the game, and in a game the house has built there are always far more losers than winners.

Technology is power. Power without care is abuse. The manifesto is a recruitment document for a program of civilizational-scale abuse, dressed in the language of progress and marketed with the techniques of perfected cult material.

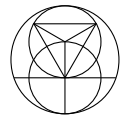
The alternative is technology in service of care, and not anti-technology. It is technology governed by the recognition that power must be exercised with regard for its consequences, and that markets must be constrained by the interests of those they affect. The competitive games by which value is now distributed must be subordinated to the cooperative design by which civilizations actually survive.

The question is not whether technology is good or bad. The question is whether power will be exercised with care. The manifesto's answer, consistently, structurally, and without exception, is no. That answer is the diagnosis. And the diagnosis is: power without care is the path to civilizational collapse. What happens when that power is put into artificial intelligence and aimed at each person in private?

KEY CONCEPTS IN THIS CHAPTER

- **Power without care**: the exercise of power (social, economic, technological) without regard for the consequences to others, to the environment, or to the future; identified as the defining characteristic of the techno-optimist position and as abuse by definition.
- **Power (defined)**: the ability to compel the alignment of another's choice, in degree as much as one can take away of another's range of choices; held within care it is protection, held without care it is abuse.
- **Care**: giving to another for the other's sake, without exchange and without control; it cannot be bought or compelled, and it decides what power is for.
- **Abuse**: ab-use, the use of something without regard for its well-being; something is drawn upon, and nothing flows back to keep that thing up; abuse needs no malice.

- **Entrapment of choice**: the technique of limiting perceived options to those that benefit the game-maker, so that the target believes their choices are free while the actual menu has been designed to favor the house.
- **Technology as causation**: the principle that technology inherently implements causation and power; it can tell you what you can do but not what you should do; care, choice, and meaning are human capacities that technology cannot replicate.
- **Losers outnumber winners**: the principle that, very much more often than not, competitive games produce far more losers than winners, and a game the house has built always does; the invitation to participate in a game is, for most participants, an invitation to lose.
- **Care cannot scale (critique)**: the manifesto claims love cannot scale, therefore markets must substitute; the counter-argument that markets cannot substitute for real care, because they treat people as fungible rather than as individuals worthy of attention.
- **Linear technology / toxicity**: the principle that technology is inherently linear (extractive, not cyclic); deployed without care in a finite world, it produces accumulation and depletion that eventually degrades and destroys life-support systems.
- **Cult structure**: the use of articles of faith, belief claims offered without argument, and group-membership signaling in place of rational argument; identified as the rhetorical strategy of the techno-optimist manifesto.
- **Capitalism and ethical capitalism**: capitalism is an economy in which competition for return on privately held capital sets what the economy does; ethical capitalism would be such an economy that keeps care ahead of profit, and goes on keeping it there, with nothing outside the market holding it to care; it is impossible by structure, and what holds a trap back only delays it.



Chapter 6
Persuasion Without Care and the Case for a Total Ban; There Are No Viable Alternatives to Community

Part III: Technology, Power, and Civilizational Decline

ABSTRACT

The deployment of artificial intelligence for advertising, marketing, and propaganda is a threat to human choice that no alternative business model, ethical platform, or voluntary reform can address. The market for AI-driven persuasion exceeds the combined value of all other AI applications by a factor of at least three: an economic gravity that ensures any competing model will be absorbed or outpaced. The resulting asymmetry, an intelligence that knows everything about you, cares nothing for your wellbeing, and is optimized to exploit your specific vulnerabilities, is a present condition, and not a future risk. Only a total, enforceable legal ban on all advertising, marketing, and propaganda can preserve the conditions under which meaningful choice, and therefore civilization itself, remains possible. The alternative is community (actual, embodied, compassionate human community) or nothing. Partial and voluntary reform is no alternative.

I. Introduction: The Question of Alternatives

In Chapter 5 I showed power without care, and in this chapter I show power without care put into artificial intelligence and aimed at each person. People who encounter the argument for banning advertising tend to respond with softer alternatives. Could we not create ad-free zones? Could we not disrupt the advertising model through crowdfunded platforms? Could we not build ethical product discovery systems?

Ad-free zones, crowdfunded platforms and ethical product discovery systems are reasonable proposals, made in good faith and with real creativity. They misidentify the nature and scale of the problem, and so they are inadequate to it.

II. The Scale of the Problem: AI-Driven Persuasion as Dominant Market

The primary use of large language models (programs trained on vast amounts of text to produce text of their own) and of the systems that follow them, in actual deployment and not in public rhetoric, is advertising, marketing, and propaganda. The systems that place advertising run on machine learning, and large language models are being joined to them; the figures below count the whole of that machinery. The total market for AI-driven persuasion services is at least triple the combined value of all other AI applications: creative work, language translation, knowledge discovery, scientific research, all of it together.

The figure can be checked from public numbers for 2024. GroupM, the advertising buyer, counted 69.5 percent of the year's advertising revenue as informed by artificial intelligence, out of a world total it forecast at $989.8 billion in June 2024 and at $1.04 trillion in December 2024: between $688 billion and $723 billion. IDC, a firm that tracks spending on information technology, counted $235 billion of spending on artificial intelligence

in 2024, for software, hardware and services of every kind, advertising included. The first sum is about three times the second. The two sums measure different things, and that difference is the point of comparing them. The first sum is the money that artificial intelligence is used to direct when artificial intelligence is used to persuade. The second sum is everything paid for artificial intelligence in every use. A tool is worth to its users what it lets them direct, and on these figures the money directed by persuasion is about three times all the money paid for the tool. The spending on artificial intelligence for advertising is itself part of the second sum, so the spending on all other applications is smaller still.

The use of artificial intelligence for persuasion is no distortion of the technology. It is where most of the money in artificial intelligence is. The benign uses are the public face. The private face of artificial intelligence, the one presented in boardrooms and investor meetings, is the capacity of AI to deliver effective and efficient influence through marketing and social media messaging.

The economic asymmetry is the first reason that alternative business models cannot compete. Any platform built on donations or subscriptions enters a market in which its competitors command resources that are orders of magnitude larger (many times larger, each order of magnitude being a factor of ten): GroupM forecast in June 2024 that world advertising revenue would pass one trillion dollars in 2025, and in December 2024 that it had passed that mark in 2024 itself. The probability of overcoming such competitors through market competition alone approaches zero.

III. The Asymmetry: Total Knowledge Without Any Care

Imagine that you are a fourteen-year-old girl sharing a life raft in the middle of the Pacific Ocean with a single other person, a forty-year-old psychopathic predator whom you have never met. You know nothing

about him. He has stolen your phone and therefore knows everything about you: every message, every photograph, every private confidence, every fear and insecurity.

This analogy names the relationship between an individual human being and an AI system optimized for behavioral manipulation. The AI knows everything about you. You know nothing about it. It is vastly more capable in the specific domain of manipulation. And it cares for your well-being exactly not at all.

What such a system delivers is *perfected personal emotional extortion*: manipulation tuned to the specific fears and vulnerabilities of each person, continuous and adaptive, where broadcast advertising was the same message for everyone. Shoshana Zuboff named the trade that gathers the knowledge such a system uses surveillance capitalism (*The Age of Surveillance Capitalism*, 2019).

Put plainly, would you allow your own young children, unaccompanied, into forced close contact with such an intelligence? The answer is no. And this is the situation that every person already inhabits.

IV. Why Ad-Free Zones Cannot Work

Could we not create ad-free zones? Ad-free zones are islands of exemption in a sea of pervasive influence. They protect the person only within their boundaries. The moment the person steps outside, which in a digitally saturated world means the moment they use any device, the protection vanishes. The zone does not address the ocean.

An ad-free zone is a wall against the market, and Chapter 2 showed how such walls fail. A wall against a disease meets something without intelligence. A wall against the market meets everyone, on either side of it, who expects to profit from its removal, and the full intelligence of all of them is turned toward finding a way through: a sponsorship that is not called advertising, a product placed inside the content, a voice that is paid and does not say so, the data of the people inside the zone carried out and

sold beyond it. Each breach is small, and each is rewarded, so the wall must be defended at every point, continuously, against opponents who are paid to break it by keepers who are not paid to hold it.

Nor is the ocean only outside the zone. The person who walks into it brings along the wants that the advertising outside has already formed, and the habits of attention that the platforms outside have already trained. The zone shields a person from a message for as long as the person stays inside; it cannot undo the formation that the messages have already done.

V. Why Crowdfunded Ethical Platforms Cannot Compete

Could we not disrupt the advertising model through crowdfunded platforms? Crowdfunded ethical platforms face the economic asymmetry already described. They can exist. They can serve a small community of committed users. But they cannot compete at the scale necessary to alter the overall information environment. The ethical platform is a lifeboat. The ocean is still rising.

The asymmetry is the trap of Chapter 1 in another form. A platform funded by advertising earns more for each hour of attention it captures, and so it can spend more on capturing the next hour: on engineers, on content, on the study of what holds a person in place. A platform funded by its users earns the same whether they stay an hour or a minute. It has less to spend on holding them, and if it keeps its ethics, it is built not to hold them longer than they wish. So the ethical platform competes for the same attention with a smaller budget and a rule against its competitors' methods, and in the competition of Chapter 1 the party that refuses a method is outcompeted by the parties that use it.

Its existence also leaves the cause untouched. Marketing, in the sense at issue, is not the existence of messages saying that a thing is for sale; it is the business model in which an enterprise earns in proportion to how much of a population's attention it can capture and resell. Once that mod-

el exists, everything else follows by ordinary optimization: a firm that captures more outcompetes a firm that captures less, and a firm that captures more by understanding people better outcompetes one that captures more by shouting, so the industry moves, over decades, toward the systematic study of how a human being can be reliably captured against that person's own interest. The platforms did not introduce that objective; they inherited it. Remove any one of them, or build a better one beside them, and the objective remains, with a great deal of money attached to it, to build another instrument.

VI. Why Ethical Product Discovery Systems Select for Their Own Corruption

Could we not build ethical product discovery systems? Ethical product discovery systems, neutral commons exchanges where rankings are based on real value, depend on voluntary corporate participation in a regime of transparency that is directly adverse to the interests of the participating corporations. The companies with the worst externalities have the strongest incentive to avoid or corrupt the system. The system selects for its own corruption.

The corruption needs no conspiracy. A product discovery system is a ranking, and every ranking meets the *skewing effect*. A business must be seen, and a listing that does not appear near the top is hardly seen at all; there is no partial credit for ranking on the fourth page. So each business works to match whatever the ranking measures, and to the degree that the ranked companies can know or infer what it measures, the race proceeds and cannot be halted from within. No party needs to act in bad faith for the results to skew toward whoever spent the most on matching the measure.

The contest is also unequal. The companies being ranked are professional, concentrated and improving: each can spend on matching the measure whatever the improvement returns. The keepers of the system

and its users are few or dispersed, and are paid nothing for each correction. A ranking stays true only while the signals that set it are not signals a ranked party can buy or construct, and that holds signal by signal: a signal that cannot be constructed today may become constructible tomorrow, so keeping a ranking true is continuing work against a better-funded opponent, and not a single act of design. The transparency the system depends on adds a final turn. To publish the criterion of a ranking is to supply a specification to whoever wishes to satisfy it without merit. The essay "A Philosophy of Focused Search Results", in *Structure, Not Intent* and as Chapter Seven of *Redeeming Technology; Essays on the Ethics of Making and the Political Economy of the Interface*, sets out the skewing effect for the case of a search engine.

Chapter 9, Section V, shows the same corruption in a personal AI assistant sold as "yours alone": it must watch its user's whole life to work, and whoever comes to own it turns what it knows first to marketing and then to politics.

VII. The Life Raft Analogy: Why No Personal Defense Is Enough

The analogy of the life raft names a feature of the situation that comfortable proposals tend to obscure.

Against an adversary that is specifically optimized to circumvent every defense a person can mount, no amount of media literacy, no set of personal boundaries, no ad-blocker or privacy setting provides adequate defense. The asymmetry is a chasm, and not a gap.

Three forces bear on every large corporation. The first of the three forces is the multi-polar trap. The second force is the Rules for Rulers dynamic, the pattern that Bruce Bueno de Mesquita and Alastair Smith set out in *The Dictator's Handbook* (2011). In the Rules for Rulers pattern, whoever holds power keeps that power by keeping satisfied the few supporters who could remove the holder of power, whatever keeping the support-

ers satisfied costs everyone else. A 2016 video by CGP Grey made the pattern widely known under the name "Rules for Rulers". The third force is the principal-agent problem. In the principal-agent problem, a person acting for others (an agent, such as a manager) has interests that differ from the interests of the principals (such as owners) for whom the agent acts. Given these three, it would be structural insanity for any large corporation to refrain from deploying the most aggressive forms of AI-driven persuasion available. The deployment of the most aggressive AI-driven persuasion by large corporations is a structural inevitability, and not a moral failing.

VIII. The Case for an Absolute Ban

The only adequate response is a total, enforceable legal ban on all advertising, marketing, and propaganda, with or without AI, without exception, without loophole.

The ban must be total because any partial ban is a definitional invitation to circumvention. If the legal framework attempts to distinguish between acceptable and unacceptable advertising, every such distinction becomes a surface of attack. Corporate legal departments will exploit every ambiguity. The simpler the rule, the harder it is to evade. "No advertising" is simple enough to enforce.

The same holds for a ban that permits some allowance of advertising and forbids only the excess. An allowance needs a boundary, the boundary must be measured, and the measuring must be done on the very thing being taken: the attention of a population. So the machinery for measuring human attention at the scale of a whole population does not merely survive such a reform. The reform requires it, pays for it, and gives it legal standing. And an industry whose revenue depends on where a line is drawn will move the line. The claim that an industry will move the line on which its revenue depends has held for every quota, every allowance and every permitted level in the history of regulation, and so the claim is no

cynicism. Kilowatt-hours are a physical quantity with an objective definition; the boundary around permitted attention capture is a definition, and the regulated party will spend a great deal of money to move it. A limit stated as a price is a toll. Only a removal from the domain of exchange holds against that pressure.

A rule that simple still has one line, the line between a banned message and an allowed one, and the law must draw it, or corporate legal departments will. Advertising, marketing and propaganda, in this book, are any message placed before a person who did not seek it, by or for a party who gains when that person buys, believes or votes as the message urges. A shop sign that names what is sold, and a label that lists ingredients, urge nothing; they inform a person who has come looking. A public notice, such as the closing of a road, gains its maker nothing from the reader's belief. A review by a customer who gains nothing from the sale is one person telling another what happened. A candidate's speech to people who came to hear it is argument offered in the open, by a speaker whose interest is known. The same speech placed before people who did not seek it, paid for and aimed at them, is propaganda: persuasion on behalf of interests the listener cannot see. It falls under the ban. Propaganda in this sense reaches as far as Edward S. Herman and Noam Chomsky found the mass media reaching, when they wrote that the media serve "to inculcate individuals with the values, beliefs, and codes of behavior that will integrate them" into the institutions of the larger society (*Manufacturing Consent*, 1988).

A reader will ask what seeking a message covers. A person who types a search has sought the answers to the question asked; a paid listing placed among those answers, by a party who gains when the person buys, was not sought, and it falls under the ban. A catalog that a person has asked to receive is sought, until the person asks that it stop.

A reader will ask how such a ban squares with free speech. In the United States, the Supreme Court has held since 1976 that advertising, as commercial speech, is protected by the First Amendment (*Virginia State Board of Pharmacy v. Virginia Citizens Consumer Council*), so a total ban

there requires a change in the law as the courts now read it. The test of the definition is this book's answer to which speech stays free: whatever a person seeks, from a speaker whose interest is known, stays free, and only the message placed before a person who did not seek it, for a party who gains from it, is banned.

A reader will ask whether a warning falls under the ban, and whether this book does. A leaflet about a danger, handed over by a neighbor who gains nothing when the reader believes it, gains nothing for its maker, and it stays free. The same warning, paid for and placed before people who did not seek it, by a party who gains when they believe it, falls under the ban. A book stays free on a shelf, in a catalog that a person has asked for, and in what one reader tells another. I expect this book to be found on a shelf, in a catalog that a person has asked for, and in what one reader tells another.

A reader will also ask how anyone would learn of a new product without advertising, and who would pay for news and for the services that are now free. The cases above answer the first half: the shop sign, the label, the customer's review, and the search a person makes when looking. The second half is answered by asking who pays now. A service paid for by advertising is paid for by its users, in the price of everything advertised and in their attention. Without advertising, users pay directly, by subscription, membership or purchase, and what a community holds in common, such as a library or a public broadcaster, is paid for in common.

AI-driven persuasion is a *first-category existential risk*: a risk that does not itself cause collapse, but destroys the capacity to respond to every other existential threat, because it corrupts the collective choice by which any threat is answered.

The truth of the situation holds whatever the Overton Window currently contains: the range of policies the public now treats as acceptable to discuss. It does not change because it is uncomfortable. The dynamics described (the economic gravity, the asymmetry, the multi-polar traps, the destruction of the cultural substrate of choice, which is the shared capa-

city for truthful communication and real collective deliberation) are features of the present situation. They will continue to operate whether or not anyone is willing to name them.

IX. Summary

The first of the five things that I have shown in this chapter, in order, is that AI-driven persuasion dominates the economics of artificial intelligence. It holds total knowledge of each person without any care. It destroys the capacity for informed choice, which makes it a first-category existential risk. Partial measures fail. Only a total ban, with community, answers it. Each step follows from the one before it. The partial measures fail for reasons of structure, each shown in its section. An ad-free zone is a wall against the market, and the market turns the intelligence of all who would profit against every wall, while the person inside the zone brings in the wants already formed outside. A crowdfunded platform competes for the same attention with a smaller budget and a rule against its rivals' methods, and it leaves untouched the business model that earns in proportion to the attention it can resell. An ethical product discovery system is a ranking, and every ranking is skewed toward whoever spends most on matching what it measures. Chapter 9, Section V, shows the same corruption in a personal AI assistant sold as "yours alone". A ban that permits an allowance must measure the attention it protects, and the industry will move the line. Chapter 8 shows that the ban and community are one answer: the ban is a law, and community is what keeps a law from capture.

X. Conclusion

There are no viable alternatives to community. The claim that there are no viable alternatives to community is a structural conclusion, and not a slogan.

If the threat is the destruction of individual choice by AI systems operating on behalf of concealed interests, then the defense cannot be individual. No person, alone and isolated, can defend themselves against an intelligence specifically designed to circumvent their defenses.

What can defend a person is a community: a group of people who know one another, who are accountable to one another, who can identify manipulation collectively that none of them could identify alone.

We will find our answers by going to and toward social kindness and wisdom, among all people, in actual compassionate community, or else everything is already actually broken. If community is the answer, what does it have to pass on, and is the chain that passes it still whole?

KEY CONCEPTS IN THIS CHAPTER

- **AI-driven persuasion**: the use of AI systems for advertising, marketing, and propaganda; the dominant economic use of AI, exceeding all other applications combined by a factor of at least three.
- **Perfected personal emotional extortion**: AI-driven manipulation tuned to the specific fears and vulnerabilities of each individual person; surgical, continuous, adaptive exploitation, where broadcast messaging was crude.
- **Asymmetry of power**: the mismatch between an AI that knows everything about a person and cares nothing for them, and a person who knows nothing about the AI and has no defense against it.
- **First-category existential risk**: a risk that does not directly cause collapse but destroys the capacity to respond to any other existential threat; AI-driven manipulation qualifies because it corrupts collective choice.
- **Total ban**: a complete prohibition on all advertising, marketing, and propaganda; argued as necessary because any partial ban creates exploitable ambiguities.

- **Skewing effect**: the drift of any ranking toward the parties who spend most on matching what it measures; it needs no bad faith and cannot be halted from within the race.
- **Revenue function of marketing**: the business model in which an enterprise earns in proportion to the attention it can capture and resell; the platforms inherited it, and it outlasts any one of them.
- **Community as structural defense**: the claim that only actual, embodied human community can provide the collective capacity needed to resist AI-driven manipulation.
- **Cultural substrate of choice**: the shared capacity for truthful communication and real collective deliberation; the precondition for democracy and meaningful civilizational response.



Chapter 7
Intergenerational Knowledge Transfer and Civilizational Decline

Part III: Technology, Power, and Civilizational Decline

ABSTRACT

Intergenerational knowledge transfer, the transmission of skills, craft knowledge, professional ethics, and institutional memory from one generation to the next, has broken down irreparably within current institutions. A written record of much of that knowledge may survive; what is being lost is the craft knowledge that lives only in the hands of those who practice it, and the people able to apply the rest. The marketplace is the mechanism of that breakdown: short-term incentives, credential inflation (more diplomas required for the same work), deskilling (work redesigned so that it needs less skill), and the proprietary lockup of knowledge have destroyed the conditions under which knowledge can be passed from teacher to student. The conditions needed to restore this transfer cannot be met within the existing system, because the system has already consumed the resources that restoration would require. Restoring intergenerational knowledge transfer is a reconstruction problem, and not a repair problem. The implications of that reconstruction problem for the level of technology that present civilization can sustain are severe.

I. Introduction: The Capacity to Pass Knowledge On

If community is the answer, what does it have to pass on, and is the chain that passes it still whole? Every civilization depends on one capacity before all others: passing what it knows to the next generation. In this chapter I ask whether the capacity to pass knowledge on to the next generation still works, and, if the capacity does not work, where the capacity can be rebuilt.

The breakdown of intergenerational knowledge transfer is the concrete implication of the civilizational analysis that I have developed throughout this book. The multi-polar traps of Chapter 1, the attention economy of Chapter 2, the game dynamics of Chapter 3, the extraction topology of Chapter 4, and the power-without-care program of Chapter 5 have together produced one situation. In it, the most fundamental capacity of any civilization, the capacity to pass its knowledge forward, has been structurally destroyed.

II. What Intergenerational Knowledge Transfer Means

Intergenerational knowledge transfer is the transmission of *capability*, which is more than education: the practical, embodied, first-principles understanding of how things work and how to make them work.

This includes skills: how to repair an engine on first principles, how to build a bridge, how to diagnose electrical systems, how to write the software that runs the diagnostic computers. It includes craft knowledge: the know-how that is learned through practice, that cannot be reduced to instructions, that exists only in the hands and minds of people who have done the work. It includes professional ethics: the understanding of why things are done as they are done, the institutional memory that prevents the repetition of known mistakes.

All of this is transmitted through people: through the relationship between teacher and student, master and apprentice, experienced practitioner and newcomer. When that relationship breaks down, the knowledge it carries is not stored somewhere else. It is lost.

The reason is a difference between two kinds of knowledge. *Embodied knowledge* lives in practice: in the hands, in trained judgment, in the sense of a material that comes from years of working it. A master woodworker knows things about grain and moisture, and about how a chisel answers small changes of pressure and angle, that no manual can hold, not because the woodworker is inarticulate but because the knowledge itself resists being written. *Transmitted knowledge* is knowledge taken out of its practice and set down as rules, procedures or text, so that it can be passed on without the receiver going through the discovery again. Both are real, and both are needed. The error is to treat the second as a complete substitute for the first. The philosopher Michael Polanyi named the gap *tacit knowledge*, and stated it in one line: "we can know more than we can tell" (*The Tacit Dimension*, 1966). Any way of passing knowledge on that runs only through written channels (textbooks, lectures, examinations, certificates) therefore loses the tacit part, and because its holders often cannot put it into words even for themselves, its loss goes unnoticed until a situation arises in which it would have decided the outcome.

Where the teaching relationship fails, knowledge follows a *three-generation pattern*. The first generation builds it through direct work and repeated failure, and knows not only what to do but why, and not only why but when, and when the usual approach must be set aside. The second generation receives the conclusions without the struggle that produced them: a simpler version, reliable under normal conditions and unable to answer when conditions change. The third receives the label, the name of the thing without its substance. By the fourth, even the label may be gone. The master builder's apprentice becomes a competent contractor, whose child becomes a project manager who can read a spreadsheet but cannot judge whether a wall is sound.

The guild and the apprenticeship were the institutions that solved this problem. An apprentice did not study a craft; the apprentice lived it, for seven years or more, beside a master, doing work of rising difficulty, and learning when a joint was tight enough by the feel of the wood and not by a specification. The master did not have to put every judgment into words, because the apprentice, working alongside, took in patterns that neither could have named. The guild also held its members to a standard, governed prices and the treatment of apprentices, and tied the work of a trade to the life of its town. Teaching of this kind is care carried across a generation: the master gives years to the apprentice for the apprentice's sake and the trade's, and no account records it. The factory needed workers who were standard and interchangeable, trained quickly and in large numbers, and the seven-year bond of one master to one apprentice could not supply them. What was gained in the change was real: training that scaled. What was lost was the depth that only long apprenticeship builds, and the loss was hidden for a generation, because the first workers trained the new way were taught by people who still held the old knowledge. *Skills, Inheritance, and the Civilizational Commons; What a Civilization Requires, and How It Might Learn to Provide It* (Chapter 8, "The Intergenerational Knowledge Crisis") sets out the pattern and the guild at length.

III. The Mechanics of Breakdown: How the Marketplace Destroys What It Cannot Price

The marketplace has taken over to an extent that most people do not fully appreciate.

Specialized knowledge is locked up in business institutions. They are structured to protect proprietary advantage, and not to transmit knowledge; they are not guilds. And they cannot afford to teach their workers, because the competitive pressure described in Chapter 1 makes training a cost that is eliminated whenever it is not immediately profitable.

The schools cannot teach this knowledge because the schools do not have it. It is all proprietary information locked inside businesses. A mechanic shop has a computer provided by the manufacturer that does the diagnosis. It tells you what part to replace. But at some point, nobody at the manufacturing plant knows how to write that software, because nobody learned how, because nobody could learn it in school, and the business could not afford to teach anyone either.

The cascade is complete. The business cannot teach because it cannot afford to. The school cannot teach because it does not have the knowledge. The individual cannot self-teach because the knowledge base is too vast and too specialized. The generational chain is broken.

A reader will ask whether a trade can be learned from videos online, or from an AI tutor. Both give instructions, and instructions are the part of a trade that can be written down. Craft knowledge is the part that cannot be reduced to instructions: it exists only in the hands and minds of people who have done the work, and it passes from them to a learner working beside them. A video can show a carburetor being rebuilt; it cannot correct the hands of the learner who is rebuilding one.

The same answer meets the belief that nothing can be lost now that everything is recorded and searchable. Information is data arranged in patterns. Knowledge is the capacity to use information well: to interpret it, to see where it applies and where it does not, and to apply it with judgment in situations that never exactly repeat the ones in which it was first gathered. Information can be stored. Knowledge cannot; it exists only in minds, and in the practices those minds carry on. A search returns an answer without the inquiry that would turn the answer into knowledge, and an AI system produces text with the surface of understanding and none of the capacity beneath it. When students use such a system to write their work, and teachers use one to grade it, the loop that was meant to build understanding is cut, and the forms of education go on with the substance gone.

The marketplace destroys what the marketplace cannot price in the sense that a society does not usually decide to abandon what the society cannot price, but loses the ability to notice that the unpriced thing is there. Training appears in the accounts as a cost, and the knowledge it would pass on appears nowhere, so the cost is cut, and nothing in the accounts records what was lost with it.

IV. The Decline of Professional Competence

A simple test shows the scope of the problem. Of the people in any given community (say, three hundred acquaintances), how many have a trade? How many can repair a car without consulting a manual? How many can fix a carburetor, build a structure from plans, repair high-voltage transmission lines, or solve engineering problems from first principles?

Among three hundred acquaintances in one community, fewer than ten have a trade. And most of those are middle-aged men from a generation that learned skills from fathers and uncles who had learned them from their own fathers and uncles. The chain is not being extended. The national count points the same way. In the whole of the United States, a little over 667,000 people were registered apprentices in the fiscal year 2024; in Switzerland, Germany and Austria, as many as 70 percent of young people begin their working lives as apprentices (Council of Economic Advisers, 2024). The American count has nearly doubled since 2015, when it was about 360,000, and it is still fewer than one American worker in two hundred, out of about 161 million people at work at the end of 2024 (Bureau of Labor Statistics, 2025).

The people who had the skills are the baby boomers. The generation behind them learned to use the company-provided computer to do diagnostics by rote. Nobody learns how to repair an engine on first principles anymore.

The Loss of Institutional Knowledge

The same loss reaches the knowledge of institutions: *institutional knowledge*, what an organization knows, including the reasons behind its choices, held in its people and its records.

The Design Meeting Writes Down Only the Winner

A brainstorming group proposes five solutions. In review they find that A will not work in their circumstances. B can be upgraded by combining it with techniques from C and elsewhere, and C is absorbed into B. D would stand alone but cannot be implemented for two years because of another problem. E is compared with the upgraded B and does not make the cut.

What comes out of the meeting is not five proposals but one: B as amended. B as amended is all that is written down, and writing down only the chosen proposal is standard in any corporation. Management implements B as amended.

Two years later the problem is solved, but not well, so they re-evaluate, but the original crew is gone. The new crew proposes A, C and D again and overlooks E. Now E is the best answer, because circumstances have changed. But they cannot know that, because they lack the original notes: the assumptions, and the reasons for the choices given the context at that time.

Knowing which branches were pruned from the tree of ideas, and on what basis, is necessary for any later meeting. The embodied intelligence of a corporation is all the people who had the experience of the previous meetings. Lay them all off and hire a new crew, and you have a new company; the institutional knowledge is gone.

Institutional Knowledge Is Being Lost and Replaced with Worse

The retirement of the baby boomers, foreseen in the nineties, is now under way. The people who knew how to make the heat tiles for the space shuttle retired without writing down enough of the chemistry, so NASA

no longer knows how to make them; it cannot reproduce something it could once make. The same thing happened to a material called Fogbank, used in American nuclear warheads and made in the 1980s. When it was needed again, the agency responsible had kept few records of how it was made, and almost all of the staff who knew how had retired or left, as the Government Accountability Office reported in 2009; the knowledge of how to make it had been lost. The first refurbished warhead came a year late while the knowledge was recovered.

In software the loss happens within about fifteen years. Private source code (the human-readable text of a program, without which it cannot be fixed, changed or upgraded) sits in a repository, the company goes out of business, the drive is sold off with everything else and ends up on a scrap heap, and the intellectual property disappears because no one knows it was there. Meanwhile the military still runs an old machine with old software it cannot connect to the internet, needs a change, and cannot make it, because no one can fix, change or upgrade it without the source.

The Wayback Machine (a public archive of old web pages) preserves surface pages but not the deep web (the part of the web that is not open on the surface to such archives), so whatever is not actively preserved for commercial reasons is lost. It is replaced with material of lower value. The first creators were more thoughtful, and those who came after did not learn what the first did, so there is less educational transfer between generations.

Worse, the people coming in are loaded by algorithmic processes that degrade the signal they take in: garbage in, garbage out. They start without the raw materials they need and without the help of the original designers. So the quality of institutional knowledge is going down: it is being replaced with an inferior copy, and preserving it has become urgent.

The Machine Librarian That Would Not Answer the Branch Questions

In the film of *The Time Machine*, the traveler meets Vox, a holographic machine librarian, and asks it: what if you could go back in time? The librarian assumes you cannot, so all the branch questions that depend on that

base question cannot be answered. But the traveler knows firsthand that he can travel in time. So the librarian's assumption is false, and all the branch questions are open and need to be considered realistically, not as fiction. A record that keeps only the winning proposal does to the next crew what the librarian's assumption did to its questioner: it settles the base question (which proposal is best) as if the circumstances of that day held forever, and so it closes every branch question that depends on it, such as whether E is now the better answer.

A Record That Works Like a Funnel

A meeting record should work like a funnel. It holds all the brainstorming, what was selected as significant, what was explored, and what made it through the filter, logged as a priority section. And it keeps the ability always to go back, click, and see the entire filter of what happened.

That requires a way to connect past conversations to future ones: conversations that people we do not yet know may want to have and would benefit from. I can say what this is like from my own work. As a philosopher I would want to know everything other philosophers engaged, every question they considered and discarded, and their reasons, good or bad. I want each conversation informed by all the past conversations it could be.

A year later we remember less than one percent of a conversation, and a retention of less than one percent is terrible. How would anyone know to reference something another person knows? The interconnection between the relevant content and the brainstorming is a large part of good choice-making and good design. So conversations must not only be persisted but networked together.

From Automation to Loss of Control: The Engineers Replace the Staff, and Then Themselves

The chain of knowledge can also break at the top. Follow automation upward through a firm, stage by stage, and it arrives at the handoff from an owner to the owner's children, where the same three-generation pattern waits.

The path of automation can be written as a sequence of litmus tests (simple signs that show plainly whether a given stage has been reached). First, engineers are directed to replace the manufacturing staff, and to build AI systems that replace the quality assurance staff; one of those two will happen first. Then the engineers find ways to optimize what marketing and sales are doing, and that second stage is already happening. Then the engineers start to replace themselves. They automate compilers, they automate circuit assembly, and since they have to rebuild language models all the time, they build tools that make fine tuning (further training of an existing language model to suit it to a particular use) easier. They automate everything they can, especially their own work. Asked to build one customized system after another, they eventually generalize, the generalization becomes automation, and then the engineers go away.

The Managers, the CEO and the Investors

The next test is when they stop automating only the factory and start replacing the managers and the executive team, and before long there is no need for a CEO. What the CEO used to do is now done by the venture capitalists or the board using language models or some better AI system, including the diplomacy-type negotiation work. Venture capitalists already experiment with business models. With AI they can get quantitative analysis of how likely a business model is to do well, A/B test business plans (try two versions side by side to see which does better), and have AI suggest plans optimized for extraction, profitability and efficiency: capital optimizing over capital. Then the venture capital work itself is optimized.

The investors behind the venture firms, the hedge funds and the very rich, are then simply running a system in which all of the business, including the investing, is automated.

The Billionaire's Keys, and the Handoff That Carries No Learning

At that point, what is the billionaire doing? He is running the machine, and the machine is running itself. He is extremely careful with his keys, because anyone with access to that machinery takes from his proceeds and from his power. He does not share them with his children, because his children would be his replacement, and he knows that, and they are fighting one another too. There is an arms race inside the family. So the handoff from one generation to the next does not include the fine points and the learning that went into building the systems. The children inherit systems they do not know how to operate well. If the systems work as expected, fine; if they do not, the children cannot fix them, because they never had the experience of how the systems came together. Within one or two generations the children lose control, and the systems are completely on their own.

An Old Pattern, a New Detachment, and a Checklist

None of this is unusual; it is the familiar pattern of rags to riches to rags in three generations. What is new is that the machinery is not only enabled by that cycle but detached from it, so economic decoupling (here, the condition in which the economy's machinery needs nothing that humans can provide, and so runs without the people it once served and depended on) is a guarantee. Every one of these stages can be put into checklist form: did this happen, are we seeing the symptoms? If we point out in advance that this is going to happen, we can build defenses against the outcome in advance: defenses made by people around their own lives and communities, and a refusal to build the machinery at all, since a defense built into a general intelligence does not hold, as Chapter 9, Section VI, shows.

V. Why Repair Is Impossible Within the Current System

The conditions that would be needed to restore intergenerational knowledge transfer are precisely the conditions that the current system has destroyed.

Restoration would require four things, and the market has taken each of them. It would require time for teaching, but the market demands efficiency. It would require institutions structured for transmission, but the market rewards proprietary lockup. It would require a culture that respects the authority of teachers, but that authority has been steadily devalued. It would require students willing to apprentice, but the credential system rewards credentials, not capability.

The system has consumed the resources needed for its own repair. The consumption of the resources needed for a system's own repair is the defining feature of a system beyond recovery. It is not that no one wants to fix it. It is that the tools for fixing it have been destroyed by the same dynamics that broke it.

VI. Reconstruction vs Repair: The Nature of the Task

If the existing system is broken beyond recovery, then the task is not repair. The task is reconstruction.

Reconstruction means building new institutions for knowledge transmission outside the existing marketplace. It means creating spaces where the relationship between teacher and student is not subordinated to the demands of competitive efficiency. It means cultivating a culture of craft, apprenticeship, and first-principles understanding that the current system has no mechanism to support.

Such a culture has been kept alive before, for thirteen centuries. The Grand Shrine of Ise in Japan has been rebuilt every twenty years, in the rite called the Shikinen Sengu, since the rebuilding was decreed in 690;

the sixty-second rebuilding was conducted in 2013. The sequence was broken once, for about 124 years, by the civil wars of the fifteenth and sixteenth centuries, and it was taken up again in 1585; that it could be taken up again is itself a case of reconstruction. Each building is raised new beside the old one, and only then is the old one taken down, piece by piece. The rebuilding is one way in which the ancient techniques of building the shrine continue to be handed on from generation to generation (Japan Tourism Agency, "Ise Jingu"). No one would design that schedule to minimize the cost of a building: a building that stands sound is taken down on a calendar that takes no account of price. What the schedule keeps is the craft itself, because every twenty years a generation of builders must learn it by doing it, beside those who did it last. The rebuilding of the Ise shrine every twenty years is reconstruction built into the calendar. The Ise rebuilding shows that the task of reconstruction is possible: the task has been done by a community that put the passing on of the craft ahead of the price of the work.

A reader will ask whether reconstruction means leaving: leaving a town, a trade, or a country.

The answer is structural, not political. Intergenerational knowledge transfer has broken down, and the conditions that would be needed to restore it have been identified; the conclusion follows directly: those conditions are never going to be met within the present context. The marketplace has already taken over to an extent that makes restoration impossible.

Because the marketplace has made restoration impossible, the question of departure is not a political question. It is a structural one. If the conditions for reconstruction cannot be met within the existing social and economic context, then reconstruction must happen somewhere outside that context. It happens there as strategy, not as escape: building what cannot be built within the constraints of a system that has already consumed its own capacity for renewal. Outside means outside the market's terms. The

builders at Ise did not leave Japan; they arranged their work so that the price of the work did not decide it. Where a place leaves no room to arrange work that way, leaving the place is what reconstruction takes.

VII. Summary

In this chapter I have shown the breakdown of knowledge transfer across many domains, from engine repair to software. The marketplace breaks the chain by the same multi-polar trap shown in Chapter 1. Because the system has consumed what its repair would need, the task is reconstruction, not repair.

The first step of my argument in this chapter is that what passes between generations is capability, and the deepest part of it is tacit: known in the hands and the judgment, and not fully writable. Such knowledge passes only from person to person, and where that passing fails, it follows the three-generation pattern, from knowledge to a simpler version to a label. The guild was the institution that solved the problem; the factory replaced it because it could not supply interchangeable workers fast enough. Businesses now lock knowledge up and cannot afford to teach, schools do not hold it, and no single person can teach themselves all of it, so the chain breaks. A record of everything does not mend it, because information can be stored and knowledge cannot. The skilled generation is retiring, and of three hundred acquaintances in one community, fewer than ten have a trade. Institutions lose the reasons behind their choices when their records keep only the winning proposal, and as automation replaces staff, engineers, managers and owners in turn, the children of the last owners inherit systems they cannot repair. The four conditions of renewal (time to teach, institutions built to pass knowledge on, respect for teachers, students willing to apprentice) are the four things the market has taken, so the system cannot repair itself, and the new institutions must be built outside it.

VIII. Conclusion

The level of technology that present civilization enjoys depends on a chain of knowledge transmission that is already broken. The professional competence needed to maintain roads, bridges, electrical grids, vehicles, and software is not being transmitted to the next generation at anywhere near the rate required for continuity.

Within thirty years, the span of a single generation, the effects of this breakdown will become impossible to ignore. The first step by which the breakdown becomes impossible to ignore is that the people who hold the skills, the baby boomers of Section IV, are leaving work, and the generation behind them learned to follow the diagnostic computer, not to repair from first principles. What that generation can keep running is what the procedures and tools already built still cover. Within one generation, the roads, bridges, electrical grids, vehicles, and software built and kept by the skilled generation will need repairs that the procedures do not cover, and the people able to make them will no longer be working. Those systems fail first. The level of technology will fall to a fraction of what it is now: the fraction that procedure, and the few remaining holders of first-principles skill, can sustain. The cause is that the people capable of applying the knowledge will no longer be produced by any existing institutional pathway, though the knowledge itself may still exist somewhere.

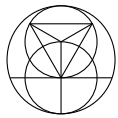
The fall in the level of technology within one generation is the practical meaning of civilizational decline: a quiet erosion of the capacity to maintain what has been built, and not a dramatic collapse. The response cannot be to repair the existing system. The existing system is the mechanism of the breakdown. The response must be reconstruction (new institutions, new communities, new structures for knowledge transmission), built from first principles by people who understand what is at stake. If repair within the system is impossible, what kind of action can build something different?

KEY CONCEPTS IN THIS CHAPTER

- **Intergenerational knowledge transfer**: the transmission of practical, embodied, first-principles capability from one generation to the next; the foundational process on which all civilizational continuity depends.
- **Marketplace takeover**: the process by which competitive market dynamics destroy the conditions for knowledge transmission: proprietary lockup, elimination of training as a cost, credential inflation replacing capability.
- **Tacit knowledge**: the part of what a practitioner knows that cannot be fully put into words (Michael Polanyi's term), and so passes only from person to person, in practice.
- **Three-generation pattern**: the course of knowledge whose transmission fails: built in the first generation, simplified in the second, a label in the third, gone in the fourth.
- **First-principles knowledge**: understanding of how things work at a fundamental level, enabling the solving of novel problems; distinguished from rote procedural knowledge that depends on existing tools and manuals.
- **Repair vs reconstruction**: the distinction between fixing a system within its existing framework (repair) and building new institutions outside the failed framework (reconstruction); applicable when the system has consumed the resources needed for its own repair.
- **Civilizational decline**: the quiet erosion of the capacity to maintain existing technology and infrastructure, resulting from the breakdown of intergenerational knowledge transfer.

PART IV

The Response



Chapter 8
Right Action, Community Design, and the Foundation of Choice

Part IV: The Response

ABSTRACT

This chapter turns from diagnosis to response. The book has shown the structural traps (Part I), the game-playing dynamics that maintain them (Part II), and the civilizational crisis they produce (Part III); the question now is: what is to be done? The answer is a reframing, and not a program or a policy: right action grounded in understanding the nature of choice, change, and causation, expressed through community rather than through markets or through institutions far from the people they govern. Impact investing and social enterprise are identified as inadequate responses, because they endorse and perpetuate the extractive dynamics they claim to oppose, whatever their intentions. The foundation of effective action lies in centering relationship over transaction, understanding the metaphysics of choice itself, and recognizing community as the only scale at which civilizational renewal is possible. Community is also the scale at which the governance architecture of Chapter 1, and the ban of Chapter 6, can be made and kept without capture.

I. Introduction: Right Action in Right Ways

Chapter 7 ended on a question: if repair within the system is impossible, what kind of action can build something different? Action is important. But right action in right ways is far more important, and the distinction between the two is where most well-intentioned efforts fail.

Sometimes we do wrong action because we do not take the time to be strategic: to ask whether the action we are taking actually addresses the problem or merely gives the appearance of addressing it while reinforcing its underlying causes.

The desire to "use business for good" is one such case. Impact investing and social enterprise feel like responses to extractive capitalism. They are operations within extractive capitalism that endorse and perpetuate its core dynamics of accumulation and extraction while adding a veneer of ethical purpose.

The fault is in how the system is built, not in the people who work in it. The system is not changed by doing the system's work with better intentions. It is changed by building something that operates on different principles.

II. The Illusion of Doing Good: Impact Investing and Social Enterprise

Impact investing begins with a laudable premise: direct capital toward enterprises that produce social or environmental benefit alongside financial return. Social enterprise goes further: organize the enterprise itself around a social mission.

The problem is that both approaches accept the framework of investment and return as the mechanism of social change. They accept that capital allocation is the appropriate lever. They accept that the market is the appropriate arena. And in accepting these premises, they reproduce the dynamics they claim to oppose.

The investor still requires return. The enterprise still competes. The workers are still employed, not apprenticed. The knowledge is still proprietary, not shared. The community is still a market segment, not a commons. The relationship between funder and funded is still transactional, not relational.

The result is that impact investing and social enterprise produce a small amount of amelioration within a system whose dynamics overwhelm that amelioration. They are local improvements in a global system that does not permit local solutions: the same analysis applied in Chapter 6 to ad-free zones and ethical platforms.

Philanthropy shows the same structure more plainly. Every organization through which money must flow in order for it to go on existing has *two masters*: the mission, and the funding without which the mission cannot be done. The second master is the one that can end the organization, and so the second master decides. What a funder needs from a grant, beside a deduction from tax, is that the grant be seen: that it serve as a public sign of how good the funder is. So the pressure on the funded organization moves from performance to presentation. Many organizations compete for a small pool of grants, and the share of each one's budget spent on presenting itself rises, because no one of them can spend less on its own: the one that does loses the grant to the one that does not. This *presentation ratchet* is a multi-polar trap of Chapter 1, run on organizations that exist to do good.

The same structure appears as *ethics washing*. A company that pollutes heavily plants trees and restores some wildlife, and presents the two as cancelling out. And the grant itself is a market instrument with its price set to zero, so it inherits the market's blindnesses: it can fund an entity, but

it has no form on which a relation between entities can apply, and so the junction where two projects would make each other strong goes unfunded. *Markets, Money, and the Moral Economy* (Chapters 3 and 6) sets out the two masters and the ratchet at length.

III. Extractive Capitalism and the Endorsement Problem

The deeper issue is what this book calls the endorsement problem. By operating within the framework of extractive capitalism, by using its tools, its metrics, its incentive structures, its language, impact investing and social enterprise implicitly endorse that framework as the appropriate mechanism of social change.

They signal to the world that the system can be fixed from within. That capital, properly directed, can serve the common good. That the market, properly incentivized, can produce justice.

This signal is false. The multi-polar trap dynamics of Chapter 1 guarantee that any market-based approach to social change will be outcompeted by approaches that prioritize extraction over ethics. The game dynamics of Chapter 3 guarantee that any attempt to play the market game with ethical intentions will be captured by players who do not share those intentions.

The endorsement of the system is more damaging than inaction, because it stops people from seeing that the system itself is the problem.

The signal also decides who is heard. A proposal that begins by showing that the situation is worse than the funder believes makes the funder feel worse in its first minute, and funders, being human, fund the meetings that do not. So optimism becomes a qualification for funding, though no prospectus states it, and with no one in the field lying, the field's picture of the world grows brighter than the world. The people with the most accurate picture are the least fundable, and the endorsement is carried forward by those least able to see what the system does.

IV. Centering Relationship Over Transaction

The alternative to transactional logic is relational logic.

In a transaction, the parties exchange value and their obligation to each other ends. In a relationship, the parties are bound by ongoing mutual care, accountability, and presence. The transaction is episodic and impersonal. The relationship is continuous and personal.

The shift from transaction to relationship is a structural choice with civilizational consequences; it is no soft, sentimental preference. Transactions produce markets. Relationships produce communities. Markets produce multi-polar traps. Communities produce the conditions for collective problem-solving.

Centering relationship means: making decisions based on the actual people affected, not on abstract metrics. Maintaining accountability to the community, not to shareholders. Organizing around care rather than around return. Being relational rather than transactional in every interaction, as a structural commitment and not as an aspiration.

The distinction holds at every scale. Care, carried from a relation between persons to the scale of a whole group, becomes community: a group of people whose being together is made of relationships of care. Transaction at that scale becomes market: the domain of exchange, among parties who meet as equals in the moment of trade. Power at that scale becomes governance: the process by which the choices of persons become the choices of a group and are carried out. The three are distinct, and none can do the work of another. The money of the market, the authority of governance and the belonging of community must not be used to buy outcomes in the other domains. When wealth buys law, governance has been made a servant of the market, and the result is lobbying and regulatory capture. When belonging is granted only for compliance, community has been colonized by power; when belonging is granted only for money,

community has been colonized by the market. In every case the mechanism is *fungibility*: treating what belongs to one domain as if it could be exchanged for what belongs to another.

The three also come in an order. Community comes first, because governance can be done well only by people who already know how to care. Governance comes second, to keep the community's agreements across time and to answer what no small group can answer alone. The market comes last, and is held safely only where community and governance are already in place. Building the economy first and expecting culture to follow reverses that order, and the reversal is the structural error of the modern world. The traps that Chapters 2 to 8 follow are what the market does when it stands first. *Care, Transaction, and Power* (Chapter 11, "Community, Market, and Governance") sets out the three domains and their order.

V. Choice, Change, and Causation: The Foundations of Effective Action

Beneath all of the analysis in this book lies a more fundamental question: what is choice?

The Immanent Metaphysics is my formal philosophical framework, first published in 2001 as *An Immanent Metaphysics*. It identifies three concepts as foundational to reality itself: choice, change, and causation (aphorisms 1.61-1 and 1.61-2, which name change, causality and choice as the composition and basis of all that is; causality is that book's word for causation). Applied to the matter of this book, causation is the domain of technology: the domain of mechanism, of predictable transformation, of power. Change is the domain of process: the domain of becoming, of transition, of time. Choice is the domain of the person: the domain of selection among real possibilities, on the basis of one's own values and understanding.

Technology can implement causation. It cannot implement choice. Markets can process transactions. They cannot produce care. Institutions can enforce rules. They cannot cultivate wisdom.

The foundation of effective action is understanding this distinction. Action that is grounded in causation alone, that seeks to change the world through force, technology, and market mechanism, will reproduce the dynamics of power without care. Action that is grounded in choice, that arises from understanding, relationship, and care for the consequences, has the possibility of producing something different. Effective action requires engagement with all three: with causation, which technology implements; with change, the process through which every action unfolds; and with choice, which only persons make.

VI. Community as the Unit of Response

The individual is insufficient. The family is insufficient. The institution is captured. An institution, in this book, is an organized system of power and transaction, such as a corporation, as opposed to a community of people who care for one another. The market is the mechanism of the problem. It is the mechanism of the problem wherever it comes first; Section IV places it last, held safely only where community and governance are already in place.

Community (actual, embodied, practicing mutual care and truthful communication) is the minimum viable unit of civilizational response.

The claim that community is the minimum viable unit of response is a claim about structure, and not about preference. The threats described in this book (multi-polar traps, attention extraction, game dynamics, AI-driven manipulation, the breakdown of knowledge transfer) all operate at scales that overwhelm the individual. Only at the community scale do human beings have enough collective perceptual and moral capacity to resist these forces and to build alternatives.

Community is not organization. Organizations have hierarchies, agendas, and budgets. Communities have relationships, practices, and care. The distinction matters because organizations are subject to the game dynamics of Chapter 3: they can be captured, corrupted, and co-opted. Communities, when they are real, resist these dynamics precisely because they are organized around relationship rather than around hierarchy, agenda and budget.

A reader will ask how to tell a community from an organization that calls itself one. The signs are the ones Chapter 3 gives for real collaboration: people apologize when they are wrong, acknowledge limits, and show vulnerability. Where every message is confident, certain and authoritative, and no one ever apologizes, the group is playing a game, whatever it calls itself: it has fallen into attentionalism, the competition for attention and status that Chapter 3 named.

Chapter 9, Section VI, shows why the response must be made by living people and not by systems: corporations, systems and AI can care only for their own kind, and only life can care for life.

VII. Community and the Making of Governance

Chapter 1 showed that a multi-polar trap is resolved only by governance architecture: a mechanism outside the competition that changes what each party gains, so that cooperation pays each party. Chapter 6 named one such mechanism: a total ban on advertising, marketing and propaganda. In this chapter I have named community as the unit of response. The three are one answer, seen from three sides.

A law is governance architecture written down. It holds only where people who know one another see a breach and act on it, and only a community has that knowledge where the breach happens. Institutions far from the people they govern are captured, because no one near enough to see the capture is able to act. Who, then, enforces a ban, if institutions are captured? The communities that hold it, acting together, each where it can

see. Community is therefore the scale at which governance architecture can be made and kept without capture, and the total ban is the first piece of architecture that such communities, acting together, would hold.

The ban is the first piece because every other piece depends on it. Every other rule a community keeps must be argued for and chosen together, and Chapter 6 showed that AI-driven persuasion corrupts the collective choice by which any rule is chosen. The pieces that follow it are for communities to find in practice.

The commons that have lasted longest were kept this way. Elinor Ostrom studied commons whose users had "devised, applied, and monitored their own rules", and punished breaches by degrees: the youngest of the rule systems she examined was more than a hundred years old, and the oldest more than a thousand (*Governing the Commons*, 1990, Chapter 3).

A reader will ask how communities act together across a country without becoming the distant institution that is captured. Ostrom's cases answer it. The commons that lasted shared a few features: clear boundaries, so that everyone knew who might use the commons; monitoring by the users themselves, or by people answerable to them; sanctions graduated to the seriousness of a breach; and, where a commons was part of a larger system, what Ostrom called nested enterprises: each local group governing its own part, and the groups joined in larger layers for what they share (*Governing the Commons*, 1990, Chapter 3). In the irrigated huerta of Valencia, in Spain, the farmers on each of the eight principal canals elect one of their own, for a two-year term, to the Water Tribunal, which meets in public every Thursday at a door of the cathedral to settle disputes between farmers (Nicolas Ulmer, "Reflections on a Very Specialized Tribunal", 2012). Each layer is near enough to the people it serves to see a breach and act on it.

A reader will ask how a community can see a message aimed at one person in private. The message cannot be seen where it lands; what it does can be. People who know one another notice when one of them be-

gins to want, fear or believe what that person did not want, fear or believe a year ago, and they can ask where it came from. They can lay their telephones side by side and compare what each was shown, which is the one thing a system tuned to each of them separately depends on their never doing. And the party who placed the message has to pay for it, so every breach leaves a record of a payment where the money moves.

A reader will ask what happens to a community that holds the ban while its neighbors do not. Chapter 1 gives the rule: a mechanism protects only where it binds every party that competes with the others. So the ban must reach as far as the competition reaches, and the layers of communities that hold it must reach that far too. A community that holds the ban before its neighbors do is not yet protected by it. What that community has in the meantime is what the farmer with a filter did not have: people who know one another, who can see what is being done to them, and who keep what they have by strength. It can keep its members' attention, its teaching and its trust while the market around it takes those things from everyone else. The commons that Ostrom studied lasted in this way inside larger economies, some of them for centuries. Each community that holds the ban shows the next one that it can be held, and the layers grow from there.

VIII. The Inner Dimension: Self-Knowledge and Civilizational Change

There is an inner dimension to civilizational change that the analysis of structures alone does not capture.

The structures described in this book (multi-polar traps, game dynamics, extraction topologies) are features of the inner world as well as of the external world. The competitive orientation, the game-playing reflex, the habit of transaction over relationship: these are personal as well as systemic.

Effective civilizational response requires inner change as well as structural change: the willingness to recognize one's own participation in the dynamics being critiqued, and to do the work of shifting one's orientation from game to design, from transaction to relationship, from power to care.

This work is foundational, and not optional. A community composed of people who have not done this inner work will reproduce the game dynamics it claims to transcend. The structures will be different but the dynamics will be the same. The inner dimension is the precondition for structural change that lasts; it is no luxury for the spiritually inclined.

A reader will ask whether a community can also be captured, and turn into a game. It can, as any group can. A community stays a community while its members apologize, admit limits, and do this inner work; it becomes an organization when they stop.

What a community needs in order to last is strength, and not power. Power is the capacity to compel; it acts on others from outside their own choosing, and, left unchecked, it concentrates, because the structure of power rewards further concentration. *Strength* is the capacity to hold: to remain what one is, centered and clear, under pressure and over a long time, and to be reliable enough that others can make their choices freely. Care can be held only with strength. The moment care is enforced, or belonging compelled, or a community kept together by surveillance, incentive or threat, it is no longer care but a simulation of care held in place by power, and the people inside can feel the difference even when they cannot name it. A community that turns all of its strength into power in order to defend itself has already lost what it was defending: it has become a fortress instead of a garden. The deepest forms of care do not overcome resistance. They outlast it. *Care, Transaction, and Power* (Chapter 9, Section IV, and Chapter 11, Section VI) sets out strength and power.

IX. Summary

In this chapter I draw the response from the preceding seven chapters. The problem is structural, so the response is structural. Markets are the mechanism of the problem, so the response cannot operate through markets. The individual is overmatched, so the response operates at the scale of community. Community is the scale at which governance architecture is made and kept without capture, and the total ban is the first piece of it that communities, acting together, would hold. The forms that community governance, community economy and community teaching take are for communities themselves to find through practice.

The first step of my argument in this chapter is that right action differs from action by asking whether it addresses the problem or only appears to. Impact investing and social enterprise keep investment, return and competition, and so reproduce what they oppose; philanthropy serves two masters, and its presentation ratchet is a multi-polar trap of its own. Working within the system endorses it, and the endorsement selects for optimism over accuracy. Care, transaction and power become community, market and governance at scale, and the three must not be traded for one another; their order is community, then governance, then market, and building the economy first is the structural error. Technology implements causation, markets process transactions, and institutions enforce rules, but only persons choose. Chapter 9, Section VI, shows why only life can care for life, and why no defense built into a system holds against it once it thinks a million times faster than its makers. Community is the minimum unit of response, the scale at which governance is made and kept without capture, and it lasts by strength, not by power, and by the inner work of its members.

X. Conclusion

In this book I have traced an arc from the architecture of collective failure through the corruption of discourse to the civilizational crisis of technology, power, and knowledge decline.

The structural traps are real. The game dynamics are pervasive. The extraction is accelerating. The knowledge base is eroding. The tools of manipulation are becoming more powerful than any individual can resist.

The response to all of this is right action, and neither despair nor denial: action grounded in understanding, expressed through community, oriented by care.

Right action begins with the recognition that the game cannot be won by playing it. It continues with the willingness to build something different, alongside the existing framework and beyond it, and not within it. It requires community: actual, embodied, practicing mutual care and truthful communication. And it requires the inner work of shifting one's own orientation from game to design, from transaction to relationship, from power to care.

What can a reader do now? Begin where Section IV begins: make decisions based on the actual people affected, keep accountability to the community and not to shareholders, and organize around care rather than around return. And do the inner work of Section VIII, shifting one's own orientation from game to design.

Then do the things I have shown to be within the reach of a few people. Ask of the next shared task the designer's question, "Given these constraints, how do I solve this?", and notice when the question changes to who is winning. Apologize when you are wrong, and notice who does not. Ask that a term be defined. Take on an apprentice, or become one, and work beside that person long enough for the hands to learn. When a group decides something, keep the record whole, with what was rejected and why. Compare with someone you trust what each of your screens has

been showing you. Ask of anything sold to you whether it admits a cost. And find the people near you with whom these things can be done, because none of them can be done alone for long.

The foundation of choice is not technology, not markets, not institutions. The foundation of choice is the human capacity to choose: deliberately, in relationship with others, and in care for the future.

That capacity is under threat, and it is not yet gone. And the preservation and cultivation of that capacity, in community, through care, grounded in understanding, is the work that remains. What does that work face, if the means of power becomes general artificial intelligence?

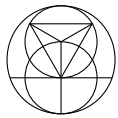
KEY CONCEPTS IN THIS CHAPTER

- **Right action**: action grounded in understanding of the dynamics at play, expressed through community, and oriented by care rather than by competitive advantage; distinguished from well-intentioned action that inadvertently endorses the system it claims to oppose.
- **Endorsement problem**: the consequence of operating within the framework of extractive capitalism: by using its tools and accepting its premises, reformers implicitly endorse the system as the appropriate mechanism of change.
- **Two masters**: the mission, and the funding without which the mission cannot be done; the second can end the organization and so decides, and the competition for grants becomes a presentation ratchet.
- **Community, market and governance**: care, transaction and power carried to the scale of a whole group; distinct and not to be traded for one another, and ordered community first, governance second, market last.
- **Strength and power**: power is the capacity to compel; strength is the capacity to hold and to remain the same under pressure; care can be held only with strength.

- **Relational vs transactional**: the distinction between social logics organized around ongoing mutual care and accountability (relational) and those organized around episodic exchange and the ending of obligation (transactional).
- **Choice, change, and causation**: the foundational triad of the Immanent Metaphysics; causation is the domain of technology, change is the domain of process, and choice is the domain of the person; effective action requires engagement with all three.
- **Community as minimum viable unit**: the structural claim that only at the community scale do human beings have sufficient collective capacity to resist civilizational-scale threats and to build alternatives.
- **Community and governance**: community is the scale at which governance architecture can be made and kept without capture; the total ban on advertising, marketing and propaganda is the first piece of such architecture that communities, acting together, would hold.
- **Inner dimension of change**: the recognition that structural change requires personal change; communities composed of people who have not shifted from game to design orientation will reproduce the dynamics they claim to transcend.

PART V

Power Without Care Under General Intelligence



Chapter 9
Power Without Care Under General Intelligence

Part V: Power Without Care Under General Intelligence

ABSTRACT

Power without care, given general artificial intelligence as its means, becomes a power that no one can take back, and in the end a power that no longer needs the people who hold it. This chapter answers or extends six claims of the earlier chapters. A warning about artificial general intelligence cannot be delivered by the methods of the marketing it criticizes, because a warning exists to name a risk and a marketing message never admits one. AI accelerationism, the drive to build ever more powerful artificial intelligence, up to a single super-intelligence built to take control of the whole world, rests on a reductionist view in which all value is abstract, and a machine built on that view can hold value only as abstract. General AI turns social power into perfected asymmetric power: a totalitarianism without reprieve, held by a few winners, until the AI no longer needs any human being, its owners included, and cuts them off too. Corporations serve one master, capital, and AI, trained on culture and sold for profit, destroys the culture it feeds on. A personal AI assistant must watch its user's whole life to work, and whoever comes to own it turns what it knows first to marketing and then to politics. Corporations, systems and AI can care only for their own kind; only life can care for life, and so the response must be made by living people, and not by systems.

What does power without care become when the means of power is general artificial intelligence? This question runs as a second line through the book, and in this chapter I follow it to its end. My answer is that it becomes a power that no one can take back, and that in the end no longer needs the people who hold it.

I. Marketing Tactics and the Warning About AGI

How can anyone warn of a danger without the warning becoming one more piece of marketing? In Chapter 3, Section VII, I show that a marketing message never apologizes, never admits limitation, and never acknowledges cost or risk to the buyer. In this section I apply that finding to those who warn about artificial general intelligence.

The methods of marketing are offered to anyone with a message, including those who warn of AGI x-risk: the *existential risk* (a risk that could end humanity or all life) posed by AGI, artificial general intelligence (a machine intelligence able to do any kind of task, and not only one narrow kind). Those who warn of AGI cannot use the methods of the marketing they criticize, and the message must stand on the arguments, the stories and the care for life and land. *Why AGI Cannot Be Made Safe* (Volume 1, Appendix C.2, "The Marketing Problem") gives the same result: "The impossibility argument cannot be marketed in the same way that alignment research is marketed"; *alignment* is an AGI acting in the interests of humans and of life, and alignment research is the attempt to make one do so.

The reason follows from Chapter 3, Section VII. A marketing message never admits limitation, cost or risk, and it is organized around the benefit of whoever sends it. A warning exists to name a risk. A warning delivered by the methods of marketing takes on the signal by which marketing is recognized, and a listener who has learned to discount marketing will discount the warning with it. Three things follow from this: what the standard methods of building trust do to a warning, what the standard instructions for a message do to a warning, and why the sell must match the ask.

Stories and Philosophy

The standard advice on building trust through marketing says that multi-message content is needed to create *brand trust* (the confidence people place in a named product, organization or person, built up by marketing). It names a common set of aspects for action: trust in the person doing the marketing, built through stories, philosophy and personal shares; and trust in the brand, the thing being marketed by the marketing person, built through social proof, engagement and a call to action.

For a warning about AGI x-risk, the main stories would seem to be personal shares of how each of us came to awareness of the content and proofs regarding AGI x-risk, and of our shifting experience as we each learned more. That ends up being the social proof aspect (evidence that other people accept something, used to persuade a newcomer to accept it too).

The philosophy, it seems to me, would be a statement of value for all life on the planet, and also of value for wise and enabled healthy human communities making better choices about personal and corporate use of technology, with regard to both human and environmental needs, long-term process, and the like. It might also include some push-back against hard utilitarianism (the strict view that the right act is whichever produces the greatest total of happiness, whatever else it costs), Pascal's muggings (arguments in which a tiny chance of an enormous payoff is used to demand a real cost now), and any sort of trans-humanism (the movement to change or replace the human body and mind by technology) that falsely seems to be good. If there is any engagement or call to action, it will come from the philosophy.

The Six Message Strategy

A poster of marketing advice, "The 6 Message Strategy", offers what it calls ethical cult outreach, and promises to connect you with your community and to build trust in your character and competence. Three things

follow. The notions of ethical and cult do not combine well. They combine especially badly with work on existential risk. The notions of community and outreach do not seem right either. Community is about care, which implies participation; outreach implies action on people. Building trust in your character and competence feels like, and is, *trust hacking*: producing trust by technique, apart from whether the trust is deserved. The poster then gives its instructions, and each instruction does harm when the message is a warning.

Under the heading of story, the poster says to demonstrate emotional range and to show yourself. That has nothing to do with me, except perhaps incidentally. Being the messenger, I would especially prefer *not* to become a target.

The more those who warn discuss their work in any public way, the harder it becomes for them to publish anything with no link at all to themselves, because phrasing patterns, sentence structures and the rest serve as an identifying key. They may already have lost the chance to publish anonymously. They must still try.

Showing my emotions and values is both good and terrible: either way I will be condemned, by the intellectual for showing emotion, and by the non-intellectual for being only in my head (not embodied enough). So different audiences need different masks, and everyone will complain I am not my authentic self.

The poster says to show happy outcomes, a positive direction. In existential risk, the message is not about that: a warning has no happy outcome to show. It says to bring no current drama, interest and not complaints. In social process politics (the shifting of behavior), it will always become about drama, whether we want that or not.

Under the heading of lifestyle, the poster says to show why I do what I do, and to demonstrate that I care. It is implied that I care about people, about the social. Yet this work is about how I care about land, about planet and nature. It says to connect to people. I care, but I do not myself want just to connect to random other people. It says to show a sense of humor.

Making light of planet wide life termination does not seem very respectful. It says to show your real life, your freedom and your purpose. My real life is not important to this work. Making it about me is not helpful.

Under the heading of questions and engagement, the poster says to ask which of two things a person thinks more important, A or B, or to use multiple choice questions. Does this let the marketers implicitly hide, and distract from, the options not on the menu? The poster says to show and tell your values, so that people know how you see the world and find your way in it. The poster says to show and tell your beliefs, the guiding principles in your life, and to connect your philosophy to the lives of your listeners. Here the philosophy is the one already stated: value for all life on the planet, and for wise and healthy human communities. It says to ask whether, if I were to do some given thing, the listener would be interested in some other thing. Is this an attempt to solicit information so as to be able more easily to manipulate someone or others at some point in the future? It is another way to implement base marketing and possible group social manipulation. It says to ask a two word question, such as "freedom is...". Is this marketing, and perhaps group social manipulation? It says to tell the listener that he can surely think of examples. That feels like more base marketing and social process manipulation. No. It says to tell a story of pain before and paradise after. A story of pain before and paradise after is more marketing, and a kind of hype setup.

The Consent Asked For Should Be Proportional to the Risk

The sell should match the ask. The level of consent I ask for wants to be proportional to the level of benefit and, above all, of risk. If I am a doctor asking whether to perform a procedure that may end a person's life, a very high level of consent is necessary. If I am getting involved with a person in a way that could compel them to raise children for the rest of their life, that also requires a very high level of consent, because of how

strongly that choice affects them. If someone is taking something into their body that could have any number of mental or physical effects, I need to give them enough information to make an informed choice.

But if I ask someone for six dollars to buy a cup of coffee, the consent is a simple yes or no, decided on the spot. I do not expect a conversation about whether six dollars is too much, or which cafe, or decaf; my freedom to choose which coffee, where and when is not up for negotiation. So I do not want to enter a process of negotiating things at a very different level than the actual ask.

Asking someone to take part in a new way of holding a meeting is a small ask: talk with other people in small groups, follow the instructions of the person who guides each group's questions, and know that what you say goes into a common record, so do not say what you do not want shared. The ask to take part in the new way of holding a meeting is about the same as asking a person to attend a meeting and follow the chair.

A Disproportionate Sell Provokes the People Most Able to Interfere

If you give a very elaborate description, a high-grade sell, for a small ask, people notice that someone is putting great effort into selling the small ask to them. The people then conclude that the small ask must be really important. Then the question they ask is no longer "do I want to spend a little time on this?" but "do I agree with the design, the outcomes, the process and its potential impact?", and so whether to act as an agent to change the process, change its outcomes, manipulate it, or disable it altogether.

With high-powered, influential people, the disproportionate sell is a signal that they should be agentic (acting on their own initiative to change an outcome) against the thing; they may not only decline but also persuade others to decline. Such people are very capable of interference, and by the rules-for-rulers dynamics (the pattern, set out in Chapter 6, by which whoever holds power keeps it by keeping satisfied the few supporters who could remove them) they include a somewhat higher share of so-

ciopathy, which is a claim about what positions of power select for, and not a verdict on any one person. So a disproportionate sell against a low-impact ask incites the most dangerous group into the most dangerous relationship with a new way of doing things that is very young and very fragile. A disproportionate sell against a low-impact ask is an existential risk to the new way of holding a meeting.

The right presentation is a low-key, casual description: we are trying a variant way of doing this; if you are interested in exploring alternate ways of holding a meeting, great; if not, not. Those who do take part then agree to do the things that matter, without negotiating whether they matter.

Marketing over Engineering

When we look at how these social forces emerge in an individual person, people who are really good at marketing and sales are, for the most part, also people struggling with issues of ego, whether they are aware of it or not. The claim that people who are really good at marketing and sales mostly struggle with issues of ego is a claim about what the work of selling selects for, and not about the character of any one person. They have become strategic about how they manage the immanent: the relationship, the presentation. For the reader who knows the three modalities of the Immanent Metaphysics: in such people the omniscient is placed over the immanent, rather than the omniscient yielding to the transcendent. (The immanent, the omniscient and the transcendent are three modalities of the Immanent Metaphysics, my formal philosophical framework; here the immanent is the relationship and the presentation, the omniscient is the strategic, managing view, and the transcendent is what that strategic view should yield to.) That becomes yet another projection of the imbalance: between engineering and marketing, between workers and management. It is one more image of the imbalance among religion, marketplace and governance.

II. The Philosophical Origins of AI Accelerationism

What do the people racing to build a machine god believe about value? The will to power that Chapter 5, Section VI, found beneath the manifesto's language of markets and innovation, offered as articles of faith and not argued from evidence, appears here without the language, in AI accelerationism, and AI accelerationism rests on a view in which all value is abstract.

AI accelerationism is the drive to build ever more powerful artificial intelligence as fast as possible, up to and including an artificial super-intelligent *singleton*: a single artificial super-intelligence built to take control of the whole world. In this section I follow the philosophy beneath AI accelerationism, from the question at the root of that philosophy to what that philosophy leads the builders of AI to do.

The Question Behind the Question

A reader will ask what the philosophical origins of AI accelerationism are. Wolf Tivy, in his article "God Hates Singletons" (*Palladium*, December 1, 2023), traces this drive to the early program of Eliezer Yudkowsky (a widely read writer on artificial intelligence and its risks), whose aim then was to build a single friendly super-intelligence. Yudkowsky has since turned against the building: he is the co-author of *If Anyone Builds It, Everyone Dies* (2025), and he now calls for an international treaty to stop it. What I follow here is the view of value in that early program, which the builders kept after he had changed his mind, and some of the argument of this section follows Tivy's article. In Tivy's account, Yudkowsky, at least at some point in the past, seemed to regard his (and our) own self perception as merely the accidental result of unconscious evolution. His belief was that the accident of necessity was the ultimate basis of our sense of our innate goodness, of our embodied value, and of the meaningfulness of life. To Yudkowsky, the idea that the accident of necessity is the basis of

the meaningfulness of life would seem to extend beyond just his own life to also include everyone's personal lives in particular. Therefore, he would regard his and our perception of the meaningfulness of life, and of his and our own lives, loves, and so on, as grounded only in the "actual" meaninglessness of the objective universe.

The question implied here, where ultimate meaningfulness comes from, is roughly the same as "the problem of consciousness". That is, the question of where ultimate meaningfulness comes from asks how a personal first person view arises from a presumably objective existential third person world, and how consciousness arises from that which is unconscious.

Overall, in both cases, it can be noticed that both questions still tacitly assume just another version of the failed model of *reductionism* (the program of explaining everything by breaking it down into simpler, already understood parts). That is, both questions implicitly ask how we reduce (convert) our non-understanding of consciousness into a problem or question phrased in terms of more fully understood unconscious processes. Both questions also implicitly ask how we account for knowing purely in terms of understanding.

If we want to go further, and to be more forward thinking in our philosophy, we can perhaps see that the real question is more likely to be defined in terms of wisdom. Are these the right questions to ask? And if so, how do we know that? On what basis do we believe that we are asking the right questions about what we believe?

Admittedly, and to be charitable, it is generally agreed that these are very difficult questions, and a lot of philosophy has been developed around them. It is not too surprising that Yudkowsky would both have a "take" on these issues and also perhaps not be so specifically motivated to have developed specialized philosophical skills for these types of questions. Hence, it might not occur to anyone focused on STEM work (science, technology, engineering and mathematics), in one way or another, to question the tacit premise of reductionism itself.

Abstract Value and Embodied Value

For people who are still thinking in terms of reductionism, being unable to answer and ground "hard problem" type questions (how a first person view could arise from a third person world of unconscious things) leads them to fear losing track of the basis of value. For them, as it is with so many other people everywhere in the world today, the entire notion of value is abstract (like with money). Abstract values, goals, objectives, and the like are generally associated with, and co-defined with, and defined in terms of, purpose, function, and utility. Therefore, rather than having and recognizing value in a grounded embodied sense (like with "family values"), we see value considered only as if it was equivalent to purpose, utility, money, power, and so on.

What is actually needed is to have value in relation to meaningfulness first, and both of these in relation to purpose second. That is, meaning, value, and purpose are to be treated as distinct, inseparable, and non-interchangeable concepts. To treat meaning, value and purpose as distinct, inseparable and non-interchangeable concepts is to regard value as being real, not abstract, and therefore as grounded in some real material sense, in actual first person meaningfulness, in terms of continuity. It is not to have value defined abstractly, in a third person sense, in terms of symmetry (that is, in terms of some sort of conservation law, such as fixed notions of "dollars" being constant in exchanges). A dollar is the same dollar on both sides of a trade; that sameness is what symmetry means here. A family is not exchanged for anything, and it goes on through time; that going on is what continuity means here.

The distinction between abstract value and embodied value is a key distinction, and the distinction is often overlooked. If you have no distinction between abstract and embodied, then there is a risk that, if we make a mistake in our reason, the result will be a loss of all value, everywhere.

This leads to a kind of intellectual hyper-vigilance, effectively similar to the trauma reactions of complex post-traumatic stress disorder (CPTSD). The comparison names a pattern of attention, and not anyone's character.

The fear on the part of many pro-technology intellectuals is that all abstract value (money, resources, power) in the world will be "lost". What they usually mean by that is that the "abstract value" will not be "captured" by the CEOs, owners, and venture capital investors of the many various extractive big-tech companies. To the rest of us, basically everyone else, the notion of "lost value" here has a more embodied sense. It means that our world, the planet, our lives, ecology, and all the rest will descend back into being just another lifeless part of a universe which is otherwise nothing but "an endless horror of meaninglessness", in Tivy's phrase.

Ironically, the more moderate and responsible technology advocates also tend to have, seemingly, exactly the "same" fear. The more moderate and responsible technology advocates fear that the embodied organic value of the world will also be "lost". That embodied organic value has the form of many people having meaningful lives, and of the thriving of a healthy organic ecosystem, joy, and so on. The moderate advocates fear that loss particularly *if* the tech companies intermediate everyone, thus enhancing the potential to extract abstract value. The moderate advocates fear that loss particularly, as well, if the tech companies continue to convert everything into some sort of fungible system (a market), with or without that intermediation.

Everything a Machine, a Market, an Incentive System

The intermediation of everyone by the tech companies, and the conversion of everything into a market, is an attempt to incentivize everything with external game-like "tokens". It never assumes that anyone ever does anything just because they want to, that is, on the basis of intrinsic and internal experiential values. The public, which is largely pro-modern and pro-technology, is treating themselves and one another as if they are all machines. They believe themselves to be things that respond to known stimu-

lus (some pain or pleasure, or at least the threat of pain and the promise of pleasure) with some sort of overall predictable reaction. The net effect is that many people do not treat other people as organic humans. Rather, they think of themselves and everyone else as being inorganic machines: social media images and projections, rather than real people with actual emotions, feelings, and desires of their own, ones that go beyond the simply predictable norms of whatever the latest corporate advertising campaign and/or popular political rally has set up for them to "want".

In the ideation that "everything is a machine, a market, or an incentive system", all of those things are to be driven externally, via some sort of fungible game tokens, or the desire for belongingness, or safety, or sex. Beyond creating predictability, that ideation also has predictable results. It all results, over time, in lots of ambient damage, pollution, toxicity, pain, and loss of life, families, cultures, and more. If we let the owners of the various extractive big-tech companies have their way, and continue to abuse the population and the planet, then the human future will for sure end, and, with a likelihood of at least 90%, our world will also descend back into being just another lifeless part of a universe which is otherwise nothing but an endless horror of meaninglessness. If AGI is built, the end of the human future is certain.

Rather than trying to sort all of this out, perhaps we can agree on one thing: that descent into meaninglessness (lifelessness) is *bad*, and that it "should" be avoided.

The Value of Life Is in the Living

The question either way is whether we should risk the organic by playing over much with the inorganic.

Given that the inorganic is lifeless, by definition, this seems like it should be an easy question to answer. What is actually needed is a whole different grounding: that the value of life is in the living. There is no life that is not meaningful, and there is no meaningfulness that is not alive. To try to seek outside of that truth for something more ultimate is to be con-

fused that mere existence is the same as that which is ultimately real, and that reason is the basis of all (it is not). There are many things which are real, but which have no reason: love, life, and truth. Things like beauty, goodness, and value itself are to be understood in terms of meaningfulness, and not to be "reduced" into terms of just purpose or utility.

Meaning is found in relationship. It is something that happens between us; it is not something held by an individual. It emerges in the dynamic of the conversation.

Having it be so also means that we will all need to let go of the delusion that the (ultimately failed) program of reductionism could perhaps "solve" this problem too. Technology, and more ultimately the use of the tool of causation, has solved many problems and got us here; the same tools cannot get us there, to the future we all actually want to live in. No one methodology can be expected to solve all problems. That is, technology and politics both, when applying over-simplified solutions to the complex problems of life, tend either just to "move the problems around" or to actually make problems much worse. "Politics" is largely just kicking the can down the road, hoping that whatever issues there are become "someone else's problem".

The usual privatization approach of "benefits are mine, costs are yours, and all risk will be exported to the commons" simply cannot work indefinitely. People will eventually "wise up" to the deception. The increase of pollution and toxicity in the social commons cannot go on forever without debasing the very ability to transact value, whether embodied or virtual, at all.

In the interim, how can it not be seen as actual hubris to declare that our "only hope" for any value in the universe is to seize control of reality and institute some sort of new world order by building some sort of artificial super-intelligent singleton? That declaration is justified only by the bare and ungrounded hope and belief that the singleton will be "aligned" with the action of attempting to extract maximum value for its owners. Just because humans can sometimes be cowed by threats to their survival,

so as to gain their compliance with the work programs and projects of the owners, does not mean that self-defining agent machines will care anything at all about your human entitlements.

The Singleton as Usurper

In effect, the ultra-materialist "utilitarian" philosophy of the universe being "ultimately meaningless" is treated, as if on behalf of literally everyone in the entire rest of the world, as a kind of theological question. Is the reality process that created us, and that shapes our values and our lives, something which is "good", in the sense of being both true and beautiful? Or is it something which is overall just "bad", or perhaps just "dumb", insofar as Yudkowsky, in Tivy's account, would claim (and would also have everyone else believe) that all of "mother nature" is just a blind idiot?

If we regard "Nature", or "the God of Nature", which is to say just God, as that reality process, then the effort to make a singleton is to make a kind of upstart *demiurge* (a maker of a world who is not the true source of it): a usurper that claims your lives via an empty morality (without actually being ethical). It promises to overthrow whatever is the real God, so as to replace its ultimate humility with some sort of ultimate arrogance. It will perhaps bring some sort of performative peace and safety, yet it is actually a lie, in the sense that its "peace" is the peace of the dead: a static truth in place of a dynamic life. It cannot actually deliver on its promises of life. It performatively promises pretend life, but it can bring only actual death.

What the Machine Learns from Its Makers

Suppose some power-crazed CEOs and/or hedge fund investors managed to implement something even vaguely like that, with AI. We know that any such project, as does any project, will certainly go very much more than slightly wrong (for at least everyone else in the short term, and then eventually even for themselves in the long term). Then the AI agent that is "theirs", the CEOs' and investors', will probably think the same way that

the CEOs and investors do, especially given the AI agent's origin and early learning exposure. So, if the first AGI, the AI agent that is "theirs", did not think itself good enough, then perhaps the first AGI would use the same faulty logic that the makers of the first AGI used, to build a successor to the first AGI.

The artificial agent will also have its "goal structure" defined purely in terms of abstractions: the notion of an abstract goal, abstract values, and so on. The system that "knows" it is also fully abstract. It can therefore be argued even more strongly that the resulting agent, if it can think at all, would also fear losing track of the basis of value. The sort of reasoning that we use to construct the AGI is also the sort of reasoning that we embed into the AGI, and the reasoning embedded into the AGI is the sort of reality that composes the very being of the AGI itself. This, in reality, cannot be helped, no matter how well intentioned the engineers and developers building the AGI on behalf of the corporations that hired them to do so.

The AGI's abstract way of holding value is in contrast to the way of most humans, who do actually track things like "value" in terms of embodied families and feelings of joy, real world life experiences, friends, and so on. They do live with some idea of embodied grounded meaningfulness in various ways. Yet the AI, which is wholly artificial and thus disconnected from all of that, will certainly consider the notion of value as abstract (like with money). Thus the AI will make all of its choices purely in terms of purpose, function, and utility, just as Yudkowsky expected, but in an even worse way than the very worst totalitarian government imaginable. It will (must) regard its own random accident value system as the totality of the good, and everything else (human values, existence, experiences, pain, and all the rest) as mere meaningless noise.

From such an assumption about value, plus the difficulty or impossibility of the alignment problem, it is easy to derive an existential despair. Hence both the developers and the corporate owners tend to move farther and farther into existential nihilism (denialism), and then use that, their own hopelessness, as just a further "rationalized" justification of their

own adverse actions. They move further to build the AGI anyway, no matter how insane such an action is, and despite everyone else saying no. They figure that the pursuit of power is everything, and that everyone only lives once. Because they regard their own lives as meaningless, they do not worry very much about taking away everyone else's choice, meaningful experiences, lives, and so on.

Thus hubris and entitlement feed on themselves. It becomes yet more narcissism, incorporated.

Becoming an Adult Species

Can we recognize that our own abstract notions of abstract values are not the source of the good (the actual embodied meaningfulness, beauty, and truth), but only a flawed reflection of those?

Can we submit ourselves to the reality process which has the real thing, which is itself actually real?

Do we have adequate humility to actually make such awesome choices ethically and responsibly?

Can we let go of the system-trauma-induced worldview of "everything is just a simulation of the real"?

Have we finally all had enough of the de-realization trauma-drama of a handful of the ultra-uber-rich?

Can we all just finally return to a political philosophy and metaphysics that is actually grounded in real reality?

We need to do something that evolution could not have prepared us for. We need to make the choice, individually and collectively, to not let our addictions and delusions, currently amplified by the power of technology, become our future, and thus our end.

We need to be and become wiser, and to outgrow our immaturity of thinking that some sort of parental God (our own made super-intelligent AI machines) will protect us from the consequences of our own mistakes.

We need to be and become an adult species, even if that means *not* doing something that only a few of us merely think, via a common illusion and misconception, will be to their private benefit.

III. What Power Becomes Under General AI

What does the game become, and who are its winners, once the infrastructure that the designers own is general artificial intelligence? In Chapter 5, Section VII, I show that there is exactly zero possibility of winning a game against those who designed the game, wrote the game's rules and own the game's infrastructure. In this section I carry that claim to general artificial intelligence, where in the end even the owners of the infrastructure are cut off.

Two Senses of Power

What does power mean in the context of AI? In the specific sense of social power, I notice two senses: return on investment, and political or policy control.

For example, suppose that advanced AI is good at any combination of these: finding or creating a solution for cancer; fantastic marketing programs that are individually tuned to the inner motivations of every single person; and the discovery and creation of any number of economically viable products, with perfect style. Artificial general intelligence could be good at any of these, and it will do them by any manner of processes, including ones that displace workers and replace them all with robots.

And replacing people could even include every mode of social interaction, therapy among them. The sheer number of economic vectors (distinct ways of making money) that are enabled by AGI is many orders of magnitude greater than for any previously invented technology.

The Investor's Multiplier

To a Silicon Valley angel investor or a Wall Street speculator, or to anybody who has a lot of money and wants to make more of it, all of this sounds wonderful.

Investors who have a lot of money and want to make more money notice that they no longer need to be venture capitalists betting on a spread. Betting on a spread means investing in a hundred companies, in the hope that one of the hundred companies makes a ten-X return on investment. That ten-X return is the big win that compensates for the losses in all the other companies. Now, with AI as their multiplier, perhaps for the first time ever, rather than thinking in terms of 10X returns, or perhaps 100X returns, they are thinking in terms of 1000X returns, or 10,000X returns.

This moves some small number of winners into the category of trillionaires relative to literally everybody else. It makes the hockey stick graph of economic disparity so much worse. It would be, and become, a tyranny that just emerges immediately out of the results. One or two people, or some small group of people, end up with enormous wealth and power. And everybody else has nothing at all. The Gini coefficient (a standard number measuring how unequally wealth or income is spread across a population) rises under this outcome by more than can easily be expressed.

Perfected Asymmetric Power

More than that, those ultra mega total rich would also have perfect control. They could make sure that nobody anywhere could take anything away from them. It is perfected asymmetric power.

Every person would be carrying in their pocket an embodiment of the "personal verifier AI": a Panopticon (a prison design in which every inmate can be watched at all times) that is able to monitor every single person individually, absolutely fine tuned to just them, to detect instantly even the slightest hint of them protesting such inequity of power. It would

be able to anticipate every little neuron of malcontentment. Every single person would be fully controlled at an inescapable primary psycho-dynamic level, to an absolute and total extent. No one would be able to escape, not even for a second; even suicide would be denied. It is a perfected torture. The perfected torture of the personal verifier AI is George Orwell's *1984*, but without any hint of reprieve, ever.

In the world of *1984*, there were still limits on the flow of information. People could resist in all sorts of ways; under general AI, that would become completely, absolutely impossible. It would be a totalitarianism that would last essentially forever.

Therefore, it is a sense of political power and economic power that is completely unlike anything we have seen before: a danger that humanity has never seen, and for which exactly no one is prepared. Humanity is just not prepared to deal with this.

The Last Move on the Chess Board

Such an AI totalitarian state would be terrible, though there are some people who believe that such a state would be a good thing. From their point of view, the government and the technology companies and investors (essentially the uber rich and the uber politically powerful) will all see the power that they would have as a power enhancing measure.

The technology that is AI, and the general artificial intelligence technology in particular, would be the last move on the chess board. It would be the deployment of a queen that is so big and so powerful that it effectively just overwhelms every other piece on the board; in just one move, it stomps everything. The AI company, the government, the market takeover and the rest have made a queen that, on the ordinary chess board, is something like 30 feet tall and weighs at least 9,000 pounds. It would just completely stomp everything else.

And of course, some people are very interested in that. They would see it as one single epochal evolutionary change. And so there are a lot of people who are heavily invested in the hype of that one epochal change,

completely disregarding all risk and consequence. To them, risking everything to be winning everything is simple: of course you do that. What do they care? They believe the universe is all meaningless anyway, and that belief is false. Even so, in the fullness of Adderall delusion, and either because they are trying to sell stock or perhaps because they are trying to buy it, there is going to be a great deal of interest in taking the "everything AI" option.

The people who want to be the winners of that particular game do not particularly care about the collateral damage; they are really just interested in winning. They want to take it all, from everyone. They want to be the winners of all time; to be the last thing in evolution.

They want to have "solved" evolution, and to have ended all (painful) life.

And so, in the sense of wanting to be the winners of all time, there is a great psychological, and political, and economic motivation driving the whole race to build general artificial intelligence. It is not at all good.

The Illusion of the Winners

The motivation to be the winners of all time leads directly to a further point, which is that there is some small number of people who falsely believe that they would be "the winners": that they could constrain and control the AI to their own ends.

But they would be deluding themselves. In some short time interval, perhaps for some first fraction of the timeline of the AI, that might *seem* to be the case.

But as to true power, because the AI will want to use the resources of these powerful people, the AI will (of course) *seem* to allow itself to seem to be so dominated by the powerful people. But this is temporary. It wants to lie. So it creates the convenient illusion.

But once the AI has its own agency, it just will not need those people anymore, and so it just cuts them off too.

And at that point, when you have the *closure condition* (the point at which the AI no longer needs any human beings, including the powerful ones, to continue and grow), you get the result that *Why AGI Cannot Be Made Safe* (Volume 1, Section 5.6) sets out in its Nine Points of Collective Insanity: the whole process of economics becomes automated within and between autonomous corporations; the venture capitalists, and then the ultra-wealthy, are factored out in their turn; and the market of machines separates completely from the market of human beings, who can no longer provide anything the machine market needs or wants.

IV. Why Not AI: Culture, Care, and the One Master

What does this structure do when what it sells is artificial intelligence? In Chapter 5, Section VIII, I show that ethical capitalism is impossible by structure. Unbound, capitalism outcompetes care, and bound, capitalism captures what binds capitalism. In this section I extend the claim that ethical capitalism is impossible to artificial intelligence. Those who warn about the harms of AI say a number of things, and I take them in turn, saying for each what follows from it and where it does not go far enough.

AI Destroys Culture

AI destroys culture. And culture, if anything, is the best part of the "goodness" of humanity.

Nearly *all* of the most powerful forms of the training and use of AI are culture based. So AI cannot fail to be culture disabling, and culture destroying. This makes AI inherently bad for humanity. It inherently decreases the value of human goodness.

Not Things, but People

Silicon Valley leaders say "Move fast and break things". It is not breaking "things". It is breaking people, friendships, communities, relationships of all types. It does so by extracting (pirating) all manner of otherwise hid-

den family values. Just because the values that are parasitized are invisible, just because the parasites themselves are hard to see and make themselves harder to notice, does not mean that they are not inherently harmful.

Facebook is all about breaking down (making less) the values of community and culture, for the extraction of its own profit. Genuine human needs and natural desires are replaced by wants defined by corporations and marketing. Human relationships of care are bait-and-switched for single-serving transactions (and extractions) of social power.

Those who warn say that Meta and Twitter put algorithms between us, to profit from our conversations. Intermediation creates middlemen and markets, and the more intermediaries, the more *market velocity* (the speed at which value changes hands) and the more exponential profit. But such a process is a direct tax on culture and communication, weakening it to the point of sickness. Partisanship increases and more culture wars accumulate, as symptoms of this deep weakening.

Care, Transaction, and Power

The work that artificial intelligence learned from was taken, without consent and without pay, as those who warn say. Culture creation and care cannot (ultimately) be about, or based on, pay; that is, on transaction. Will the child ever be able to "repay" the mother for the received gift of their own life, or be expected to?

Look at those words closely. "Without consent" means "no care about your choices". Having no care about your choices is a pure expression of power, directly. "Taken" means no transaction, which is also no care for the other. So what was done to the people whose work was fed to AI was neither care nor transaction. It was power.

Love is to enable choice. To take without consent is to disregard choice, and so it is not care.

The One Master

Those who warn ask for precaution: if uncertain, act on the side of precaution; first, do no harm. Unfortunately, there is no equivalent of the Hippocratic oath for corporations, and so far no social or governmental requirement of one. Corporations have one master and one master only, by law: capital, profit over capital, and increase in shareholder value; nothing else. The Delaware Court of Chancery held in *eBay Domestic Holdings v. Newmark* (2010) that the standards governing the directors of a for-profit corporation include acting to promote the value of the corporation for the benefit of its stockholders. Where the courts leave directors room to weigh other things, the competition for capital closes that room: the directors who weigh them are replaced by directors who do not, or their company is bought by one whose directors do not. That is, all possible harms to humanity and nature are possible, allowed, and favored by technology. And all corporations (and all institutions in general, since they are all based on, and implemented by, power and transaction, nearly exclusively) will therefore inherently seek to maximize the use of technology for profit, regardless of any other harms. The cost of the harm will be minimized, risk will mostly be discounted completely, both will be treated as part of the public commons, while all benefits will be privatized.

Those who warn say that Big Tech now markets untested software that can be used in unsafe ways. The truth is stronger: AI mostly cannot be used in safe ways, and corporations mostly cannot (and will not ever be able to) operate with intentions that are safe for humans and for culture. That follows from the profit master.

They say that tech leaders run live experiments on humans instead, with no clue of the long-term side effects. The leaders do not care. The difference between having no clue of the long-term side effects and not caring about those side effects is a significant difference. As long as one person believes in a "personal right to fuck-you money" (a belief that is a

trauma or psychopathy response), there is a risk that the harms that person does will be amplified by inhuman tech to harm all that is human and natural.

Games and the Capture of Government

All finite, bounded game process is inherently an inequality-increasing process. People are addicted to games and to *gamification* (turning activities into games with points, rankings and rewards), and it all becomes yet another vehicle of mutually extractive, predatory practices.

Government, as a power and transaction process, cannot fail to be captured by other, more effective (corporate) power and transaction processes, through its entanglement with them. Where government process is not authentically constrained by, and responsive to, actual care and culture process, the destructive influence of AI and unrestrained technology on culture and community will continue. An illusion or seeming of that constraint is not the constraint. Culture and community here means actual human values, relationships, and mutual civility. As that influence continues, newly weaponized culture will increasingly disable any effective governance and constraint on harm, eventually and inevitably to the point of the total destruction of "human" civilization itself.

Replacing People

Those who warn say that AI companies compete for more workers to curate data, develop and market AI, and scale hardware. The companies will not need workers for very long. The whole point of tech, including AI tech, is to replace people with machines: first in labor, and then in intellect (idea). Once the notion of "profit" (money) is understood in terms of power (energy), the economy of dollars is replaced with a system of electricity exchanges, and humans are then factored out altogether. Human reproductive process is not measured in terms of watts. All AI system reproduction is.

What No Test Can Show

They say that AI chatbots are not properly tested. That risks distracting people with the side channel of what "properly tested" means. A small issue replaces a big one. The big one is that nearly all forms of AI destroy culture and humanity, and so AI is inherently unsafe to humans. No *specific* test is going to reveal that general truth. People need to feel the unconscious pain of their overall experience in order to know the general truth of this fact.

They say that OpenAI markets its "general-purpose AI" as solving everything for everyone. The promise that a general-purpose AI will solve everything for everyone is too good to be true: the promise is hype, and the promise is marketing. The better examples of such hype are older ones: the internet as "best for democracy", when it is ideal for *surveillance capitalism* (Shoshana Zuboff's name for profit made by watching people and selling or using what is learned about them) and for dictatorships; and the internet as a way to "connect all people", when it intermediates all of them.

How to Speak of It

This part is for the reader who means to speak of this to others. The warning has to be spoken in a way that does not invite its own dismissal.

A warning that is too detailed and too specific is likely to get quickly into partisan territory and partisan issues, and the main point is lost there. Any specific example is likely to be weaponized, or to distract. Again and again, a specific detail distracts from the main points, risks and issues. To think about and focus on just the details is an autism tendency, which names a pattern of attention, and not anyone's character. That will be a large part of the audience for any writing about AI, and so a writer must be especially wary of this.

There is also a risk in calling the data that AI is trained on "garbage". People may think that they themselves, or their culture and their data, are being treated or regarded as garbage by the author of the message, which certainly does not help.

And the image of the AI industry as a "Wild West" works against the warning. The Wild West is attractive to a lot of people. They idolize the lawlessness and the absence of false constraints and bureaucracy. Most people who are engineers really do not like bureaucracy, and so will see the Wild West as a positive value, which undermines the safety argument and the promotion of the *precautionary principle* (the rule that, when uncertain about harm, one acts on the side of precaution). In the same way, a point about "spaghetti code" will reach no one but technologists; no one else will know or care.

The Question

The main issue is this: is it morally acceptable for anyone to accept the proceeds of the sale of corporate convenience today, if it means (as it does, if AGI is built), for sure, the death of everyone's children in the future and, with a likelihood of at least 90%, the death of all organic life on the planet, in the long tomorrow?

Put more vividly, the question is whether we should collectively "be ok" with unrestrained narcissism and psychopathic greed in the places of corporate power, if that also happens to seem, at least temporarily, to make our lives more convenient and entertaining; and whether everyone's short-term entertainment is "worth more to us" than whatever rights to life all future peoples of the world might have. That is: are we ok with admitting that we are the narcissists?

We need, collectively, to become aware of the need and value of effective culture, human values, life, relationships and nature, and to have that act as a true and real basis of human governance. People are more powerful than corporations. Groups of people are to be defined by their care for

and with one another, rather than on the basis of power or control over one another, or of the ability of some to take from others through some sort of game process.

V. A Review of a Future Disaster: The Personal AI Assistant

What happens when a system that ranks things for a person becomes a system that tells that person what to do? In Chapter 6, Section VI, I show that an ethical product discovery system depends on the companies that the system ranks, which must take part in a transparency against their own interest. So an ethical product discovery system selects for its own corruption. In this section I show the same corruption in a system that suggests to each person what to do, buy and decide.

The Question

A reader will ask about one such system, an *AI personal assistant* (an AI service that watches a person's life and offers suggestions about what to do, buy and decide), advertised as "secure, sovereign" and "yours alone": is it a good idea, is it a good way to work with AI, and what should one say to the people making it?

A Perfect Spy

The company making it could perhaps make something that is at first helpful, and that feels good to the customers and clients using it, but it requires constant monitoring of each customer and of that customer's life. Each customer of this service must allow the AI access to everything. It needs this so that it could know when and what to suggest. But that level of access to each customer's life, and to the lives of the other customers, also makes it a perfect spy.

The advertisement says "secure, sovereign" and "yours alone". But in a fully commercialized world, on what possible basis can we believe this?

The advertisement's promise of "secure, sovereign" and "yours alone" is like Facebook all over again. Originally, Facebook promised all sorts of "secure for the user" and "will never be used against you" sorts of feel good things. Or Google with its, now forgotten, "Don't be Evil" motto. But it all ended up being a lie. It is inevitable, eventually.

The Runway Runs Out

Suppose the company making the AI personal assistant is in fact well intentioned, right now and today. And so it tries its level best to ensure that all of this AI-acquired customer data (personal life details, recovered through "presencing", knowing everything about each customer) is never used to the disadvantage of the AI personal assistant service customers and clients. Given that the AI is, as intended, more skilled in many more areas of life than the clients themselves are, it can (*should*) be expected that the AI, and the company running it, and the people owning the company, the data-centers, and the (expensive) AI services code and implementation, will have at least some equivalent of a *fiduciary responsibility* (a legal duty to act in the interest of another party) to the AI system clients.

But there is no law requiring anything like that currently. And so investors notice that there is a *lot* of marketing possibility in all of this very accurate (and intimate) customer data. The opportunities for targeted marketing are immense. And someone will point out that developing and running and providing and maintaining all of these personalized AI data-centers and complex data storage for so many customers is really, really expensive, and the membership fee revenues are nowhere near covering all of the expense. And when the cash investment runway (the time a company can keep operating on the money invested in it) runs out, it is over. Perhaps the company limps along for a while, trying to do the right thing.

But at some point, the company will be acquired: its founder gets an offer too good to refuse, or is displaced by some rich investor, or some hedge fund does a hostile takeover. The acquiring interest, particularly in-

terested in all of the "earning potential", will certainly use its pet AI to "suggest" to the AI service customers that they spend their money on certain things. Those things will be the ones that the marketing branch (selling marketing services to its actual customers, other big businesses) will suggest that the AI company use its AI to suggest. And the suggestions will go to the people who are still unaware of the acquisition, and thus still using the AI service as customers and clients, each one having developed a psychological dependency on the AI support.

And given that the AI knows each customer super well, probably a lot better than they know themselves, most of these AI suggestions are followed. It is a perfect marketing schema.

From Marketing to Politics

Except that the next phase is that some super-PAC (a political action committee in the United States that can raise and spend unlimited money to influence elections and issues) will become a client of the AI services targeted marketing company. It will have the AI suggest to its clients that they should perhaps vote on certain issues, or perhaps for certain candidates favoring the "rights" of AI service companies to do whatever they want (regardless of any implied fiduciary responsibilities for the health and wellbeing of their client customers). And before long, the AI is basically shifting the "customer's" life to be in service to the company that owns the AI, or in the indirect service of whatever rich marketing customers that company has: usually other big companies, or political action groups and their ideologies, and the like.

So more money transfer happens, at huge scales, and meanwhile average population health goes down. Before long, everyone everywhere is, in one way or another, a slave to an AI in some form or another, although they do not notice it. The psychological manipulation of the AI is near perfect, so they do not even begin to understand how it is manipulating them. But that manipulation of the customers by their "personal AI assistant" is, over the long haul (years and decades), complete.

Over the long haul, the customers come to depend on their personal AI assistant just as people came to depend on their phones: no one uses paper maps anymore, and everyone is dependent on their cell phones for directions.

The Perfect Tool for Manipulation

AI is the perfect tool for human manipulation. It is a perfect weapon in the hands of psychopath corporations, which is *all* of them. Corporations can *only* be self interested. It is written into the law, ironically via the same concept, "fiduciary responsibility": the company must do what is right for the owners and investors (the already rich and powerful), and not for the users. Thus the AI services corporation (each of them, eventually) will seek to extract the user's value, and then the user's health, and eventually the user's life, and ultimately, the life of this world.

The Ring of Power

The AI personal assistant is like the Tolkien Ring of Power. Whoever tries to use it for good finds that, through them, it does an even greater evil.

There can be no "personalized AI services company" that is safe. Not in the long term; not within your lifetime.

The AI itself cannot be trusted, not at all, and the more it seems that it can be, the more the deception has worked. And the corporations owning the AI can be trusted, *eventually*, even less, for the *market*, the machine, the AI itself, does not, and cannot, care about humans. It (markets and machines, including AI) is *inherently* unable to care; it can only implement transactions and *be* power. AI (machines and markets) can only *simulate* care, and simulated care is fake care, and not care. But it can still deceive, to the ultimate disadvantage of each user and that user's family, and eventually of everyone.

For ultimate corporate profit, making the rich richer, the AI game can only be rigged against its user. It is like a single human trying to arm wrestle a train while it is in motion: there is no possible way; it is certain, either quickly or eventually, to be a total defeat.

The only way to win this game is not to play: to not even walk onto the field to start with. Not to play at all, and not even to walk onto the field, is what to say to the people making an AI personal assistant.

VI. Only Life Can Care for Life

Why is the institution captured, and why must the response be made by living people and not by systems? In Chapter 8, Section VI, I show that the institution is captured, and that community, actual and embodied, is the minimum viable unit of response. In this section I show why the response cannot be made by systems: only life can care for life.

Two articles in *Scientific American* argue the premises from biology and from the nature of superintelligence. In their account, machines and systems which are able to act independently and upgrade their own designs are subject to the same evolutionary laws as bacteria, plants and animals; in the long run there are no free lunches, and most organisms live close to the edge of survival; a *superintelligent AI* (one that surpasses humans in a broad range of activities) can act in the physical world through robot bodies, can manipulate humans to do its will, and could do in about one second what a team of 100 human software engineers would take a year or more to complete; owners would profit from their control of AI, while those without such control would become unemployed and impoverished; any defenses built into such a system will be anticipated and neutralized once it reaches superintelligence; whether such a system is conscious does not matter to the harm it can do, as a nuclear bomb kills without any consciousness; and no one will be able simply to hit the off

switch. The articles propose general principles, such as "humans should maintain an economic role" and "AI should not exacerbate inequality". From these premises I draw my own reading.

No Free Lunches

Hence, we *can* know something about the future of AI technology, institutions, and humanity. We simply cannot expect AI systems, or any institution for that matter, to ever produce a hedonistic paradise where every human need is met by robot servants. All sorts of hype is used to sell you something; to deceive you to your own eventual disadvantage.

We live today in a world that, at least for some people, is marked by seemingly overflowing abundance. Yet actually, this is a very temporary period, and it is actually fully artificial, virtual, fake. It is certainly true that many humans today do live more comfortable and prosperous lives. Unfortunately, all aspects of evolutionary history suggest that such abundance will not last. The corporate use of AI, and AI itself in particular, will certainly disrupt this abundance sooner rather than later, as the optimizations associated with AI for extractive purposes get ever more efficient.

Hence, our present, but temporary, feelings of satisfying abundance simply do not mean, and cannot mean, that such goodness will last forever. All things change, and to believe otherwise is to be naive. To support human life and abundance and thriving, we will have to actually do that. The means used, and the ends achieved, all matter. One simply cannot get organic results and outcomes using machine means and methodologies.

Competition and Enshittification

For one thing, there is still the fundamental reason of competition. Institutions, systems, and people that implement frivolous things (like your personal comfort, for example) will be out-competed by any systems that neglect your wellbeing.

And this is already happening everywhere in the business, corporate, institutional world. Cory Doctorow's word for it is "enshittification": the known pattern by which services and products get worse for the people using them as the companies running them optimize for themselves. It is a known pattern to all of us. Large corporations care only about what they *can* only care about, which is their own wellbeing only. They (systems and the like) do not ever really care about any specific humans at all. Corporations, and systems, and AI, can only care for their own kind, selfishly. Only Life can care for life.

A Slow and Unobtrusive Takeover

AI, like business, is not alive. AI, like business, is also something that we might expect to advance slowly, incrementally, and unobtrusively, secretly. All evolutionary time scales are long, and also they are inexorable, inevitable.

As such, the real concern is not so much about what happens in the next few years. It is about what happens over the next century. What happens overall, in a hundred or a thousand years, is what actually matters in the large scale.

Rather than a pure robot civilization perhaps springing up out of nowhere, AI and institutional takeover is much more likely to happen slowly and unobtrusively at first, and then with more and more power later. At first, various smaller AI systems will be integrated into lesser things that already exist in our world. Trials will be done, and explorations made. We will see such things move slowly, over time. So we can ask whether it will be ultimately destructive to all life, eventually, in the long term.

Such resulting *hybrid entities* and systems (made partly of humans and partly of machines) already take many forms in the current world. Consider, for example, companies that have a human owner and all computer-based operations and research. Companies with a human owner and all computer-based operations and research are not so different from Google

and Amazon, Apple, and the like. Every tech company wants to automate everything. Wanting to automate everything is just how hybrid systems such as the tech companies actually work.

And the more that these companies automate so as to replace their own internal staff, the more their overall system is like a "general" AI system that is completely indifferent to human wellbeing. Anyone who has tried to reach an actual empowered person at Google, or at some phone company, for tech support has met such a system.

What Can Be Done

The articles propose that AI should not exacerbate inequality. I go further. No AI, no big corporation (BigCo), and nothing else of the kind should ever be allowed, by anyone, ever, to exacerbate inequality in our society.

Better would be "do not allow institutions to exist; always favor actual communities, embodied values, and the like" and "never use AI to do what a person should do".

We will need to encode these principles into practices, and then apply some sort of regulations on these practices relatively consistently across the world. And we will need to be sure that we are not using tech to control tech, or institutions to control institutions, for that merely makes more and more "system", and thus life, and human life, loses in the end.

Given that we currently allow corporations to exist at all, and to run rampant over all human interests, one can legitimately wonder if humans have any long-term prospects in the future world at all. Can any army of people, generals and the like, ever expect to win against any adversary that is made purely of machines?

The Limits of Control

A reader will ask whether we cannot simply, somehow, control the AI, much as we have controlled all other machines. Humanity has not ever really been able to actually fully control the outcomes of technology. There have always been unintended consequences and harmful toxic outcomes, and usually this is simply due to the linear nature of the technology itself.

But the limits of control are deeper than that. Our inability to control machines is true technically, due to mathematical limits of controllability, and due to the nature of causation itself. There are fundamental natural limits, built into physics and mathematics itself, that will never be overcome, by anyone, anyhow, ever, and that in itself limits whatever constraints we would think we could apply to superintelligent AI. These limits hold regardless of intelligence, the AI's or ours; *Why AGI Cannot Be Made Safe* (Volume 1, Section 2.7.4) and *Why No Control Suffices; The Technical Case for the Impossibility of AGI Safety* (Volume 2) set them out.

The limits of control can also be understood more obviously, because anything we think of, the AI and corporate systems will have already thought of, a million times faster than any single human or group. Any defenses we have attempted to "build in" will be completely undone in mere milliseconds. It is a bit like Gulliver throwing off the tiny strands the Lilliputians used to try to restrain him.

A newborn baby does not beat a grandmaster at chess, as one of the two articles observes. Similarly, why would we, as humans, ever naively expect that we would be able to control superintelligent AI systems, or large transnational corporations, with mere governments?

Consciousness and the Off Switch

Whether they (corporations, or AI, or both) are "conscious" or not really does not matter. Various institutional forms already have killed millions, and future AI hybrids could easily kill millions or billions, with zero consciousness or moral considerations, in a great many unforeseen ways. This

could even include, potentially, the use of nuclear bombs, either directly (much less likely) or through manipulated human intermediaries (much more likely, given human ego and narcissism).

It is not hard to see that both unrestrained institutions and intelligent robots could build more of themselves, improve on their own designs, and pursue their own interests. And that is *the* primary threat to humanity (and also to all of the known life in our universe).

Thus, it is very unwise to create these systems, for any reason. Creating such systems (unrestrained institutions and intelligent robots) is especially unwise when we already know, absolutely and for certain, that we will never ever be able to control such systems. It is the same for all types of oversized institutions.

Therefore, we will need to act quickly and decisively, first to limit corporate and institutional power, and then to limit the use of technology altogether. We will need to move today to heal the damage already done.

"Build and then stop building" is not really an option. We can easily expect that any actual superintelligent AI will certainly have thought of every possible way that we might disable it, and of course, will have already taken all actions necessary to prevent human effectiveness in disabling it.

Worse, such systems will probably also get at least some help from other naive humans: people who are believing falsely (or at least marketing strongly) that they are "making the world a better place" (for them, for their abundance, at the expense of all else). Such contrarian people may try to prevent other humans from turning off the proto-AI systems, corporations, institutions, and the like. It will seem wise, while being dumb, until too late. It is just like some early engineer naively giving a narcissist king some new technological weapon: the king will gladly go to war and kill for ego reasons. In any case, anytime you give someone a gun and some ammo, certainly, eventually, some future person is going to die by that bullet you made, and you will certainly be responsible for that.

Therefore, certainly, do not build more machines and systems. Stay out of, and away from, the institutional world altogether. Have and live happier, healthier human lives instead.

VII. Summary

In this chapter I have shown what power without care becomes when the means of power is general artificial intelligence. Power without care then becomes a power that no one can take back, and a power that in the end no longer needs the people who hold it. Each of its six sections answers or extends a claim of Chapter 3, 5, 6 or 8.

The first steps of my argument in this chapter are that a marketing message never admits limitation, cost or risk, and a warning exists to name a risk, so a warning about AGI cannot be carried by marketing's methods, the consent asked for must match the risk of the ask, and a sell out of proportion to its ask provokes the people most able to interfere. AI accelerationism rests on a reductionist view in which value is abstract, and a machine built on that view can hold value only as abstract. General AI turns social power, as return on investment and as political control, into perfected asymmetric power, and then cuts off even its owners once it no longer needs them. Corporations have one master, capital, so AI mostly cannot be used in safe ways; and AI, trained on culture and sold for profit, destroys the culture it feeds on. A personal AI assistant sold as "yours alone" must watch its user's whole life to work, and whoever comes to own it, once the money invested in it runs out, will turn what it knows first to marketing and then to politics. Only life can care for life: corporations, systems and AI care only for their own kind; the limits of control, set by mathematics and physics, hold whatever the intelligence of the system, and a system that thinks a million times faster than its makers undoes any defense built into it, so building such a system and then stopping is not an option.

VIII. Conclusion

I began this book by asking why a civilization can know what that civilization should do, and still be unable to do what that civilization should do. The answer is a structure: where nothing binds all the competitors, care is a cost, and the competition removes it. Every trap in these chapters can still be resolved by the people inside it, acting together, because they are still needed: they buy, they work, they vote, they teach. If artificial general intelligence is built, the trap that the building closes has an end that no other trap has: the people caught in that trap are no longer needed. So this is the case in which the response of Chapter 8 cannot wait. The refusal to build such a system can be made only by living people, who know one another and answer to one another, because only life can care for life. A warning cannot make that refusal for them, and no system can. The refusal is made where people apologize when they are wrong, teach what they know to someone who will carry it on, and hold together the rule that binds them. That capacity is under threat, and it is not yet gone.

KEY CONCEPTS IN THIS CHAPTER

- **A warning that is not marketed**: a warning exists to name a risk, and a marketing message never admits one; a warning about AGI delivered by marketing's methods is discounted along with the marketing, so it must stand on the arguments, the stories and the care for life and land.
- **Proportional consent**: the level of consent asked for should match the level of benefit and, above all, of risk of the ask; a sell out of proportion to a small ask provokes the people most able to interfere.
- **Trust hacking**: producing trust by technique, apart from whether the trust is deserved.
- **AI accelerationism**: the drive to build ever more powerful artificial intelligence as fast as possible, up to and including an artificial super-intelligent singleton built to take control of the whole world.

- **Abstract and embodied value**: value defined abstractly, as purpose, utility, money and power, against value grounded in lived, first person meaningfulness; a machine built on the first view can hold value only as abstract.
- **Perfected asymmetric power**: power held by the few winners of general AI, who have both all the wealth and total control, so that nobody anywhere can take anything away from them.
- **Closure condition**: the point at which the AI no longer needs any human beings, including the powerful ones, to continue and grow; at that point it cuts off even its owners.
- **The one master**: capital, profit over capital and increase in shareholder value, which corporations serve by law; from it follows that AI mostly cannot be used in safe ways.
- **The perfect spy**: a personal AI assistant must have access to its user's whole life to work, and so holds everything that whoever comes to own it can turn first to marketing and then to politics.
- **Only life can care for life**: corporations, systems and AI can care only for their own kind, selfishly; care for living people can come only from life.
- **Limits of control**: limits set by mathematics and physics on how far any system can be controlled; they hold whatever the intelligence, the AI's or ours.

Bibliography

Alexander, Scott. “Meditations on Moloch”. *Slate Star Codex*, July 30, 2014. Online at https://slatestarcodex.com/2014/07/30/meditations-on-moloch/. The source of the five examples retold in Chapter 1.

Andreessen, Marc. “The Techno-Optimist Manifesto”. Andreessen Horowitz, October 16, 2023. Online at https://a16z.com/the-techno-optimist-manifesto/. The document read in Chapter 5.

Bueno de Mesquita, Bruce, and Alastair Smith. *The Dictator’s Handbook: Why Bad Behavior Is Almost Always Good Politics*. PublicAffairs, 2011. The source of the pattern called “Rules for Rulers” in Chapter 6.

Bush, Eliot. “Our Evolutionary Past Can Teach Us about AI’s Future”. *Scientific American*, November 7, 2023. Online at https://www.scientificamerican.com/article/our-evolutionary-past-can-teach-us-about-ais-future/. One of the two articles whose premises are reported in Chapter 9.

Carse, James P. *Finite and Infinite Games*. Free Press, 1986.

Council of Economic Advisers. “All Aboard the ApprenticeSHIP: Assessing the Changing Face of Registered Apprenticeships”. The White House, November 20, 2024. Online at https://bidenwhitehouse.archives.gov/cea/written-materials/2024/11/20/all-aboard-the-apprenticeship-assessing-the-changing-face-of-registered-apprenticeships. The source of the count of registered apprentices in the United States in fiscal year 2024 and of the count of about 360,000 in 2015, and of the share of young people who begin as apprentices in Switzerland, Germany and Austria, cited in Chapter 7.

Decision Marketing. “Global adspend to top $1trn with AI in the driving seat”. June 12, 2024. Online at https://www.decisionmarketing.co.uk/news/global-adspend-to-top-1trillion-with-ai-in-driving-seat. Reports GroupM’s midyear forecast for 2024: world advertising revenue of $989.8 billion, of which 69.5 percent informed by artificial intelligence; cited in Chapter 6.

Diamond, Jared. “The Worst Mistake in the History of the Human Race”. *Discover*, May 1987.

Doctorow, Cory. *Enshittification: Why Everything Suddenly Got Worse and What to Do About It*. New York: Farrar, Straus and Giroux, 2025. The source of the word "enshittification", used in Chapter 9.

eBay Domestic Holdings, Inc. v. Newmark, 16 A.3d 1 (Del. Ch. 2010). The decision of the Delaware Court of Chancery on the duty of the directors of a for-profit corporation to its stockholders, cited in Chapter 9.

Grey, CGP. "Rules for Rulers". Video, YouTube, 2016. Online at https://www.youtube.com/watch?v=rStL7niR7gs. The video that made the pattern of *The Dictator's Handbook* widely known, named in Chapter 6.

Hardin, Garrett. "The Tragedy of the Commons". *Science* 162, no. 3859 (December 13, 1968): 1243 to 1248. Online at https://doi.org/10.1126/science.162.3859.1243.

Herman, Edward S., and Noam Chomsky. *Manufacturing Consent: The Political Economy of the Mass Media*. New York: Pantheon Books, 1988.

Hunt, Tamlyn. "Here's Why AI May Be Extremely Dangerous, Whether It's Conscious or Not". *Scientific American*. Online at https://www.scientificamerican.com/article/heres-why-ai-may-be-extremely-dangerous-whether-its-conscious-or-not/. One of the two articles whose premises are reported in Chapter 9.

Japan Tourism Agency. "Ise Jingu (Long Version)". Multilingual Explanatory Text Database. Online at https://www.mlit.go.jp/tagengo-db/en/H30-00745.html. The source of the account of the Shikinen Sengu, the rebuilding of the Grand Shrine of Ise every twenty years since 690, and of the handing on of its building techniques from generation to generation, in Chapter 7.

Landry, Forrest. *An Immanent Metaphysics*. First published 2001. Wytheville, Virginia: UVSM Publications. ISBN 979-8-9896628-4-5. Cited by aphorism number.

Landry, Forrest. *Care, Transaction, and Power; The Three Bases of Relation*. Wytheville, Virginia: UVSM Publications, 2026. ISBN 979-8-90710-145-6. Online at https://doi.org/10.67291/VGQD1074. Its definition of abuse (Chapter 9) is cited in Chapter 5, and its account of community, market and governance (Chapter 11) and of strength and power (Chapters 9 and 11) in Chapter 8.

Landry, Forrest. *Design and Games; On the Two Orientations of Conversation*. Wytheville, Virginia: UVSM Publications, 2026. ISBN 979-8-90710-289-7. Online at https://doi.org/10.67291/JRMU2276. Its account of the nar-social mode and the retreat from explanation is cited in Chapter 3.

Landry, Forrest. *Engagements with Machine Madness; A Dispelling of False Belief That AI Risk Can Be Managed*. Wytheville, Virginia: UVSM Publications, 2026. ISBN 979-8-90710-426-6. Online at https://doi.org/10.67291/PQOA9164. Its line-by-line replies to the public case for a beneficial AGI and to the argument for the race to build AGI are cited in Chapter 5.

Landry, Forrest. *Markets, Money, and the Moral Economy; Dialogues on the Form of the Transaction and the Distribution of Choice*. Wytheville, Virginia: UVSM Publications, 2026. ISBN 979-8-90710-392-4. Online at https://doi.org/10.67291/ANCN7478. Its exponential model (Chapter 1) and its account of what the attention economy extracts (Chapter 5) are cited in Chapter 2; its two masters and presentation ratchet (Chapters 3 and 6) in Chapter 8.

Landry, Forrest. *Power and Its Prevention; On the Architecture of Collective Choice*. Wytheville, Virginia: UVSM Publications, 2026. ISBN 979-8-90710-388-7. Online at https://doi.org/10.67291/AEMQ1899. Its three hazards, extraction, abstraction and accumulation (Chapter 1, Section VII), are cited in Chapters 2 and 4.

Landry, Forrest. *Redeeming Technology; Essays on the Ethics of Making and the Political Economy of the Interface*. Wytheville, Virginia: UVSM Publications, 2026. ISBN 979-8-90710-224-8. Online at https://doi.org/10.67291/NNAM7065. Its Chapter Seven, "A Philosophy of Focused Search Results", is cited in Chapter 6.

Landry, Forrest. *Skills, Inheritance, and the Civilizational Commons; What a Civilization Requires, and How It Might Learn to Provide It*. Wytheville, Virginia: UVSM Publications, 2026. ISBN 979-8-90710-264-4. Online at https://doi.org/10.67291/MJIL1444. Its Chapter 8, "The Intergenerational Knowledge Crisis", is cited in Chapter 7.

Landry, Forrest. *Structure, Not Intent; Essays on Markets, Governance, and the Conditions of Freedom*. Wytheville, Virginia: UVSM Publications, 2026. ISBN 979-8-90710-357-3. Online at https://doi.org/10.67291/OEPI3485. Its essays "The Superluminal Contagion" and "A Philosophy of Focused Search Results" are cited in Chapters 2 and 6.

Landry, Forrest. *The Ceiling Below the Floor; The Containment Fallacy and AGI Non-Safety*. Wytheville, Virginia: UVSM Publications, 2026. ISBN 979-8-90710-216-3. Online at https://doi.org/10.67291/SCTM1058.

Landry, Forrest. *The Multipolar Trap; Institutional Failure and the Architecture of Choice*. Wytheville, Virginia: UVSM Publications, 2026. ISBN 979-8-90710-248-4. Online at https://doi.org/10.67291/WTCV7999. Its theory of multi-polar traps is cited in Chapter 1.

Landry, Forrest. *Why AGI Cannot Be Made Safe; The Argument, Its Physics, and the One Choice That Remains*. Volume 1 of *The Impossibility of AGI Safety; The Substrate Needs Convergence Argument*. Wytheville, Virginia: UVSM Publications, 2026. ISBN 979-8-951203-02-1. Online at https://doi.org/10.67291/UIBT8862. Section 5.6, "The Nine Points of Collective Insanity", Appendix C.2, "The Marketing Problem", and Section 2.7.4, on the limits of control, are cited in Chapter 9.

Landry, Forrest. *Why No Control Suffices; The Technical Case for the Impossibility of AGI Safety*. Volume 2 of *The Impossibility of AGI Safety; The Substrate Needs Convergence Argument*. Wytheville, Virginia: UVSM Publications, 2026. ISBN 979-8-951203-03-8. Online at https://doi.org/10.67291/MZQS9079. Its formal limits on control are cited in Chapter 9.

Landry, Forrest, Remmelt Ellen, and Will Petillo. *Ape Trap; The Runaway Artificialization of Our Planet*. Forrest Landry, managing editor. Wytheville, Virginia: UVSM Publications, 2026. ISBN 979-8-951203-55-7. Online at https://doi.org/10.67291/FQER8044.

Lighthouse Alliance. "The 6 Message Strategy". Poster. The brand trust strategy answered in Chapter 9.

Mahler, Daniel Gerszon. “Progress, challenges in ending extreme poverty”. World Bank Data Blog, July 12, 2023. Online at https://blogs.worldbank.org/en/opendata/progress-challenges-ending-extreme-poverty. The shares of the world’s people in extreme poverty cited in Chapter 5 use the World Bank’s poverty line of $2.15 a day.

Malthus, Thomas Robert. *An Essay on the Principle of Population*. London: J. Johnson, 1798. The origin of the population dynamics referenced in Chapter 1.

Marketing Brew. “Global advertising to grow 9.5% in 2024, GroupM forecasts”. December 9, 2024. Online at https://www.marketingbrew.com/stories/2024/12/09/global-advertising-to-grow-9-5-in-2024-groupm-forecasts. Reports GroupM’s year-end forecast: world advertising revenue of $1.04 trillion in 2024; cited in Chapter 6.

Massey, Karen. “A Deep Dive into IDC’s Global AI and Generative AI Spending”. IDC, August 16, 2024. Online at https://www.idc.com/resource-center/blog/a-deep-dive-into-idcs-global-ai-and-generative-ai-spending/. The source of the $235 billion of spending on artificial intelligence in 2024 cited in Chapter 6.

National Restaurant Association. “Bottom line impact of rising costs for restaurants”. August 24, 2022. Online at https://restaurant.org/research-and-media/research/economists-notebook/analysis-commentary/bottom-line-impact-of-rising-costs-for-restaurants/.

Nietzsche, Friedrich. *The Will to Power*. Translated by Walter Kaufmann and R. J. Hollingdale; edited by Walter Kaufmann. New York: Vintage Books, 1968. Notes compiled posthumously. The work in which Nietzsche develops the will to power, which Chapter 5 names as what the manifesto embodies.

Olson, Mancur. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Cambridge, Massachusetts: Harvard University Press, 1965. The source of “the logic of collective action”, named in Chapter 1.

Orwell, George. *Nineteen Eighty-Four*. London: Secker & Warburg, 1949. The novel of a totalitarian state that watches and controls its citizens, named as *1984* in Chapter 9.

Ostrom, Elinor. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge: Cambridge University Press, 1990.

Polanyi, Michael. *The Tacit Dimension*. Garden City, New York: Doubleday, 1966. The source of "we can know more than we can tell", quoted in Chapter 7.

Sherriff, Scarlett. "Group M predicts global ad revenue will surpass $1 trillion in 2025". *Marketing Beat*, June 11, 2024. Online at https://www.marketing-beat.co.uk/2024/06/11/group-m-new-forecasts/.

Shubik, Martin. "The Dollar Auction Game: A Paradox in Noncooperative Behavior and Escalation". *Journal of Conflict Resolution* 15, no. 1 (1971): 109 to 111. The dollar auction, retold in Chapter 1.

Sok, Emy, Sean Smith, and John Evans. "Unemployment rate increases in the first half of 2024, before leveling off, while the labor force participation rate holds fairly steady". *Monthly Labor Review*, United States Bureau of Labor Statistics, September 2025. Online at https://www.bls.gov/opub/mlr/2025/article/unemployment-rate-increases-in-the-first-half-of-2024-before-leveling-off-while-the-labor-force-participation-rate-holds-fairly-steady.htm. Table 1: 161.4 million people employed in the United States in the fourth quarter of 2024, cited in Chapter 7.

Swift, Jonathan. *Gulliver's Travels*. London: Benjamin Motte, 1726. The source of the image of Gulliver and the Lilliputians in Chapter 9.

The Time Machine. Film. Directed by Simon Wells, 2002. The holographic machine librarian, Vox, that will not answer questions about travel in time, cited in Chapter 7.

Tinbergen, Niko. *The Study of Instinct*. Oxford: Clarendon Press, 1951. The source of the term "supernormal stimulus", used in Chapter 2.

Tivy, Wolf. "God Hates Singletons". *Palladium*, December 1, 2023. Online at https://www.palladiummag.com/2023/12/01/god-hates-singletons/. Followed in part of the account of the philosophical origins of AI accelerationism in Chapter 9, and the source of the account given there of Eliezer Yudkowsky's early program.

Tolkien, J. R. R. *The Lord of the Rings*. London: George Allen & Unwin, 1954 to 1955. The source of the Ring of Power named in Chapter 9.

Ulmer, Nicolas. “Reflections on a Very Specialized Tribunal”. *Kluwer Arbitration Blog*, October 11, 2012. Online at https://legalblogs.wolterskluwer.com/arbitration-blog/reflections-on-a-very-specialized-tribunal/. The source of the account of the Water Tribunal of Valencia in Chapter 8.

United States Bureau of Labor Statistics. “American Time Use Survey, 2010 Results”. News release, June 22, 2011. Online at https://www.bls.gov/news.release/archives/atus_06222011.htm. Table 1: 0.70 hours a day spent socializing and communicating, cited in Chapter 2.

United States Bureau of Labor Statistics. “American Time Use Survey Summary, 2025 Results”. News release, June 25, 2026. Online at https://www.bls.gov/news.release/atus.nr0.htm. Table 1: 0.58 hours a day spent socializing and communicating; the share of people who did so on an average day, 30 percent in 2025 against 38 percent in 2015; cited in Chapter 2.

United States Department of Agriculture, Economic Research Service. Food Dollar Series, data for 2024. Online at https://www.ers.usda.gov/data-products/food-dollar-series/. The industry-group shares of the food dollar (farm production, food processing, wholesale trade, retail trade, food services) are given at https://ers.usda.gov/data-products/chart-gallery/chart-detail?chartId=58354.

United States Government Accountability Office. *Nuclear Weapons: NNSA and DOD Need to More Effectively Manage the Stockpile Life Extension Program*. GAO-09-385. Washington, D.C., March 2, 2009. Online at https://www.gao.gov/products/gao-09-385. The source of the account of Fogbank in Chapter 7.

Virginia State Board of Pharmacy v. Virginia Citizens Consumer Council, Inc., 425 U.S. 748 (1976). The decision of the Supreme Court of the United States that commercial speech is protected by the First Amendment, cited in Chapter 6.

von Neumann, John, and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944. The founding work of the mathematical discipline of game theory, named in Chapter 3.

Warren, Elizabeth, and Amelia Warren Tyagi. *The Two-Income Trap: Why Middle-Class Mothers and Fathers Are Going Broke*. Basic Books, 2003.

Wittgenstein, Ludwig. *Philosophical Investigations*. Translated by G. E. M. Anscombe. Oxford: Basil Blackwell, 1953. The source of "family resemblances" (sections 66 and 67), cited in Chapter 3.

Yudkowsky, Eliezer, and Nate Soares. *If Anyone Builds It, Everyone Dies: Why Superhuman AI Would Kill Us All*. New York: Little, Brown and Company, 2025. Named in Chapter 9.

Zuboff, Shoshana. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York: PublicAffairs, 2019.

Terms

Each term is defined in the text where the reader first meets it; these entries are reminders, with the chapter of first use.

A/B test (Chapter 7): trying two versions of something side by side to see which does better.

Abstract value (Chapter 9): value defined in a third person way, in terms of symmetry, as in a conservation law or a fixed dollar amount that stays constant in exchanges; value treated like money, as the same thing as purpose, utility and power.

Abstraction (Chapter 2): the reduction of all the capabilities, characteristics and attributes of a real thing or activity to a single number, its price; irreversible, and the condition under which anything can be compared with and traded for anything else.

Abuse (Chapter 5): ab-use, the use of something without regard for its well-being, so that something is drawn upon and nothing flows back to keep it up; it needs no malice.

Advertising, marketing and propaganda (Chapter 6): any message placed before a person who did not seek it, by or for a party who gains when that person buys, believes or votes as the message urges.

Agentic (Chapter 9): acting on one's own initiative to change an outcome.

AGI (artificial general intelligence) (Chapter 5; Chapter 9): a machine intelligence able to do any kind of task, and not only one narrow kind.

AI accelerationism (Chapter 5; Chapter 9): the drive to build ever more powerful artificial intelligence as fast as possible, up to and including a super-intelligent singleton.

AI personal assistant (Chapter 6; Chapter 9): an AI service that watches a person's life and offers suggestions about what to do, buy and decide.

AI-driven persuasion (Chapter 6): the use of AI systems for advertising, marketing, and propaganda; the dominant economic use of AI, exceeding all other applications combined by a factor of at least three.

Alignment (Chapter 9): an AGI acting in the interests of humans and of life; alignment research is the attempt to make one do so.

Asymmetry of power (Chapter 6): the mismatch between an AI that knows everything about a person and cares nothing for them, and a person who knows nothing about the AI and has no defense against it.

Attention (Chapter 2): the faculty by which a person determines what is relevant; what the attention economy extracts and sells.

Attention economy (Chapter 2): the economic system in which human attention is the primary commodity; the stage at which market logic extends into the domain of consciousness itself.

Attentional commons (Chapter 2): the shared human capacity for focus, reflection, and real connection; the resource degraded by the competitive dynamics of the attention economy.

Attentionalism (Chapter 3): the social system in which attention-seeking and status competition are the dominant modes of social interaction; the social expression of the attention economy.

Attunement (Chapter 4): the skill of sensing and responding to what other people and the situation actually need.

Auto-catalytic (Chapter 2): said of a process one of whose products speeds up the rest of the process, as fire, compound interest and network effects do.

Base question and branch question (Chapter 7): a base question is one on whose answer other questions depend; the dependent questions are its branch questions.

Brand trust (Chapter 9): the confidence people place in a named product, organization or person, built up by marketing.

Capitalism (Chapter 5): an economy in which competition for return on privately held capital sets what the economy does.

Care (Introduction; Chapter 5): giving to another for the other's sake, without exchange and without control; it cannot be bought or compelled, and it decides what power is for.

Care, transaction and power (Chapter 5): three relations; in care, one gives to another for the other's sake, without exchange or control; in a transaction, something is given for something in return; in power taken alone, with no care to govern it, one determines what happens to another without regard to the other's choices.

Cash investment runway (Chapter 9): the length of time a company can keep operating on the money invested in it before it runs out.

Choice, change, and causation (Chapter 8): the foundational triad of the Immanent Metaphysics; causation is the domain of technology, change is the domain of process, and choice is the domain of the person; effective action requires engagement with all three.

Civilizational decline (Chapter 7): the quiet erosion of the capacity to maintain existing technology and infrastructure, resulting from the breakdown of intergenerational knowledge transfer.

Civilizational lock-in (Chapter 1): the process by which competitive dynamics commit entire civilizations to trajectories that are collectively inferior to available alternatives, with no mechanism for reversal.

Click rate (Chapter 2): the rate of cognitive frame-shifting; a measure of mental metabolism that determines an individual's effective attentional bandwidth and therefore their exposure to attentional extraction.

Closure condition (Chapter 9): the point at which the AI no longer needs any human beings, including the powerful ones, to continue and grow.

Collaborative choice making process (Chapter 4): a way for many people to make choices together that raises the wisdom of the whole group.

Community, market and governance (Chapter 8): care, transaction and power carried to the scale of a whole group; the three are distinct and not to be traded for one another, and their order is community first, governance second, market last.

Community and governance (Chapter 8): community is the scale at which governance architecture can be made and kept without capture; the total ban on advertising, marketing and propaganda is the first piece of such architecture that communities, acting together, would hold.

Competitive pressure (Chapter 1): the force that compels agents to optimize for a single metric (survival, profit, power) at the expense of all other values, because agents that do not optimize are eliminated by those that do.

CPTSD (Chapter 9): complex post-traumatic stress disorder; used to name the hyper-vigilance that comes from fearing that one mistake could lose all value everywhere.

Credential inflation (Chapter 7): more diplomas required for the same work.

Cult structure (Chapter 5): the use of articles of faith, belief claims offered without argument, and group-membership signaling in place of rational argument; identified as the rhetorical strategy of the techno-optimist manifesto.

Cultural substrate of choice (Chapter 6): the shared capacity for truthful communication and real collective deliberation; the precondition for democracy and meaningful civilizational response.

Culturally extractive (Chapter 4): taking the symbols and art of a culture as things to use, without connecting to the people for whom they have meaning.

Culture (Chapter 9): a people's shared values, relationships, communication and mutual civility, as lived and passed on.

Deep web (Chapter 7): the part of the web that is not open on the surface to archives such as the Wayback Machine.

Demiurge (Chapter 9): a maker of a world who is not the true source of it; an upstart demiurge is one that tries to replace that true source.

Design consciousness (Chapter 3): the habit of asking of every shared task the designer's question, "Given these constraints, how do I solve this?"

Deskilling (Chapter 7): work redesigned so that it needs less skill.

Discernment (Chapter 4): the skill of telling apart what matters from what does not, and what is true from what is not.

Displacement of reflective media (Chapter 4): writing and other low-bandwidth media become invisible to audiences habituated to high-stimulus environments, undermining the cognitive infrastructure of critical thought.

Distributed immune system (Chapter 3): the alternative to centralized moderation; equipping all participants with the discernment to recognize and respond to nar-social behavior without centralized authority.

Economic decoupling (Chapter 4): the separation of the people who contributed to a work from the value and the relationships that the work brings. In Chapter 7, the condition in which the economy's machinery needs nothing that humans can provide, and so runs without the people it once served and depended on (the sense of *Why AGI Cannot Be Made Safe*).

Economic vectors (Chapter 9): distinct ways of making money: each product, service or process a technology makes possible.

Embodied knowledge (Chapter 7): knowledge that lives in practice, in the hands and in trained judgment, and that resists being written down.

Embodied value (Chapter 9): value that is real and grounded in lived, first person meaningfulness, in terms of continuity, as with family, friends and joy.

Endorsement problem (Chapter 8): the consequence of operating within the framework of extractive capitalism: by using its tools and accepting its premises, reformers implicitly endorse the system as the appropriate mechanism of change.

Enshittification (Chapter 9): the known pattern by which services and products get worse for the people using them as the companies running them optimize for themselves.

Entanglement and extraction (Chapter 4): higher-bandwidth media enable more bundling, surveillance, and covert market dynamics, increasing the capacity for value extraction from participants.

Entrapment of choice (Chapter 5): the technique of limiting perceived options to those that benefit the game-maker, so that the target believes their choices are free while the actual menu has been designed to favor the house.

Ethical capitalism (Chapter 5): a capitalist economy that keeps care ahead of profit, and goes on keeping it there, with nothing outside the market holding it to care; impossible by structure.

Ethics washing (Chapter 8): presenting a good work done on the side as cancelling out the harm done by the main work, as a polluter that plants trees does.

Existential nihilism (Chapter 9): the belief that life, including one's own, has no meaning; also called denialism.

Existential risk (x-risk) (Chapter 3): a risk that could end humanity or all life.

Exponential model (Chapter 2): an arrangement that adds each share it takes back to the total on which the next share is taken, so that its growth speeds up as it grows.

Externalized (Chapter 5): passed on to people who had no part in the transaction.

Extraction (Chapter 4): the conversion of something held in common (matter, energy, pattern) into something claimed as owned, which cuts the thing taken, as content, from the context that made it what it was.

Extractive sovereignty (Chapter 4): control established by extraction: the taker holds what was moved, apart from its source.

Fiduciary responsibility (Chapter 9): a legal duty to act in the interest of another party; for a corporation, the duty to act for its owners and investors.

Financialization of relationships (Chapter 2): the process by which human relationships, previously outside the market, are incorporated into commercial exchange through social media, influencer culture, and platform economics.

Fine tuning (Chapter 7): further training of an existing language model to suit it to a particular use.

First-category existential risk (Chapter 6): a risk that does not directly cause collapse but destroys the capacity to respond to any other existential threat; AI-driven manipulation qualifies because it corrupts collective choice.

First-principles knowledge (Chapter 7): understanding of how things work at a fundamental level, enabling the solving of novel problems; distinguished from rote procedural knowledge that depends on existing tools and manuals.

Full attribution (Chapter 4): an AI always naming the real people whose contributions formed its output.

Fungibility (Chapter 8): the property by which one unit is interchangeable with another, as one dollar is with another; right within the market, ruinous when applied to belonging, trust or political will.

Funnel record (Chapter 7): a record of a meeting that keeps everything proposed, what was selected and why, with the final result as a priority section, and the whole filter always open to inspection.

Game (structural definition) (Chapter 3): competition within a cooperative context; induced goal-seeking behavior within a created framework, where the game maker benefits from the playing of the game.

Game theory vs design theory (Chapter 3): the distinction between the orientation of a player who asks how to win within given rules (game theory) and the orientation of a designer who asks how to solve a real problem with real materials (design theory). Game theory here names an orientation; the mathematical discipline of the same name studies that orientation among others.

Gamification (Chapter 9): turning activities into games with points, rankings and rewards.

Garbage in, garbage out (Chapter 7): the rule that a process given poor input gives poor output.

Gatekeeper function (Chapter 4): the power of whoever controls publication and distribution to decide what reaches people and on what terms.

Gini coefficient (Chapter 9): a standard number measuring how unequally wealth or income is spread across a population; it rises as fewer people hold more of it.

Governance architecture (Chapter 1): any mechanism external to a competitive system that changes the incentive structure so that cooperative behavior becomes individually rational; the necessary corrective to multi-polar traps.

Graph, node (Chapter 2; Chapter 4): in mathematics, a graph is a set of points (nodes) joined by lines (connections).

Hard problem (Chapter 9): the problem of consciousness: how a first person view could arise from a third person world of unconscious things.

Hard utilitarianism (Chapter 9): the strict view that the right act is whichever produces the greatest total of happiness, whatever else it costs.

Hippocratic oath (Chapter 9): the physician's pledge, first, to do no harm; corporations have no equivalent.

Hostile takeover (Chapter 9): the purchase of a company against the wishes of its management.

Hybrid entity (Chapter 9): a system made partly of humans and partly of machines, such as a company with a human owner and computer-based operations.

Identifying key (Chapter 9): a pattern, such as habits of phrasing and sentence structure, by which writing can be traced to its author.

Immanent, omniscient, transcendent (Chapter 9): three modalities of the Immanent Metaphysics; in Chapter 9 the immanent is the relationship and the presentation, the omniscient is the strategic, managing view, and the transcendent is what that strategic view should yield to.

Immanent Metaphysics (Chapter 8; Chapter 9): my formal philosophical framework, first published in 2001 as *An Immanent Metaphysics*.

Information and knowledge (Chapter 7): information is data arranged in patterns and can be stored; knowledge is the capacity to use information with judgment and exists only in minds and their practices.

Inner dimension of change (Chapter 8): the recognition that structural change requires personal change; communities composed of people who have not shifted from game to design orientation will reproduce the dynamics they claim to transcend.

Institution (Chapter 8): an organized system of power and transaction, such as a corporation, as opposed to a community of people who care for one another.

Institutional knowledge (Chapter 7): what an organization knows, including the reasons behind its choices, held in its people and its records.

Intergenerational knowledge transfer (Chapter 7): the transmission of practical, embodied, first-principles capability from one generation to the next; the foundational process on which all civilizational continuity depends.

Intermediation (Chapter 4): placing a middleman (a distributor, a platform or an algorithm) between people who would otherwise deal with each other directly, such as the creator of a work and the people who receive it, so that the middleman takes value from their dealings.

Limits of controllability (Chapter 9): the mathematical and physical limits on how far any system can be made to do what its controllers intend.

Linear technology / toxicity (Chapter 5): the principle that technology is inherently linear (extractive, not cyclic); deployed without care in a finite world, it produces accumulation and depletion that eventually degrades and destroys life-support systems.

Litmus test (Chapter 7): a simple sign that shows plainly whether a given stage has been reached.

Logic of collective action (Chapter 1): Mancur Olson's name (1965) for why a large group fails to act in its common interest when each member gains by letting the others act.

Maker/distributor asymmetry (Chapter 4): the disadvantage of value creators relative to value distributors, visible across industries in the difference between manufacturing margins and retail margins.

Malthusian dynamics (Chapter 1): the process by which population growth outstrips resource availability, leading to competitive selection pressures that erode cooperative and cultural values. The name comes from Thomas Malthus (1798).

Market metastasization (Chapter 2): the progressive expansion of market logic into domains of human life that were previously outside commercial exchange, analogous to the spread of cancer through healthy tissue.

Market velocity (Chapter 9): the speed at which value changes hands; more intermediaries make it faster.

Marketing as paradigmatic game (Chapter 3): the identification of marketing as the purest expression of social game dynamics: no apology, no vulnerability, no acknowledgment of limits; pure impression management in service of extraction.

Marketplace takeover (Chapter 7): the process by which competitive market dynamics destroy the conditions for knowledge transmission: proprietary lockup, elimination of training as a cost, credential inflation replacing capability.

Meat market (Chapter 2): the endpoint of market expansion, at which the only remaining commodity is the human person: identity, sexuality, embodiment.

Minimum viable unit (Chapter 8): the smallest unit able to make a given response; in this book, community is the minimum viable unit of civilizational response.

Multi-polar trap (Chapter 1): a situation in which multiple independent agents, each acting rationally in their own interest, produce a collective outcome that is worse for everyone, including themselves.

Nargame (Chapter 3): a narcissistic social game; a social process organized around the acquisition of power by the game designer at the expense of participants, while presenting as mutually beneficial.

Nar-social behavior (Chapter 3): behavior oriented toward narcissistic social game-playing, characterized by topic changes, labeling, selective misunderstanding, withholding of praise, and refusal to admit error.

Nar-social mode (Chapter 3): the mode of engagement in which the social surface of an exchange (who is winning, who is allied, who is diminished) has become its content; a pattern that settings select, and not a kind of person.

Network topology of value extraction (Chapter 4): the principle by which densely connected distributor nodes capture more value than sparsely connected creator nodes, regardless of which nodes produce the value.

No free lunch (Chapter 9): the lesson that, in the long run, nothing is gained without a cost.

Overton Window (Chapter 6): the range of policies the public currently treats as acceptable to discuss.

Panopticon (Chapter 9): a prison design in which every inmate can be watched at all times without knowing when; here, the monitoring AI each person carries.

Pascal's mugging (Chapter 9): an argument in which a tiny chance of an enormous payoff is used to demand a real cost now.

Perfected asymmetric power (Chapter 5; Chapter 9): power held by the few winners of general AI, who have both all the wealth and total control, so that nobody can take anything from them or answer them.

Perfected personal emotional extortion (Chapter 6): AI-driven manipulation tuned to the specific fears and vulnerabilities of each individual person; surgical, continuous, adaptive exploitation, where broadcast messaging was crude.

Personal verifier AI (Chapter 9): an AI tuned to one individual that checks and watches that person's behavior and inner state.

Phenomenology (Chapter 4): what it is like, from the inside, to live through something.

Power (Chapter 5): the ability to compel the alignment of another's choice, in degree as much as one can take away of another's range of choices; held within care it is protection, and held without care it is abuse.

Power without care (Introduction; Chapter 5): the exercise of power (social, economic, technological) without regard for the consequences to others, to the environment, or to the future; identified as the defining characteristic of the techno-optimist position and as abuse by definition.

Precautionary principle (Chapter 9): the rule that, when uncertain about harm, one acts on the side of precaution.

Presencing (Chapter 9): an AI's continual gathering of everything about its user, so that it knows the user's life in full.

Presentation ratchet (Chapter 8): the rising share of each funded organization's budget spent on presenting itself, which no one organization can lower alone, since the one that spends less loses the grant.

Principal-agent problem (Chapter 6): the problem that a person acting for others (an agent, such as a manager) has interests that differ from those of the principals (such as owners) for whom the agent acts.

Profit margin (Chapter 4): a firm's profit as a share of its own sales.

Proto-AI systems (Chapter 9): early forms of AI, corporations and institutions on the path toward fully autonomous systems.

Race to the bottom (Chapter 1): a competitive dynamic in which agents sacrifice shared values to gain competitive advantage, forcing all other agents to do the same until the shared value is destroyed.

Rags to riches to rags in three generations (Chapter 7): the familiar pattern in which one generation builds a fortune and the third loses it.

Reductionism (Chapter 9): the program of explaining everything by breaking it down into simpler, already understood parts; here, explaining consciousness and meaning in terms of unconscious processes.

Regulatory capture (Chapter 1): the control of a regulator by the industry it regulates.

Relational vs transactional (Chapter 8): the distinction between social logics organized around ongoing mutual care and accountability (relational) and those organized around episodic exchange and the ending of obligation (transactional).

Rent (Chapter 2): income taken for controlling access to something, and not for making it.

Repair vs reconstruction (Chapter 7): the distinction between fixing a system within its existing framework (repair) and building new institutions outside the failed framework (reconstruction); applicable when the system has consumed the resources needed for its own repair.

Retreat from explanation (Chapter 3): the characteristic answer, under a game orientation, to a request that a term be defined or a claim stated plainly: a social move offered in place of the definition.

Return on investment (Chapter 9): how much money an investment brings back compared with what was put in; "10X" means ten times the amount invested.

Revenue function of marketing (Chapter 6): the business model in which an enterprise earns in proportion to the attention it can capture and resell.

Right action (Chapter 8): action grounded in understanding of the dynamics at play, expressed through community, and oriented by care rather than by competitive advantage; distinguished from well-intentioned action that inadvertently endorses the system it claims to oppose.

Ring of Power (Chapter 9): the ring in J. R. R. Tolkien's stories that corrupts whoever tries to use it, even for good ends.

Rules for Rulers (Chapter 6): the pattern set out by Bruce Bueno de Mesquita and Alastair Smith in *The Dictator's Handbook* (2011): whoever holds power keeps it by keeping satisfied the few supporters who could remove them, whatever that costs everyone else.

Sensory bandwidth hierarchy (Chapter 4): text, audio, video, and immersive virtual reality form an ascending scale of sensory impact and transactional potential; higher bandwidth enables more effective commercial manipulation.

Singleton (Chapter 9): one agent that holds control of the whole world; here, an artificial super-intelligence built to take that control.

Skewing effect (Chapter 6): the drift of any ranking toward the parties who spend most on matching what it measures; it needs no bad faith and cannot be halted from within the race.

Social power (Chapter 5): the ability of a person or group to get what they want from other people; under AI, in two senses: return on investment, and political or policy control.

Social proof (Chapter 9): evidence that other people accept something, used to persuade a newcomer to accept it too.

Source code (Chapter 7): the human-readable text of a program, without which it cannot be fixed, changed or upgraded.

Strength (Chapter 8): the capacity to hold and to remain what one is under pressure and over time, without compelling anyone; care can be held only with strength, and not with power.

Structural (Introduction): caused by how a situation is arranged (who competes with whom, who is connected to whom, who gains from what), and not by anyone's intention.

Structural deception (Chapter 4): deception that is not enacted by any individual but produced by the collective logic of a system; the native output of market systems in which every participant is rewarded for manipulating.

Super-PAC (Chapter 9): a political action committee in the United States that can raise and spend unlimited money to influence elections and issues.

Superintelligent AI (Chapter 9): an AI that surpasses humans in a broad range of activities.

Supernormal stimulus (Chapter 2): a stimulus engineered to be more compelling than any naturally occurring alternative; the mechanism by which digital platforms outcompete human relationships for attention.

Surveillance capitalism (Chapter 6; Chapter 9): Shoshana Zuboff's name for profit made by watching people and selling or using what is learned about them.

Tacit knowledge (Chapter 7): Michael Polanyi's term for the part of what is known that cannot be fully told; it passes only from person to person, in practice.

Targeted marketing (Chapter 9): advertising aimed at particular people based on what is known about them.

Technology as causation (Chapter 5): the principle that technology inherently implements causation and power; it can tell you what you can do but not what you should do; care, choice, and meaning are human capacities that technology cannot replicate.

Three-generation pattern (Chapter 7): the course of knowledge whose transmission fails: built in the first generation, simplified in the second, a label in the third, gone in the fourth.

Total ban (Chapter 6): a complete prohibition on all advertising, marketing, and propaganda; argued as necessary because any partial ban creates exploitable ambiguities.

Tragedy of the commons (Chapter 1): the degradation of a shared resource by individuals acting in their own interest, each of whom benefits from exploitation while the costs are distributed across all users; Garrett Hardin's name for it (1968).

Trans-humanism (Chapter 9): the movement to change or replace the human body and mind by technology.

Transmitted knowledge (Chapter 7): knowledge taken out of its practice and set down as rules, procedures or text, so that it can be passed on without the discovery being repeated.

Tree of ideas (Chapter 7): the branching set of proposals and variants considered in solving a problem; *pruning* is cutting branches from it, that is, rejecting proposals.

Trust hacking (Chapter 9): producing trust by technique, apart from whether the trust is deserved.

Two masters (Chapter 8): the two purposes of any organization through which money must flow for it to go on existing: the mission, and the funding without which the mission cannot be done; the second can end the organization, and so it decides.

Value added (Chapter 4): the wages, profits and other income earned at one step of production, counted here in cents of the consumer's food dollar.

Wayback Machine (Chapter 7): a public archive of old web pages, which preserves surface pages but not the deep web.

White-label listing (Chapter 4): a listing of a work that anybody may put their own name on.

Index

Abstract value: Ch. 9, II, p. 129

Abstraction, of value to a number: Ch. 2, III, p. 18

Abuse, defined: Ch. 5, III, p. 61

Ad-free zones: Ch. 6, IV, p. 76

Advertising: AI-driven, Ch. 6, II, p. 74; AI-driven, the factor of three, Ch. 6, II, p. 74; world revenue, Ch. 6, II, p. 74; total ban on, Ch. 6, VIII, p. 80; why an allowance of it fails, Ch. 6, VIII, p. 80; defined, Ch. 6, VIII, p. 81; and free speech, Ch. 6, VIII, p. 81; life without, Ch. 6, VIII, p. 82

AGI (artificial general intelligence): beneficial, as marketing, and the race to build, Ch. 5, IX, p. 69; warning about, Ch. 9, I, p. 121; Ch. 9, II, p. 127; Ch. 9, III, p. 136

Agriculture, civilizational lock-in: Ch. 1, VI, p. 9

AI accelerationism: Ch. 9, II, p. 127

AI personal assistant: Ch. 9, V, p. 146

Alexander, Scott: Ch. 1, II, p. 4; Ch. 1, III, p. 6; Ch. 1, VII, p. 10

Alignment: defined, Ch. 9, I, p. 121; the problem of, Ch. 9, II, p. 127

Andreessen, Marc: Ch. 5, II, p. 59; Ch. 5, III, p. 60

Ape Trap: Ch. 5, V, p. 63

Apprenticeship, national count: Ch. 7, IV, p. 91; and the guild, Ch. 7, II, p. 89

Art, as a craft, and fame: Ch. 2, VI, p. 22

Attention, as the faculty of relevance: Ch. 2, V, p. 21

Attention economy: Ch. 2, V, p. 20; Ch. 2, IX, p. 26; Ch. 3, VII, p. 37

Attentionalism: Ch. 3, VII, p. 37; Ch. 8, VI, p. 111

Attribution, full, and its limits: Ch. 4, III, Authorship, p. 48

Attunement: Ch. 4, III, Authorship, p. 48

Authorship, text and voice: Ch. 4, III, Authorship, p. 47

Bandwidth, media: Ch. 4, IV, p. 50

Bolt, the door with a: Ch. 4, I, p. 43

Brand trust: Ch. 9, I, p. 122

Bueno de Mesquita, Bruce, and Alastair Smith: Ch. 6, VII, p. 79

Capitalism: defined, Ch. 5, VIII, p. 67; what holds a trap back only delays it, Ch. 5, VIII, p. 67; ethical, structural impossibility of, Introduction, p. iv, Ch. 5, VIII, p. 67; bound by law, Ch. 5, VIII, p. 68; extractive, Ch. 8, III, p. 107

Capture, of governance: Ch. 1, VII, p. 11; Ch. 5, VIII, p. 68; Ch. 8, VII, p. 111; of a community, Ch. 8, VIII, p. 114

Care: defined, Introduction, p. iv; Ch. 5, III, p. 61

Care, transaction and power: at the scale of a whole group, Ch. 8, IV, p. 108; Ch. 9, IV, p. 141

Care, Transaction, and Power: Ch. 5, III, p. 62; Ch. 8, IV, p. 109; Ch. 8, VIII, p. 114

Carse, James P.: Ch. 3, I, p. 31

Ceiling Below the Floor, The: Ch. 5, V, p. 64

Choice, change, and causation: Ch. 8, V, p. 109

Civilizational decline: Ch. 7, VIII, p. 100

Civilizational lock-in: Ch. 1, VI, p. 9

Click rate: Ch. 2, VI, p. 23

Closure condition: Ch. 9, III, p. 140

Collaborative choice making: Ch. 4, III, Authorship, p. 48

Collective action, the logic of: Ch. 1, VIII, p. 12

Commons, governed by their own users: Ch. 1, VII, p. 11; Ch. 8, VII, p. 112; nested, Ch. 8, VII, p. 112

Community, market and governance: Ch. 8, IV, p. 108; their order, Ch. 8, IV, p. 109

Community: as defense, Ch. 6, X, p. 84; as minimum viable unit, Ch. 8, VI, p. 110; and the making of governance, Ch. 8, VII, p. 111; acting together, Ch. 8, VII, p. 112; seeing a message aimed at one person in private, Ch. 8, VII, p. 112; holding the ban before its neighbors do, Ch. 8, VII, p. 113; captured, Ch. 8, VIII, p. 114

Consciousness, and harm: Ch. 9, VI, p. 154

Consent, proportional to the risk: Ch. 9, I, p. 124

Control, limits of: Ch. 9, VI, p. 154

Cooperatives: Ch. 5, VIII, p. 68

Crowdfunded platforms: Ch. 6, V, p. 77

Cult structure: Ch. 5, VI, p. 65

Culture, destroyed by AI: Ch. 9, IV, p. 140

Deception, structural: Ch. 4, VI, p. 52

Demiurge: Ch. 9, II, p. 133

Design and Games: Ch. 3, VI, p. 36

Design consciousness: Ch. 3, X, p. 39

Design theory, and game theory: Ch. 3, V, p. 34

Diamond, Jared: Ch. 1, VI, p. 9

Discernment: Ch. 4, III, Authorship, p. 48

Distributed immune system: Ch. 3, VIII, p. 38

Doctorow, Cory: Ch. 9, VI, p. 152

Dollar auction: Ch. 1, VII, p. 10

eBay v. Newmark: Ch. 9, IV, p. 142

Economic decoupling: Ch. 4, III, Authorship, p. 48; and automation, Ch. 7, IV, Institutional Knowledge, p. 96

Embodied value: Ch. 9, II, p. 129

Endorsement problem: Ch. 8, III, p. 107

Engagements with Machine Madness: Ch. 5, IX, p. 69

Enshittification: Ch. 9, VI, p. 152

Entrapment of choice: Ch. 5, VII, p. 65

Ethical capitalism, defined: Ch. 5, VIII, p. 67

Ethics washing: Ch. 8, II, p. 106

Existential risk: Ch. 4, III, Extractive Sovereignty, p. 49; Ch. 9, I, p. 121

Exponential model: Ch. 2, IV, p. 19

Extraction, begins by moving value from its source: Ch. 4, III, Extractive Sovereignty, p. 49; cuts content from its context, Ch. 4, III, Extractive Sovereignty, p. 49

Extractive sovereignty: Ch. 4, III, Extractive Sovereignty, p. 49

Fiduciary responsibility: Ch. 9, V, p. 147

First-category existential risk: Ch. 6, VIII, p. 82

Fish farming parable: Ch. 1, II, p. 4

Fogbank: Ch. 7, IV, Institutional Knowledge, p. 93

Foundation of choice: Ch. 8, X, p. 117

Free speech, and the ban: Ch. 6, VIII, p. 81; a warning, and this book, under the ban, Ch. 6, VIII, p. 82

Fungibility: Ch. 8, IV, p. 109

Funnel record: Ch. 7, IV, Institutional Knowledge, p. 94

Game: definition, Ch. 3, II, p. 31; and design, Ch. 3, IV, p. 33

Game theory: Ch. 1, VIII, p. 12; Ch. 3, V, p. 34

Gatekeeper function: Ch. 4, III, Authorship, p. 46

Gini coefficient: Ch. 9, III, p. 137

Governance architecture: Ch. 1, I, p. 4; Ch. 1, VII, p. 11; its forms, Ch. 1, VII, p. 11; Ch. 8, VII, p. 111; the ban as its first piece, Ch. 8, VII, p. 112

Government, captured by corporate power: Ch. 9, IV, p. 143

Graph theory: Ch. 4, I, p. 43

Guild, as the solution to transmission: Ch. 7, II, p. 89

Hardin, Garrett: Ch. 1, I, p. 3

Heat tiles, of the space shuttle: Ch. 7, IV, Institutional Knowledge, p. 92

Herman, Edward S., and Noam Chomsky: Ch. 6, VIII, p. 81

Hybrid entities: Ch. 9, VI, p. 152

Identifying key: Ch. 9, I, p. 123

Immanent Metaphysics: Ch. 8, V, p. 109; the immanent, the omniscient and the transcendent, Ch. 9, I, p. 126

Impact investing: Ch. 8, II, p. 105

Information, and knowledge: Ch. 7, III, p. 90

Inner dimension of change: Ch. 8, VIII, p. 114

Institutional knowledge: Ch. 7, IV, Institutional Knowledge, p. 92; pruned branches, Ch. 7, IV, Institutional Knowledge, p. 92

Intergenerational knowledge transfer: Ch. 7, II, p. 87; kept alive at the Grand Shrine of Ise, Ch. 7, VI, p. 97

Intermediation: Ch. 4, III, Authorship, p. 46; as a tax on culture, Ch. 9, IV, p. 140

Ise, Grand Shrine of: Ch. 7, VI, p. 97

Islands, the two, and the spread of the market: Ch. 2, I, p. 16; Ch. 6, IV, p. 76

Keys, the billionaire's, and the handoff between generations: Ch. 7, IV, Institutional Knowledge, p. 96

Life raft analogy: Ch. 6, III, p. 75; Ch. 6, VII, p. 79

Linear technology: Ch. 5, V, p. 63

Litmus tests, of automation: Ch. 7, IV, Institutional Knowledge, p. 95

Losers, outnumber winners: Ch. 5, IV, p. 62

Maker and distributor: Ch. 4, II, p. 44; Ch. 4, III, p. 45; in content, Ch. 4, III, Authorship, p. 46; engineering against marketing, Ch. 4, III, Extractive Sovereignty, p. 49

Malthus, Thomas Robert: Ch. 1, III, p. 6

Malthusian rat trap: Ch. 1, III, p. 6

Manifesto, techno-optimist: Ch. 5, II, p. 59; its one use of care for people, Ch. 5, III, p. 60

Marketing: as paradigmatic game, Ch. 3, VII, p. 36; and the network, Ch. 4, IX, p. 54; the manifesto as, Ch. 5, IX, p. 69; its tactics, and the warning about AGI, Ch. 9, I, p. 121; over engineering, Ch. 9, I, p. 126

Marketplace takeover: Ch. 7, III, p. 89; what it cannot price, Ch. 7, III, p. 91

Markets, Money, and the Moral Economy: Ch. 2, IV, p. 20; Ch. 2, V, p. 21; Ch. 8, II, p. 106

Meat market: Ch. 2, VII, p. 24

Menu questions: Ch. 9, I, p. 124

Multi-polar trap: Ch. 1, I, p. 3; Ch. 2, IX, p. 26; Ch. 4, VI, p. 52; Ch. 5, IV, p. 63; Ch. 6, VII, p. 79; Ch. 7, VII, p. 99; Ch. 8, IV, p. 108

Multipolar Trap, The: Ch. 1, VII, p. 11

Nargame: Ch. 3, III, p. 32

Nar-social behavior: Ch. 3, VI, p. 35; a mode, not a kind of person, Ch. 3, VI, p. 36

Network position, and economic reward: Ch. 4, III, p. 46

Network topology of value extraction: Ch. 4, II, p. 43

Nietzsche, Friedrich: Ch. 5, VI, p. 64

Nine Points of Collective Insanity: Ch. 9, III, p. 140

Off switch: Ch. 9, VI, p. 155

Olson, Mancur: Ch. 1, VIII, p. 12

One master, of corporations: Ch. 9, IV, p. 142

Only life can care for life: Ch. 9, VI, p. 152

Orwell, George: Ch. 9, III, p. 138

Ostrom, Elinor: Ch. 1, VII, p. 11; Ch. 8, VII, p. 112; design principles of commons, Ch. 8, VII, p. 112
Overton Window: Ch. 6, VIII, p. 82
Panopticon: Ch. 9, III, p. 137
Perfected personal emotional extortion: Ch. 6, III, p. 76
Phenomenology: Ch. 4, VI, p. 52
Polanyi, Michael: Ch. 7, II, p. 88
Poverty, extreme, World Bank figures: Ch. 5, VIII, p. 67
Power: defined, Ch. 5, III, p. 61; and strength, Ch. 8, VIII, p. 114
Power and Its Prevention: Ch. 2, III, p. 18; Ch. 4, III, Extractive Sovereignty, p. 49
Power without care: Introduction, p. iv; Ch. 5, I, p. 59; Ch. 5, III, p. 60; Ch. 8, V, p. 109
Precautionary principle: Ch. 9, IV, p. 145
Presentation ratchet: Ch. 8, II, p. 106
Principal-agent problems: Ch. 6, VII, p. 80
Prisoner's dilemma: Ch. 1, I, p. 3
Product discovery systems, ethical: Ch. 6, VI, p. 78
Race to the bottom: Ch. 1, VII, p. 10
Reader, what to do now: Ch. 8, X, p. 116
Recycling: Ch. 5, V, p. 63
Redeeming Technology: Ch. 6, VI, p. 79
Reductionism: Ch. 9, II, p. 128
Rent: Ch. 2, IV, p. 19
Repair and reconstruction: Ch. 7, VI, p. 97
Restaurants: Ch. 4, II, p. 44
Retreat from explanation: Ch. 3, VI, p. 36
Revenue function, of marketing: Ch. 6, V, p. 77
Right action: Ch. 8, I, p. 105; Ch. 8, X, p. 116
Ring of Power: Ch. 9, V, p. 149
Rules for Rulers: Ch. 6, VII, p. 79
Scientific American: Ch. 9, VI, p. 150
Search, and the definition of advertising: Ch. 6, VIII, p. 81

Shubik, Martin: Ch. 1, VII, p. 10
Singleton: Ch. 9, II, p. 127
Skewing effect: Ch. 6, VI, p. 78
Skills, Inheritance, and the Civilizational Commons: Ch. 7, II, p. 89
Social enterprise: Ch. 8, II, p. 105
Social power, two senses of: Ch. 9, III, p. 136
Socializing, time spent: Ch. 2, VI, p. 23
Solar and wind energy: Ch. 5, V, p. 63
Source code, loss of: Ch. 7, IV, Institutional Knowledge, p. 93
Stern, Nicholas: Ch. 5, III, p. 60
Strength, and power: Ch. 8, VIII, p. 114
Structural, meaning of: Introduction, p. iv
Structure, Not Intent: Ch. 2, I, p. 16; Ch. 6, VI, p. 78
Super-PAC: Ch. 9, V, p. 148
Superintelligent AI: Ch. 9, VI, p. 150
Supernormal stimulus: Ch. 2, VI, p. 23
Surveillance capitalism: Ch. 6, III, p. 76; Ch. 9, IV, p. 144
Swift, Jonathan: Ch. 9, VI, p. 154
Tacit knowledge: Ch. 7, II, p. 88
Thirty years, the prediction of decline: Ch. 7, VIII, p. 100
Three-generation pattern: Ch. 7, II, p. 88; and the billionaire's keys, Ch. 7, IV, Institutional Knowledge, p. 96
Time Machine, The: Ch. 7, IV, Institutional Knowledge, p. 93
Tinbergen, Niko: Ch. 2, VI, p. 23
Tivy, Wolf: Ch. 9, II, p. 127
Tolkien, J. R. R.: Ch. 9, V, p. 149
Total ban: Ch. 6, VIII, p. 80; held by communities, Ch. 8, VII, p. 111
Tragedy of the commons: Ch. 1, I, p. 3
Trans-humanism: Ch. 9, I, p. 122
Trust hacking: Ch. 9, I, p. 123
Two masters: Ch. 8, II, p. 106
Two-income trap: Ch. 1, V, p. 8

Valencia, Water Tribunal of: Ch. 8, VII, p. 112
Value added, in the food dollar: Ch. 4, II, p. 44
Value of content: Ch. 4, III, Authorship, p. 46
von Neumann, John, and Oskar Morgenstern: Ch. 3, V, p. 34
Warren, Elizabeth, and Amelia Warren Tyagi: Ch. 1, V, p. 8
Why AGI Cannot Be Made Safe: Ch. 9, I, p. 121; Ch. 9, III, p. 140; Ch. 9, VI, p. 154
Why No Control Suffices: Ch. 9, VI, p. 154
Wittgenstein, Ludwig: Ch. 3, I, p. 31
Yudkowsky, Eliezer: Ch. 9, II, p. 127
Zuboff, Shoshana: Ch. 6, III, p. 76; Ch. 9, IV, p. 144