A BOOK ABOUT THE BETTER NEXT

The Mathematics of Civilization

An Explainer of Nature's Human Value Structure

Why humans waking up wanting better, shared memory, agency, civility, and correction will raise both the floor and the ceiling

RICHARD POLK

Built from the August 2026 manuscript foundation edited by Arbad Tahir

Reading edition · 23 October 2026

The Mathematics of Civilization
© 2026 Richard Polk. All rights reserved.
Public Review Edition — Version 1, Oct. 23, 2026

This manuscript is shared for personal reading, discussion, criticism, and suggested corrections. Readers may download and print a copy for personal use and share links to the manuscript at AskRichard.xyz.

Except as permitted by copyright law, including fair use, or by written permission, no other reproduction, redistribution, sale, translation, or adaptation is authorized. Brief quotations with attribution and a link to AskRichard.xyz are welcome for reviews and discussion.

Questions, factual corrections, contrary evidence, and proposed improvements are invited through AskRichard.xyz. This edition is subject to revision.

NOTE

What ‘Mathematics’ Means Here

A claim about direction, not a claim of infallibility

This book uses the word mathematics in its oldest practical sense: a disciplined search for patterns that remain coherent when we look backward, describe the present, and reason forward. It does not claim to have discovered a new law of physics, whatever the title suggests.

It does not say that history proceeds in a straight line, that every invention is good, or that humanity is guaranteed a painless future. Civilizations collapse. Wars reverse progress. Epidemics expose fragility. Technologies enlarge cruelty as readily as kindness when people design them that way. I’ve lived through enough of the last century not to need convincing.

The narrower and more defensible claim is this: across long periods, systems that help more people survive, learn, cooperate, create, exchange, and correct error tend to generate more durable value than systems that depend on keeping most people weak, ignorant, frightened, or excluded. That recurring selection pressure is what I call the mathematics of better.

The word better, left alone, will do whatever an author wants it to. Better does not mean perfect, uniform, painless, or morally pure. Better means that, of the choices on the table, one leaves people more able to run their own lives, for longer, and leaves more room for the next improvement.

A clean water system is better than contaminated water. Why, though? Not because clean water is more admirable. Because the people drinking it go to work, raise children who live long enough to argue with them, and can spend the day on something besides

staying alive. You don't measure the water. You measure what people do with the time it gives them, and the same test runs on a school, a court, a power grid, a market, or a piece of software.

Better is comparative, testable, revisable, and unfinished. Where a claim fails, it's the claim that's wrong, not the definition.

One more word, because Part IV and the last page of this book are named for it. NEXT is not a synonym for better. NEXT is just what comes next: whatever inheritance, choice, constraint, technology, nature, and time produce, whether anyone wanted it or not. A better next is a judgment made after the view has been widened. Civilization doesn't have to be good to continue, which is why the difference matters. To continue it has to keep enough memory, enough agency, and enough room to correct itself that another arrangement is still possible.

Freedom creates the vectors. Shared visibility reveals the resultant direction.

AUTHOR'S NOTE

The Provenance of Visibility

How this book was made

People ask whether the machines wrote this book. Parts of it, and I'll tell you which parts. They didn't decide anything in it.

For most of my life a question that ran across economics and history and engineering and medicine and energy meant months in a library and a dozen specialists who were never going to be in the same room. This time they were, more or less, and the room was open whenever I was. I'm in bed by nine, and the worrying starts at half past four in the morning. I used ChatGPT, Copilot and Claude. I'd ask, they'd answer, I'd ask why they were so sure, and often the answer changed.

An example: an early draft said 826 million people live in extreme poverty. It was wrong in an instructive way. Different institutions were publishing different figures from different poverty lines and data vintages, and my manuscript had quietly turned several defensible numbers into one indefensible one. The correction was not that the machines knew nothing. It was that confidence without provenance is not enough.

There were other corrections. Dwight D. Eisenhower's farewell speech runs about sixteen minutes, not the shorter figure often repeated. A crowd in Leipzig I'd called a doubling went from about 70,000 to about 120,000, which isn't doubling, and this is a book with mathematics in the title! A Clean Air Act figure I had used turned out to be the midpoint of a very large range.

I've sold things and been sold things all my life, and I know what confidence without a receipt looks like, from both sides of the counter. Claims moved when somebody put them against real

documents.

Fear of AI is a mistake. The ultimate AI outcome will not be AI killing us. What it will kill is our lies, starting with the ones we tell ourselves about how sure we are. I don't say that because the danger is imaginary. A machine that can show me where my beliefs do not match the evidence can also be used to move my beliefs to suit somebody else. Same capability, pointed the other way. Which way it points depends on how the thing is built, what evidence it is allowed to see, and whether a person is still the one closing the loop. The mistake is treating the danger as destiny. AI may make some lies harder to maintain. It can also manufacture them cheaply.

Arbad Tahir is the other reason the book is what it is. He's a young engineer in Islamabad, Pakistan. I'm a retailer in Boulder, Colorado, born in 1949, so do the arithmetic. I'm Jewish and he's Muslim, and neither of us found that worth mentioning until other people did.

Our collaboration began with a cold call from Islamabad to Boulder. He was looking for work, and I liked his tone. At the time I had a stalled project called BoulderBubble.com, meant to show the people of Boulder, and anyone else who cared to look, more of what was interesting and useful about the place. So instead of ending the call, I offered him the chance to work on it. That small decision became a four-person Boulder-Islamabad lab, though none of us would have called it that at the start.

The first experiment that mattered was OptimismAI.app. We gathered the writings of ten of history's great thinkers, Aristotle, Immanuel Kant and Albert Einstein among them, and asked what each could add to optimism as a way of looking at a question. We weren't trying to make them agree, and we certainly weren't trying to manufacture a cheerful answer. We wanted to know whether very different minds, set side by side, could improve a question before anyone answered it, and then, once there was an answer, whether the system could ask what was missing, what contradicted it, and what

ought to be asked next. We called that persistent question refinement. The answer stopped being the end of the job and became the material for the next, better question, and that habit has stayed underneath almost everything we've done since. I came to put it this way: questions are the compression algorithm of intelligence.

Then came AskRichard.xyz. We gathered roughly 40,000 documents of my comments. The obvious use was to look things up: what did Richard say about this? The more interesting possibility was underneath. Across thousands of comments on unrelated subjects, could a machine find what I keep coming back to, what I weigh against what, where I contradict myself, what I mean by better? And could that, more than the words themselves, help answer a question I'd never been asked? I put the ambition less academically: could it answer as though it were me with a thousand PhDs? Not because the machine becomes me. It doesn't. The hope was a machine that knows far more than I do, reasoning through a picture of my judgment that I can see and correct. That's where the idea that became Value-Structured Intelligence started to take shape.

AskTheFounders.us made the experiment harder and better. Now there was no single writer. Jefferson, Hamilton, Madison, Adams, Franklin and the rest disagreed with each other, contradicted themselves, and sometimes failed spectacularly to live up to their own principles. A set of values could no longer be a personality profile or a book of favorite quotations. It had to show competing principles, changing circumstances, and what happened when values collided. We didn't set out knowing VSI was where we'd end up. One question produced the next.

Then the public-information problem became impossible to ignore. A machine can say something persuasively without showing why anyone should believe it. What took me several models, several sources and one stubborn engineer to do shouldn't be a hobby. It should sit next to what you're reading and show you where a claim came from, what disagrees with it, how sure anyone has a right to be,

and what to ask next, and then get out of your way, because you're the one who has to live with the decision. We are trying to build that. It's called IsItTrue.xyz. Of everything we've tried, I think making checking ordinary may matter most to the public. Making and spreading a claim has already become pedestrian. Checking one hasn't.

The experiments are unfinished, and they don't prove anything this book argues. They're what happened when four people in Boulder and Islamabad took pieces of the argument seriously enough to try them.

The method became part of the evidence. That's why it gets these pages.

PREFACE

The View from 1949

A lifetime inside the acceleration

Being born in 1949 has allowed me to watch an unusual amount of the future arrive.

There was no television in the house when I was born. By the time I was twenty, a camera was sending pictures back from the surface of the Moon. Each new thing was astonishing for a while. That is the part nobody warns you about, and the reason so few people can tell you what year their life changed. The future does not announce itself. It shows up, gets argued over, gets cheaper, and then becomes the thing you have stopped noticing.

As a child I belonged to a retailing family that sold the appliances of liberation: televisions, refrigerators, washing machines, dishwashers, and other machines that quietly returned hours of life to ordinary families. We sold the television sets on which people later watched human beings walk on the Moon, and when Neil Armstrong actually stepped onto it, the company gave away 40,000 maps of the place so that a family in Chicago could go outside afterward and look for the spot. I grew up among stories of oversized watermelons tied to automobile roofs, crowded stores, imperfect advertising, and the wonderfully human effort to persuade somebody that tomorrow could be easier than yesterday. Writer Ann Paden preserved that company's history in *I Bought It at Polk Bros.: The Story of an American Retailing Phenomenon*.

What that world taught me, it taught me slowly and without my noticing. A washing machine is not a convenience. It converts a Monday spent at a washtub into a Monday spent somewhere else, a redistribution of a person's life disguised as an appliance. The

advertising said as much, and the fact that it was advertising doesn't make it false. It makes it a claim somebody should check. Most of this book is me, checking.

THE THING NOBODY NOTICED

In April 1956, a trucking man named Malcolm McLean loaded fifty-eight metal boxes onto a converted tanker in Newark, N.J., and sent it to Houston.[1] There was no ceremony. For as long as there had been ships, longshoremen had carried cargo aboard piece by piece, sacks and barrels and crates, and the cost of getting something from a factory into a ship's hold helped decide where it could economically go. It was just what shipping cost. McLean's boxes made loading a matter of cranes instead of men's backs, and within a generation they had helped reorganize factories, ports, cities, and eventually the world economy.

Nobody watching that ship leave Newark thought they were watching anything historic, or the start of a system that would one day put a telephone together from parts made on several continents. That is how NEXT usually arrives. You see the box. You don't see the world the box makes possible.

I'd guess that a large part of what our stores sold eventually traveled in one of those boxes. I never thought about the boxes. I thought about the televisions.

MY LIFETIME IS A DATA SET

I was born November 8, 1949. If you want to know whether human conditions can change, one human lifetime turns out to be a fairly useful place to look.

The world I entered held about 2.5 billion people. Today it holds more than 8 billion. That fact alone could have supported the darkest predictions about population: more than three times as many people

needing food, shelter, water, energy, medicine, schooling, and everything else a human life consumes. So look at what happened at the same time.

Around 1950, worldwide life expectancy at birth was about 46 years. The latest United Nations figure is 73.3 years.[2]

In 1950, roughly 23 percent of children born worldwide died before reaching age five. By 2024 the global rate was about 3.7 percent.[3] That is the statistic I have the hardest time treating as a statistic. It means a mother is holding a child who, on the odds of the year I was born, had a much greater chance of not being there. Multiply that by the Earth. Vaccination did some of it, and antibiotics, clean water, sanitation, oral rehydration salts, nutrition, electricity, refrigeration, schools, roads, hospitals and better information did the rest, along with a great many people keeping records. None of it required humans to become a different species. It required us to find something that worked, remember it, correct it, and make more of it available.

Literacy tells the same story. In the years just after the Second World War, UNESCO reported that more than 44 percent of adults worldwide could not read or write. By 2024, global adult literacy was about 88 percent.[4] A person who can read can check a contract before signing it, read the dosage on a medicine bottle, follow the instructions for fixing a pump, compare one political claim with another, study something nobody in the family has studied before, and answer a lie in print. Reading moves you from the people things are decided for toward the people who can at least go and find out.

Extreme poverty is harder to compare across seventy-six years, because definitions, surveys, price adjustments and coverage all change. The direction isn't subtle. The World Bank's current historical reconstruction estimates that around 60 percent of humanity lived in

extreme poverty in 1950. Today the share is roughly one person in ten.[5] That is still hundreds of millions of people, and still an unacceptable floor. It is not a law of nature.

THE NUMBERS DO NOT SAY WHAT PEOPLE THINK THEY SAY

They don't prove that humanity is good, or that history bends toward anything on its own, or that the next seventy-six years will be better. Wars reverse progress. Governments destroy things. Technologies can enlarge cruelty as efficiently as kindness. Markets can widen opportunity and hide costs somebody else is paying. Human beings can hold more information than any generation in history and still believe something ridiculous before breakfast. I have done it myself.

The numbers prove something narrower.

Conditions we once treated as permanent can change.

Children dying was normal. Illiteracy was normal. Extreme poverty was normal. A life that averaged forty-some years was normal. They were conditions, and conditions can change. That is a smaller claim than optimism, and a considerably more useful one.

THE PAST WAS NOT WRONG

We should therefore look at history without romance and without contempt. The past was not wrong. It was the path that made visible a better next. Fire, walls, kings, empires, factories, fossil fuels, national armies, broadcast media, and centralized expertise solved problems earlier structures could not solve, and created new concentrations of power and new harms in the solving. The obligation of every generation is not to hate the staircase beneath its feet. It is to build the next step.

WHERE THE ACCELERATION WENT

Today the acceleration is no longer merely mechanical. Electrons are arranged as energy, knowledge, memory, and communication.

Point a smartphone at a tree and a child can learn what kind it is, where it came from, and what eats it. Farmers can see weather before it arrives. At two in the morning, which is when people need it, a patient can ask a machine to explain what a specialist said that afternoon. None of this replaces the forest, the teacher, the farmer, the doctor, or the patient's own judgment. Built badly, the same tools can manipulate, spy, exclude, frighten, and manufacture lies at a scale no advertising department I ever knew could have imagined. Built differently, they give a person more say.

Inside one lifetime we went from selling machines that gave people back an hour to building machines that can show a person, in seconds, connections across thousands of years of human memory. That doesn't make the machine wise. It changes what a human being can see.

As a boy in Evanston, Illinois, I once asked the school librarian for a book about how something really worked. She gave me one about a robot named Timmy, who was friendly and didn't work, so I asked again. This time she gave me a book on photovoltaics, which in those years meant a switch that opened a door when you broke a beam of light. That book led me, by way of a science fair, to a peanut farmer in Americus, Ga., where I was told the telephone company had put some new solar cells on a line to see if they would survive the weather. I typed him letters on a machine whose keys struck unevenly, making each letter a different shade. He wrote back. I asked him why the solar cells mattered. He didn't talk about the sun. He talked about his family: when the power went out, he could still call his relatives.

I have remembered that answer for most of my life. The technology was interesting. What it gave somebody back was the point.

THE FUTURE WAS UPSTAIRS AND I MISSED IT

For some of the important years in the development of GPU computing, NVIDIA leased the upper floor of my building on Pearl Street in Boulder. I sold shoes downstairs. The people upstairs kept telling me that what they were building was becoming more than a gaming chip. I liked them. You could feel the energy in that lab. I did not buy the stock. I did not buy Google either. I have a remarkable ability to stand close to the future without realizing what it's worth.

I've said since, only half joking, that artificial intelligence was being invented upstairs from my shoe store while I sold shoes underneath it and kept the roof in good repair. The joke is better than the historical claim. The useful part is that something important was happening almost directly over my head and I didn't understand how important it was. Nobody does, at the time. The future is much easier to recognize after it has become ordinary.

What I can do instead is look for what has worked before, and it keeps coming back to the same short list.

More memory. More visibility. More agency. More correction.

Not always, not everywhere, not automatically. Often enough to be worth investigating.

I am grateful for the tension between pride and gratitude. Pride says human beings built all this. Gratitude remembers that none of us built it alone. Every tool contains the labor of ancestors, strangers, teachers, rivals, and children not yet born when the work began. The world I entered in 1949 was not an inferior version of ours. It was our inheritance at that moment. People used what they had, learned, made mistakes, corrected some of them, and passed the result forward. Our job is the same.

I don't know what NEXT will be. Nobody does. But after one lifetime in which the human population more than tripled while lives lengthened, children survived, literacy spread, and extreme poverty fell, I think we've earned the right to ask a better question than whether everything is falling apart.

What have we learned that makes a Better NEXT possible?

And then:

What should we do with it?

THE EQUATION

Inherited capacity x individual choice x shared correction = civilization's next step.

CONTENTS

Note: What ‘Mathematics’ Means Here

Author’s Note: The Provenance of Visibility

Preface: The View from 1949

PART I - THE EVIDENCE BEHIND US

1. The Human Constant
2. The Persistent Pursuit of a Better Next
3. Scarcity Was a Stage, Not a Destiny
4. Eight Billion Mornings
5. Good Is More Profitable Than Bad

PART II - THE STRUCTURE OF NOW

6. Agency Is the Engine
7. Electrons Arranged as Energy, Knowledge, and Memory
8. Truth Is Infrastructure
9. The Old Mathematics of Power

PART III - THE COLLECTIVE AI HORIZON

10. AI Turns the Lights On
11. History Written by Vectors
12. Value-Structured Intelligence
13. Alignment Without Uniformity
14. Two AI-Literate Generations

PART IV - THE FLOOR AND NEXT

15. The Floor Is the Measure
16. Something to Lose
17. Nature Put Everything in the Box; Man Found Good and Bad
18. The Public-Health Economics of Peace
19. When Capability Outruns Obligation
20. The Velocity of the Ordinary
21. The Arithmetic of MORE

Epilogue: Take a Deep Breath and Smile

Coda: NEXT

Appendix: The Core Equations

Glossary

Notes and Sources

Acknowledgments: The Vectors That Made This Possible

PART I

The Evidence Behind Us

The future is credible only when it explains the path already traveled.

When considering a pattern, the first question is whether the pattern is real.

CHAPTER 1

The Human Constant

Eight billion mornings and the equation we have been living

Every morning, more than eight billion people wake up wanting something to be better than it was the day before.

Most of the wanting is small. Some of it is enormous, and you can't always tell which from the outside. A parent wants a child's fever to break. A shopkeeper somewhere is counting on one more customer before noon. A scientist is waiting to find out whether a result will happen twice. A refugee is watching a border. Each one is the biggest thing in the life it belongs to. Wanting a little more meaning when you already have everything and wanting clean water before the day gets hot are the same motion at different altitudes. The wants fight each other in every detail. They all point the same way.

That's the smallest thing I can ask you to agree to. Not God, not government, not history, not virtue. Just that people are wanting.

Wanting doesn't make us wise or kind. Sometimes we pick a local better that does bigger harm. Safety by pushing the neighbors around. Convenience by fouling the river. Belonging by keeping somebody out. Profit by lying. Whoever is choosing those isn't bad at arithmetic. They're doing the sum correctly, but their field of view is too small. I'd rather know that than have a moral verdict, because you can't redesign a person's character, and you very often can change what they're able to see. Call the small-field version the old mathematics: take what you can reach, guard what you took, and send the bill to whoever can't send one back. It has a long record, and it isn't finished.

Before there was an economy, a constitution, a creed, or a nation, there was a human being waking up into a condition and wanting some part of it to be better.

That is older than civilization and more universal than anything civilization has come up with since. Somebody beside a fire wanted the warmth to last till morning. Somebody watching the sky in planting week wanted the rain to come before the seed was wasted. What each of them wanted belonged to one life and one season, and the particulars go on forever. The motion underneath doesn't change.

Civilization is what happens when memory, evidence, courts, and a wider ledger slowly improve the field the wanting works in.

The common denominator of civilization is not agreement. It is people waking up wanting better.

THE PAST IS THE BRIDGE

Start there and you have to treat the past differently. Religions, families, courts, markets, schools, sciences: they're parts of the bridge out of the cave. Each carried something across a generation that would otherwise have had to start over: a memory, a restraint, an obligation, a way of belonging. I want to say be grateful for all of it, and I nearly do.

Then I hear the comeback, and it's a good one. Some of those bridges were built to keep certain people off them. Some defended situations we'd now reject, some kept out people who should have counted, and some protected authority while claiming to protect truth. A comfortable man in Boulder saying be grateful for the bridge sounds a lot like a man who was allowed to cross it.

I'm not going to argue my way out of it, because I can't. People were responsible for every one of those things, so each has people's limits in it, and some of the people limited were kept out on purpose.

Gratitude was never the hard part. What you do with it is. Ask what the situation was, what it made possible, what it cost once the bill came in, and what's worth carrying into the next situation, and the two sides stop fighting. The people who revered the U.S. Constitution amended it twenty-seven times, ten of them at once and almost immediately, because a good many states wouldn't have ratified the document otherwise. Reverence and correction in the same two years, by the same hands. You can admire a market for getting strangers to cooperate, and still fix it when its prices leave out the people who never entered the deal.

I end up here. Announcing that people in the past were wrong and we've finally got it right would be absurd; we're standing on arrangements they paid to discover. Announcing that the situations were fine would be a different absurdity, and it's the one my side of the argument is prone to. Part of what makes an inheritance good is that it leaves the heirs enough stability, knowledge, and freedom to improve it. Part of what makes an heir worth the inheritance is that he does.

We honor the bridge by crossing it, not by insisting humanity must live forever in the middle of it.

Civility can be reduced to manners, which is how it got its reputation for being optional. I mean something harder by it: staying in the room long enough for a disagreement to tell you something instead of costing you something. Under civil rules, the side that lost can come back tomorrow, a minority stays visible, an error gets corrected without burning down the building that recorded it, and this year's rival is next year's customer.

You can see how far we've come from the cave. People still fight, threaten, conquer, lie, and kill. They also sign treaties, let inspectors in, take disputes to arbitration, print dissent, hand over power after

losing, keep archives, trade across old battle lines, vaccinate strangers, and ask out loud whether an institution built for an old danger is still worth what it costs.

Hope, looked at this way, isn't sentiment added to the mathematics. Hope is the recurring human input. Memory is the storage system. Diversity supplies experiments. Consequences supply evidence. Civility keeps correction possible. Agency allows people to act on what has been learned. What comes out is next.

ANY CLAIM ABOUT THE FUTURE HAS TO EXPLAIN THE PAST

Civilization is the long process of enlarging that field of view, and it got there by specific inventions, not by anyone getting better. Take the first one that mattered.

Somewhere in what's now southern Iraq, a little over five thousand years ago, a man whose job was keeping track of a temple's grain got tired of remembering. He pressed marks into a lump of wet clay, so many jars, so many sheep, owed by whom, and let it dry. It was a receipt. But the receipt could do something nobody had ever owned before. It said the same thing after he was dead.

Writing allowed the dead to advise the living. Before the clay, everything a person learned in a lifetime had to be said out loud to someone standing nearby, or it was gone. A farmer who worked out the right week to plant carried that knowledge for many productive years, but then took it into his grave, minus whatever a son happened to absorb from him. Every generation started over again, somewhere near where the last one began. The writing on the clay broke that, not by making anyone smarter, but by letting the data outlive the mouth it came from. A clerk could be wrong. So could a Roman engineer or a medieval physician, in ways their readers could see later, and being visibly wrong across centuries is how error stops being permanent. The receipt was the beginning of being able to check.

Other advances were versions of the same job. Measurement let a claim be checked against the world instead of against a reputation. Law made power partially predictable, which is the minimum condition for planning anything. Markets let strangers coordinate without trusting each other, science made organized doubt productive rather than merely irritating, and democracy widened the set of people permitted to participate in correction. Digital networks then connected billions of minds at a cost close to zero, and artificial intelligence has now begun to increase the speed at which evidence, memory, contradiction, and perspective can be cross-correlated.

We're in the midst of that last one. It belongs on the list with the others, and the list isn't finished.

Look for the pattern in outcomes. Intentions won't show it.

More children survive. More adults can read. More people are connected to one another than at any previous moment. Renewable electricity has become dramatically cheaper within a single decade. Extreme poverty, still vast and recently stubborn, has fallen a long way from its 1990 level.[6]

None of it happened everywhere or for good, and a war, a government, or a bad decade can undo any of it. What matters is that improvement keeps becoming possible when more people know more and can act on it, and keeps stalling when they can't. That's a pattern, not a promise.

Each of those numbers is a sum. Informing people allows them to know more, act on it, and correct it when it goes wrong.

One rule governs everything from here on. It's easy to abandon without noticing, so I'd rather put it somewhere it can be held against me, and then hold it against me.

I sold shoes for years underneath the floor where NVIDIA was building the chips that artificial intelligence would come to run on, and I didn't see it. I've made my peace with that, and the peace is

honest. What was in line with my values was keeping the place those people wanted to be, and the world changed, and it didn't need me to see it coming. The future runs on the mechanism of the past, and I had the mechanism in my building and didn't recognize it. Nobody does, at the time. So my judgment about what comes next is worth as much as the mechanism I can show you. No more.

Any claim about the future has to run on the same mechanism that explains the past. If I tell you these tools will make people more capable, I owe you a time when more memory and more comparison already did that, somewhere we can go and look. Otherwise I'm not reasoning. I'm prophesying, and prophecy has an unbroken record of being wrong in the particulars while sounding right at the time.

That rule throws out most of what is said about artificial intelligence, from both sides. Nothing in the record supports a machine that ends scarcity, disease, and war, because no tool before did anything of the kind. The machine that makes people irrelevant fails the same test from the other end. Time after time, the technologies that took work away made work nobody had imagined, doing real and lasting harm to particular people in particular towns who deserved better. Both forecasts are exciting. Neither has to explain the last 4,000 years, which is exactly why they can be stated so confidently.

The test is crude and it does most of the work. If an argument about what comes next would also have to be true of writing, of printing, of electricity, and of the shipping container, it's describing a mechanism. If it only works for this technology, this year, it's describing a mood.

WE MISTAKE THE HEADING FOR THE TRACK

Every student pilot gets this lesson once, and most of them get it the hard way.

She takes off on a clear morning for a town a hundred miles due east, points the nose at it, and settles in. The air is smooth. Nothing on the panel says anything is wrong. Forty minutes out, the town isn't in view where it should be. It's off to the right, and the longer she looks the further right it drifts. She has flown exactly the heading she meant to fly the whole time. What she hadn't done was measure the wind, coming out of the north at twenty knots since she left the ground and nudging her south with every mile, so gently that nothing in the cockpit noticed. She was aimed at the destination but she was going somewhere else.

There is a fix. It is old and it looks wrong from the ground. She turns the nose into the wind, north of where she wants to go, by about eleven degrees for an airplane doing a hundred knots in that breeze, and holds it there. Now the airplane is pointed at a field she has no interest in but traveling toward the town she wants. Anyone watching from below sees a machine flying slightly sideways to where it's going. The pilot isn't confused. She has measured a force she can't see and included it. Pilots keep two words for this, heading and track, and a first month of training is mostly spent learning that they aren't the same thing.

The mathematics isn't complicated. Wherever there are people, forces are pushing in directions that don't agree. None of them goes away because whoever is at the controls finds it inconvenient. The forces add up. They add up to where the country is actually going, whatever its official heading says.

We already have such a resultant, eight billion people wide. We see it about as well as the student saw the wind, because our information is scattered, bent, late, or owned by somebody, and so we make the mistake the pilot trained out of in her first month. We confuse the heading with the track.

That one error produces most of the others. We mistake the loudest vector for the strongest one. We mistake a ruler for a population, a platform for a public, a trend for a truth, or a weapon for security. Something real was pushing sideways and nobody measured it, and when we land in the wrong place we call it unforeseeable instead of mathematics that was there all along.

Done for eight billion people at once, the pilot's fix takes four steps, short enough to carry in your head.

Give people more real choices they can act on and you get more real preferences. A preference nobody can act on is barely an idea. Enough real choices and you have vectors: eight billion directions of need, want, fear, and refusal, each one carrying a magnitude.

A vector is a sequence or list of numbers (coordinates) that capture the underlying properties of an item.

With cross-correlation, those vectors become visible, cross languages and distances, and show where they collide and where they quietly agree.

What comes out is not one mind. It is a direction. And a direction is the one thing humanity has never been able to see while it was restricted by vocabulary.

Not one machine handing down one answer. Billions of people with the same memory to draw on, every one of them free to choose, refuse, and disagree. A public library, not a state broadcaster. Both hand out information at scale. Only one of them leaves you free to go and check.

Measure better and some of the confusion improves. Very often, people who disagree about religion, government, history, and borders still want their children safe, their water clean, their work respected, and their future open. That doesn't erase the disagreement, and pretending it does is how peace talks die.

The two of us who made this book share no country, religion, age, or professional habit, which tests the same claim with two diverse places to fail or succeed. If an idea only made sense in Boulder but not Islamabad, it was too small. If a structure only worked once you hid the values inside it, it was too small. Agreeing was never the useful part. What was useful was keeping the evidence and the disagreement in view at the same time, long enough to reach the next question. I've come to look forward to being told I'm wrong by someone fifty years younger, which is not a sentence I ever expected to write.

Two people are a small model of eight billion. The difference between them was the instrument, not something to clear away before the work could start.

Alignment without uniformity is not only the destination.
It is how the argument was assembled.

The optimism on offer here is not the kind that believes everything will work out. It is a decision to examine whether better systems outperform worse ones over time, and to accept the answer.

The propositions are ordinary enough to check against what happens in your own life. Trust makes every deal cheaper. Clean energy turns a fuel bill you pay forever into a thing you own, which is why a diesel generator and a solar array of the same output stop being the same object on a balance sheet, even in a year when they cost the same to run. Peace opens the markets that war closes. Agency converts recipients into participants.

I've checked each of those against enough years to believe them, and each can fail, and each does.

Trust without checking invites fraud. Schooling without freedom teaches obedience, efficiently. Cheap energy can power an enormous amount of waste, and AI spreads propaganda faster than it spreads

anything else. So the mathematics must include the feedback. Better is what is still better after the wider costs, the longer time, and the people who were left out have been added to the account.

The next is not whatever happens next. The better next is what survives wider measurement.

Wider doesn't mean more people. Ask a bigger poll and you can get a worse answer, and a country can be unanimous because everyone was reading the same screen. What widens a measurement is how independent the inputs are and how honestly they're weighed. Majorities are easy to assemble and worth a good deal less.

THE EQUATION

Better = a choice whose benefits remain durable as more people, more consequences, and more time are included in the calculation.

CHAPTER 2

The Persistent Pursuit of a Better Next

The human species is defined less by what it has than by what it keeps trying to improve

"The natural flights of the human mind are not from pleasure to pleasure, but from hope to hope."

Samuel Johnson, The Rambler, No. 2 (1750)

Before civilization had a philosophy of progress, it had cold children, spoiled food, infected wounds, dangerous darkness, and the persistent question: Can this be better?

A cave fire was not a declaration of human supremacy. It was a local answer to darkness, cold, predators, and raw food. A roof was an answer to rain. A pot was an answer to storage. A path became a road because repetition exposed inefficiency. A mark on a wall became a symbol because memory trapped inside one brain died too easily.

Let's stay with the cold children for a while. Abstractions are easy. Objects aren't.

A NEEDLE IN THE EYE

Somewhere upward of 40,000 years ago, somebody took a splinter of bone, ground it thin, and drilled a hole through one end.[7]

Think about what that took. The hole is the hard part. Bone splits along its grain, and a hole near the end of a thin shaft is an invitation for the whole thing to fail; you have to bore it slowly, from both sides, with a stone point finer than the shaft you are ruining, and most attempts end in a broken needle and a wasted afternoon. Nobody does that for entertainment. Somebody did it because a hole

in a needle lets you pull a thread through hide, and pulling a thread through hide lets you make a seam, and a seam lets you make clothing that fits.

Draped hide keeps a person warm in the way a blanket does, which is to say imperfectly and only while they are sitting still. Fitted clothing is a different category of object. Sleeves that close at the wrist, a hood that closes around the face, layers that trap air against the skin and stay put while the body moves — that is not warmer clothing, it is a portable climate. And a portable climate is the difference between a species that lives where the weather permits and a species that decides where it is going to live.

Follow it forward. With tailored clothing, human beings occupied latitudes that had been closed to them: the cold steppe, the edge of the ice, eventually the far northeast of Asia. Populations that could survive a Siberian winter were positioned at the top of the land bridge that once joined Asia to Alaska. Everything that followed on two continents — every culture, every city, every argument about who was there first – arose downstream from the capacity to walk into cold weather and not die of it.

I've compressed that chain, and it's longer than a needle deserves. Dates are contested, the mechanism was a package of inventions rather than one, and nobody sewing a sleeve was thinking about Alaska. Nobody involved was pursuing a destiny. Somebody was cold, somebody tried something, and the something worked well enough to be worth the broken needles.

Then there's what the needle is made from: perhaps a bone from the animal that had to be killed to make the coat, using a stone tool made by someone else, according to a technique that had to be taught, by a person who learned it from someone perhaps now dead. Call the needle an invention and you have named the smallest part of it. What it is, is the visible end of an accumulation.

Nobody involved had a word for any of this. That cuts against the way we tell our own history. The person drilling that hole had no concept that they were making progress. Nothing in their world suggested that conditions were supposed to improve their life, that their grandchildren would live differently because of it, or that they were participating in anything at all. The idea of history having a direction is a recent one, a few centuries old at most, and for most of the human past it would have sounded like nonsense. Life was what it was. The seasons came around.

What they had was Tuesday. It was cold on Tuesday, the hide draped over someone kept slipping, and somebody thought there might be a way to fasten it and keep it on. That's the entire mechanism. To pursue it, nobody had to believe in a better next, and almost nobody who drove it did. All it takes is one person who prefers one condition to another and is willing to spend an afternoon finding out whether their preference can be satisfied.

I take some comfort that nothing here depends on people being visionary. Everything depends on their being irritated by things that could be better, which is the one human trait with an unbroken record.

BUT PEOPLE GET USED TO THINGS

People do get used to things. That can cause problems. Few findings in psychology are more reliable, and few facts of ordinary life more obvious. Nervous systems do it at the level of neurons — a smell vanishes in minutes, a constant hum stops being heard — and preferences do it at the level of years. A new house is astonishing for a year, and then it is simply where you live. Big things amaze for longer: the washing machine that gave back a day a week to a generation of women stayed a wonder for a good while. What wore it down was everything that came after it: the dryer, the dishwasher, the television, then the next thing and the one after that. There are so many of them now that each gets less of our attention, and it seems

to happen even quicker than it does. In the end nobody stood in front of the washing machine feeling grateful. They stood in front of it wondering why it took so long.

Read one way, that's fatal. If improvement stops registering with us, the accumulated record buys nothing, and a civilization can be carried an enormous distance only to arrive feeling precisely as aggrieved as it set out.

Read the other way, improvement is the engine. Habituation is what guarantees that the irritation regenerates. A species that stayed satisfied would have stopped at discovering fire. A next comes because the last one stopped mattering, sometimes inside a year and sometimes over a couple of generations. Grandpa believed one thing; his grandson believes another.

Comfort is another problem, a subtler one. Comfort doesn't claim the wanting stops. Its claim is that people who have enough redirect their wanting onto smaller objects, until a rich society is applying its formidable capacity for dissatisfaction to the temperature of an office or the wording of a menu.

Another problem is that people often react to others around them to formulate their opinion of what is wrong. Someone unhappy with a situation that everybody else describes as fine does not usually conclude that the situation is defective. He concludes that he is. He says nothing — as do others unhappy with the same thing — and the situation collects a unanimous public endorsement that few or none of its members agree with.

That concept has a name, preference falsification, and it accounts for the strangest feature of the world's historical record: regimes with almost no genuine support can look immovable on a Monday and be gone by Friday.[8]

Nobody's mind changed on Tuesday. The change comes when they start to share opinions about what they want to happen.

Nothing shows that the wanting stops. The wanting goes invisible, to the institution, to the historian, and above all to the other people who share it. Full magnitude, pointing where it always pointed, and showing up in none of the instruments.

That's a measurement failure. Which is lucky, because measurement is something we know how to improve.

Another problem is one I can't answer. A person can see a problem exactly, describe it accurately, know precisely who is responsible, and do nothing — because acting is expensive, the odds are bad, or the last four people who tried it were destroyed. Neither habituation nor conformity is at work. Arithmetic is, and it's the most common way for a live vector to produce no motion whatever. Nothing here dissolves it. The sum does change when the cost of finding the other people who already agree with you falls close to zero — and that cost has lately fallen further, and faster, than at any time since movable type.

SURVIVAL IS NOT A DIFFERENT STORY FROM IMPROVEMENT

Modern debates tend to file survival and improvement in separate drawers, as if one were the grim business of staying alive and the other an optional luxury for people who already have luxury. The two are the same story told at different levels of capacity.

Watch how the question evolves. Once a child survives infancy, the next question is whether the child can be taught anything. Feed a family reliably and you have not finished; you have arrived at the problem of keeping food through a winter. A village safe from raiders may discover that its real difficulty is settling a dispute between two of its own families, ideally without burning the place down. Each answer builds the platform that can lead to a more complex better, which becomes thinkable, and the impossible question of one generation ends up as the floor beneath the next one's feet.

That is why the pursuit of the question looks so different depending on where the questioners stand inside it. To someone at the bottom, the person worrying about dispute resolution looks absurdly comfortable. To someone at the top, the earlier question looks like something that was solved long ago and needn't be revisited. Both are wrong in the same way. One ladder, and every rung was somebody's whole life.

The path was neither gentle nor pure. Better tools made better harvests and better weapons, and the same forge served both. Writing preserved wisdom and taxes. Roads carried medicine and armies, usually in the same season. The state coordinated irrigation and oppression, and frequently required the second to accomplish the first. Every new power made people both freer and more dangerous, in the same motion.

The mathematics of civilization is therefore not the story of goodness automatically defeating evil. It is the story of consequences becoming visible, mistakes becoming expensive, and systems being selected or abandoned under pressure.

WHY WE KEEP GETTING UP IN THE MORNING

Human history includes famine, plague, slavery, genocide, empire, industrial pollution, total war, and the deliberate manufacture of lies. It also contains humans' repeated refusal to accept those conditions as final. After a collapse, people rebuild. After a disease is understood, they organize prevention. After rights are denied, those excluded force the definition of citizenship to widen. The recovery is never complete, and the act of recovery is still the strongest evidence there is.

So why *do* we get up? Some of it, simply, is that living systems resist death, which isn't a virtue and needs no explanation. Memory works, too: a person who remembers a better winter has a comparison, between what is and what could be, the smallest engine

capable of producing effort. Parents contribute by handing off unfinished business to their children, either the most generous thing our species does or the most burdensome, depending on the parent.

But the largest share belongs to cooperation, because it creates surplus. One person can defend a doorway. Many people can build a water system, and the water system changes the mathematics. Surplus creates time; time creates learning; learning creates better tools; better tools create more surplus. That loop is the whole of it, and it isn't a metaphor. You can watch it run in any village that gets a well.

Watch it closely, in fact, because the first thing a well produces is not water. It is hours. If the water source is an hour's walk away, somebody makes that walk two or three times a day carrying twenty kilos, and that somebody is almost always a woman or a girl. Put a well in the village and you haven't merely improved the water supply. You have handed back several hours a day to a specific person, and what she does with them is not predictable. Some of those hours go to a daughter who now attends school, because walking to fetch the water was the reason she did not. Some go into a woman's small business, because time was the only capital anyone in that household had to invest. Some hours go to rest, which is a legitimate use and needs no justification.

That well does the same mathematics as the washing machine in the Preface, arriving from the other end of the wealth distribution. The technology was never the achievement. Measure the hour it gave back, and watch what somebody does with it.

Just as readily, the loop runs backward. Fear narrows time horizons. Narrow time horizons reward extraction. Extraction weakens trust. Weak trust raises the cost of coordinating anything at all, which produces less surplus, less agency, and more fear to start

the circuit again. Civilization advances when it finds structures that reverse the direction of that loop, and collapses when it forgets which way it is turning.

NEXT IS STRUCTURAL

Next gets mistaken for novelty all the time, and the mistake is expensive. New is not necessarily better. A new weapon, a new fraud, or a new addiction can be considerably worse than what it replaced, although each will be marketed as an advance.

NEXT means whatever comes along and does better than the old way once more of the costs are counted. It may be invented long before it becomes dominant. It waits until costs fall, infrastructure appears, people learn, or the failures of the old structure become undeniable — and the waiting can be very long indeed, because what stands in the way is rarely the technology.

In 1847 a Hungarian obstetrician, Dr. Ignaz Semmelweis, working in Vienna, noticed something that should have been impossible to ignore.

The Vienna General Hospital ran two maternity clinics. Poor women were assigned to one or the other, essentially by the day of the week. In the first clinic, staffed by physicians and medical students, women died of puerperal fever at rates that ran several times higher than in the second, which was staffed by midwives. The difference was large, it was consistent, and it was notorious. Women arriving in labor begged to be sent to the second clinic. Some of them, when they found out only the first clinic was available that day, chose to give birth in the street instead — and those women, Semmelweis observed, did better than the ones who made it inside.

He worked the problem the way you'd hope anyone would. He eliminated the explanations on offer one at a time: overcrowding, climate, the position women were placed in to deliver, even the route the priest walked through the ward. None of them survived contact

with the data. Then a colleague of Semmelweis cut his finger during an autopsy and died of something that looked, in every particular, like the fever killing the women.

That was the answer. The physicians and students in the first clinic performed autopsies in the morning and deliveries afterward. The midwives did not. Something was traveling from the dead to the living on the hands of the men who were supposed to be helping. Semmelweis had no germ theory to explain it, that was decades away, so he called it cadaverous particles and ordered the one intervention the hypothesis implied. Wash your hands in chlorinated lime before examining a patient.

The death rate in his clinic dropped. And his profession rejected him. Not quietly, and not briefly. His findings were dismissed, his appointments were not renewed, and his argument was received as an insult, which in a sense it was: the claim was that doctors were carrying death to women in childbirth on their own hands, and it was made to men whose standing rested on the proposition that a gentleman's hands were clean.

Semmelweis got angrier and less politic as the years went on, which didn't help him. He died in an asylum in 1865, still largely unheard.[9] The principle he had demonstrated with a bucket of chlorine was accepted later, only after Louis Pasteur and Joseph Lister applied science to his theories. By then he was dead and so were a great many women.

The triumphant version of this story is a lie.

Everything the mathematics of better required was present in Vienna in 1847. There was an unmet need, measured to the decimal. There was memory, in the form of records kept clinic by clinic. There was an experiment, and it worked. What was missing was selection, the part where the better arrangement gets adopted because it is better. Selection failed for twenty years, and it failed for reasons that were entirely human: the finding of what caused the infection

implicated the people who would have to act on it, no theory existed to make it comfortable, and the men who could have changed the practice were the men it accused.

The better next existed. It sat, demonstrated, killing nobody, saving lives in one clinic in one city, while the institution that could have spread it to the rest of the world declined to look. Nothing about that refutes the pattern. The pattern runs slowly, against friction, in the direction it eventually runs anyway. The cost of the delay was paid in full by those who had no part in the argument.

Not the physicians who resisted, and not Semmelweis. The victims were women in labor who never heard the dispute, could not have followed it, and had no standing in it whatsoever. That's the ordinary distribution.

When a better arrangement is available and an institution declines to adopt it, the delay is rarely paid for by the people doing the declining, which is why the declining can go on so long. A cost that lands somewhere else looks like the way things are, and nobody sends an invoice.

Better structures do tend to win eventually. Somebody is always living inside the word eventually.

That is also why historical change so often looks impossible and then, suddenly, inevitable, with nothing in between. The printing press existed before mass literacy, and electric motors turned in workshops before any house was wired for them. Solar cells existed before scale and manufacturing drove costs down. Artificial intelligence existed in research laboratories before networks, data, chips, and interfaces made it ordinary. The S-curve is not magic. It is accumulated readiness becoming visible, and most of the accumulation isn't technical. Mostly it's people changing their minds about something, or dying and being replaced by people who never held the old view in the first place.

The persistent pursuit of a better next is not a slogan pasted onto history. It is a short description of how things actually change: people need something, remember what was tried, invent, and pick the better thing, except when the picking stalls. The process doesn't promise that the best option wins immediately. It says only that a more useful structure keeps attracting energy once people can perceive and access it, and that the perception is usually the slow step.

Somebody drilled a hole in a piece of bone, and thereafter, two continents filled with people. Somebody washed his hands before delivering a baby and was thrown out of the profession for it. Both of those are the same plot, and neither was obvious at the time.

Which leaves the uncomfortable half of the question. The Viennese physicians were not stupid. They were the educated and trained men of a serious institution, working from the best theory available, and they were wrong in a way that killed women for two decades while the evidence sat in their own ledgers. There's no reason to think we're any better at recognizing the thing in front of us. The difficulty is that it never looks like a thing in front of us. It looks like the way things are done.

* * *

CHAPTER 3

Scarcity Was a Stage, Not a Destiny

Why reproducibility keeps changing the economics of power

The fact of domination is not proof that human beings are born for domination. It is evidence that scarcity once made control an efficient way to coordinate limited resources.

Nothing in that sentence says people are nice. It says that the old mathematics rewarded certain behavior for a very long time, and that we've been mistaking its result for the character of the species running it.

Think about the problem a ruler might have been solving. Information moved at the speed of a horse. Food spoiled in days. Energy came out of muscle, wood, wind, falling water, and eventually a seam of coal that a miner had to be sent underground to cut. A failed harvest could erase a community in a single winter, and nobody could see it coming in time to prevent it.

Under those conditions, concentrated authority wasn't merely tempting. It worked. A king could mobilize grain, labor, walls, and armies faster than any dispersed population could deliberate its way to the same result, and speed was often the problem. Deliberation costs time. If you're the king when the raiders are four days out, a system that reaches a decision in six days is useless in that case, whatever its merits in principle. Some structures we now find indefensible were, in their period, answers. Not good answers, but answers. Grain was allocated among a population that would otherwise have fought over it, and the alternative on offer was rarely a fairer distribution, only a bloodier one.

It is easy in the present to condemn every older system as stupid or evil. That misses the deeper lesson. Through time, history's First Triumvirate, the secret political alliance formed by Caesar, Pompey, and Crassus in 60 B.C., was followed by feudal obligations, mercantilism, colonial administration, the industrial corporation, and the military-industrial complex. Each solved real coordination problems with the tools and assumptions available. The harm from them was real. Their historical function was also real. A structure can be the best available NEXT in one period and an indefensible obstacle in another.

Here's an example from my life. For thirty years we ran two shoe stores in Boulder instead of one, and because of the second, smaller store, we led the market. It was a good decision.

Then, not long ago, we closed that store, because the numbers had started saying something I didn't want to hear: two stores in that town were no longer better than one. Nothing about the store had changed. The year and the circumstances had. I don't like being wrong. Nobody does. But a decision can be right for thirty years and wrong in the thirty-first. The mathematics that made it right is the same mathematics that now makes it wrong. That's a small thing to know about a shoe store and a large thing to know about a kingdom.

There's an ugly cousin to this argument. Saying that feudalism coordinated a scarce grain supply isn't saying feudalism or its artificial scarcity was fine. Slavery didn't solve a coordination problem. For the people who owned other people, it solved a labor-cost problem, by inventing a scarcity of freedom, assigned by category, that had no basis in the physical world at all. It matters whether a structure was rationing something genuinely limited or manufacturing a limit in order to charge rent on it. Most of history contains both, often inside the same institution, and any account that can't tell them apart isn't worth having.

The mistake, then, was never that humanity passed through scarcity structures. The mistake is treating scarcity as eternal and domination as the final expression of human nature. Once energy, information and making things get cheaper, and more people can take part, the old arrangement starts wasting more than it protects, and it goes on charging for the protection.

THE BINDING CONSTRAINT

Every business has one thing that makes you lie awake. All the other problems wait their turn. Fix those and not much changes; fix the one and everything moves. Farmers call it the poorest nutrient in the soil, engineers call it the bottleneck, and for most of history it was simply whatever there wasn't enough of. There wasn't enough of the thing, wherever the thing was, so you held it, and you kept everyone else from holding it. Walls, granaries, guilds, patents on salt, the enclosure of a common field: the same move, in different centuries, with different vocabulary. Institutions that look monstrous to us were built for that condition, and each was, in its own terms, rational.

Now watch what happens when the cost of another unit collapses. The system stops optimizing for possession, because possession has stopped being the hard part. What becomes hard instead is coordination, access, timing, and trust: knowing what's needed where, moving it there, and being believed when you say you will. And this is the consequence that matters. You cannot dominate your way to coordination. Force is an excellent instrument for taking a thing from someone. It is a poor instrument for getting a thousand people to tell you the truth about what they need, and a useless one for getting them to want to.

None of that is a moral observation, only a statement about which strategies pay once the limiting factor has moved. Nobody has to get kinder. Good behavior follows profit and loss with a reliability that has embarrassed moralists for centuries.

Some goods make the shift visible. A scarce treasure produces gates, secrecy, hierarchy, and a permanent argument about allocation. Reproducible things behave almost the opposite way, because their value can rise when they're shared. A scientific result becomes more valuable once other people test it, break it, and repair it. A tool that serves a second user hasn't been taken from the first. Sell someone an education and they walk out with more capacity than they brought in, which is a strange thing for a transaction to do. And a solar panel does not own the sunlight it converts.

None of this eliminates material limits. Land still has location, minerals still have to come out of the ground, ecosystems still have boundaries that don't care what any of us believe, and human attention is as finite as it was in 1400. Abundance is not infinity. It is a shift in the most important constraint.

You can check the shift in the transition to renewable energy. The fuel for a solar panel arrives during the day without an invoice. Everything else about it costs money, the equipment, the installation, the grid connection, the financing, the storage, the coordination. But the fuel, the sun, doesn't, and it never will. Once the thing is built, the marginal cost of another captured photon approaches zero, even though the system around that photon stubbornly does not. Between 2010 and 2024, the global average cost of utility-scale solar electricity fell about 90 percent, and international agencies reported that new solar power in 2024 was on average cheaper than the lowest-cost fossil-fuel alternatives.[10]

That transition changes behavior, and it changes it the way the binding-constraint argument says it will. When energy was scarce and fuel had to be mined, transported, burned, and defended, control of the deposit created wealth, which is why so much of the twentieth century's diplomacy was, at bottom, about deposits. When energy increasingly comes from manufactured devices placed wherever sun, wind, water, or heat can be reached, the wealth moves toward

engineering, deployment, networks, storage, and intelligent timing. Possession stops being the game. The strategic resource becomes coordinated electrons.

THE OBJECTION I CANNOT DISMISS

If falling marginal cost reliably shifted power away from possession and toward participation, the last thirty years should have been the most distributive period in economic history. Copying a piece of software costs nothing you can measure. Sending a document across the planet costs less still. By that argument, the internet should have spread the power to do things more widely than anything before it.

Instead, it produced some of the most concentrated firms in the history of commerce. A handful of companies own the layers through which most of humanity now searches, communicates, transacts, and increasingly thinks. Their market values exceed the economies of substantial countries. They got there because they decreased the cost of the product. When product or service copies are cheaper, the winner of a category takes nearly all of it, and networking meant the position compounded rather than eroded. Cheaper reproduction didn't democratize industries. It made the industry winner-take-most.

The binding constraint moved. Once reproduction became cheap or free, the scarce thing itself became the coordination layer, the index, the marketplace, the shoe store replacing the cobbler, the standard, the protocol, the place where buyers and sellers can find one another. Whoever owns the new binding constraint captures the value the old one used to hold, just as thoroughly and with rather better public benefit.

There's only one question worth asking about any technology, which is never whether it makes something cheap. Ask what becomes scarce as a result, and who ends up owning that.

What keeps me optimistic about the direction is that this new kind of concentration is weaker than the old one, in a specific way you can check. Oil sits under somebody's soil and stays there. A coordination layer exists only because people keep choosing to stand in it, which means it can be replaced by a better one, competed with, regulated, forked, or walked away from, and has been, repeatedly, by firms that looked permanent at the time and now serve as cautionary examples in business schools. Owning a scarcity people can leave is a much less durable position than owning one they can't.

Weaker isn't weak, and slower to compound isn't fair. Anyone awake for the last few decades should know that a distribution that is theoretically available doesn't arrive on its own. But the direction of the pressure is real, and it runs toward participation for the same reason it always has: a coordination layer that serves more people is worth more than one that serves fewer, and that mathematics doesn't require anyone involved to be generous.

THE MORAL LESSON FOLLOWS THE ECONOMIC LESSON

That brings me to inclusion, and to why it belongs in the arithmetic rather than in the charity.

We describe inclusion as something the fortunate extend to the unfortunate, paid for by surplus and justified on moral grounds. In a system where the binding constraint is coordination rather than possession, that description is an accounting error. A healthy child is a future participant in every market that child will ever enter. Teach an adult to read and you've produced a problem solver, a worker, a customer, a voter, and a parent who will raise other people who read. Stability next door means a larger market and a quieter border, and whatever was being spent on the border becomes available for something else.

Each of those is a return, not a donation. The old mathematics asked: how much can I keep by limiting access? The emerging mathematics asks: how much more can we create by expanding

capability?

Expanding access and mandating outcomes are separate operations, and only the first follows from the arithmetic. Opening a door adds capacity. Assigning someone a position behind that door regardless of whether they can hold it is a different proposition, defended on different grounds, and it tends to damage the person it was written for. Nothing in the mathematics of abundance requires that second move. And mechanisms built to widen access get captured, reliably, by people whose interest is the mechanism rather than the widening.

The transition won't happen because anyone becomes selfless. It happens because once the tools of production and knowledge are widely distributed, participation often becomes more profitable than exclusion. And because the people running exclusion strategies eventually notice they're losing money.

MANUFACTURED SCARCITY IS NOT A HISTORICAL TOPIC

The old mathematics doesn't belong to the past. It has moved.

In a wealthy economy, the most valuable move is no longer finding something scarce, but taking something that has stopped being scarce and arranging for it to behave as though it still were. Housing is the clearest case. Nothing about construction explains why a dwelling in a productive city costs what it does; we can build faster and better than at any point in history. What explains it is that the right to build has been made scarce, deliberately and through entirely lawful means, by people who already own the housing whose value depends on the shortage continuing. That's a scarcity structure operating inside an abundance economy, and it moves a very large amount of money from young people to old ones every year without showing up on anyone's ledger as a transfer. I'm one of the old ones. I'm not proud of the arrangement.

Once you know its shape, the same architecture turns up elsewhere. Occupational licensing that protects incumbents rather than clients. Intellectual property terms extended long past any plausible incentive effect, so that a work enters the public domain about when everyone who cared about it is dead. A medicine whose molecule costs pennies and whose price is set by a patent thicket. In each case the physical constraint has relaxed and the institutional constraint has been tightened to compensate, which is the old mathematics performed with paperwork instead of walls.

Abundance has a counter-move. When a real scarcity dissolves, the people whose position depended on it don't generally give up. They go and find a chokepoint, and the chokepoints available in a rich society are mostly legal rather than physical.

So here's a question, not a rhetorical one. Whenever you meet a scarcity, of housing, of credentials, of medicine, of access to almost anything, ask what would happen to the price if satisfying the scarcity was profitable.

That question doesn't settle anything by itself. Some manufactured scarcities are defensible; a patent system with no term at all would produce less medicine, not more. But it separates the two conversations, and most public argument about prices goes badly because the two are conducted as if they were one.

Scarcity can be profitable for a while. Expanded capability compounds.

THE EQUATION

Falling marginal cost + wider access + more participants = a structural shift from guarding value to growing value.

CHAPTER 4

Eight Billion Mornings

Population is not only a burden on systems. It is also the source of every new idea.

"Population, when unchecked, increases in a geometrical ratio. Subsistence only increases in an arithmetical ratio."

Thomas Malthus, An Essay on the Principle of Population (1798)

Eight billion people are not one crowd. They are 8 billion centers of experience waking into different circumstances with a common capacity to prefer a better condition.

In 2024 the United Nations estimated 8.2 billion people on the planet, and the number has a way of stopping thought rather than starting it.[11] We hear it as a statistic, a threat, a humanitarian problem, a market, a political abstraction. Civilization, though, is lived one nervous system at a time, and it starts each day with somebody opening their eyes into a particular situation and deciding what to do about it.

So here are ten of those mornings, chosen to make the number visible, not to prove anything.

TEN MORNINGS

A rice farmer checks the monsoon forecast on a phone that didn't exist in her district ten years ago. The forecast changes one decision, whether to transplant today or wait, and that decision protects or destroys a season's income for a household of five. She has no opinion about the mathematics of civilization. She's thinking about rain.

A cardiac surgeon in São Paulo has been awake since four-thirty, reviewing imaging for a mitral valve repair scheduled at seven. The operation depends on an ultrasound protocol developed in Austria, a synthetic graft made in Minnesota, and an anesthetic technique refined across three continents over forty years. She has met none of the people whose work makes hers possible. She doesn't need to. The knowledge was cheap enough to travel.

A twelve-year-old in Nairobi opens a textbook that arrived because someone, somewhere, decided the cost of printing and shipping it was worth what the child might eventually do with it. Whether that bet pays off depends on so many later decisions, hers, her teachers', her government's, that nobody could compute the return. But the bet was placed, and the book is in her hands at seven in the morning, and the day starts from there.

A long-haul truck driver outside Frankfurt wakes in a rest-stop cab, eats something packaged, and drives a logistics chain that feeds nine countries. He covers in a day what a horse-drawn freight wagon needed the better part of a month for, and carries something like twenty times the load while doing it. He doesn't experience this as progress. He experiences it as a back that hurts and a delivery window that won't move.

A water engineer in Chennai walks a desalination plant at dawn, adjusting the pressure in a reverse-osmosis membrane that turns the Bay of Bengal into drinking water for more than a million people. Her plant and its sister to the north supply nearly a third of everything the city drinks.[12] The technology is energy-hungry, expensive, and not enough; the city still rations supply to many neighborhoods. She knows the system falls short. She also knows that ten years ago the alternative was trucks.

A subsistence farmer on the Mozambican coast wakes before light to check a fish trap. No GDP figure will ever record what he catches. If the trap is empty he'll eat less today, and so will his children, and

no institution anywhere is tracking that.

In Bangalore, a software developer opens a laptop at a kitchen table while her daughter eats breakfast. She'll spend the morning writing code for a payments company in Stockholm whose employees she has never met. Her salary, adjusted for purchasing power, is roughly a fifth of what a developer in Stockholm earns for the same work. The gap is closing, year by year, because the code doesn't know where it was written.

In rural Bolivia, a grandmother walks forty minutes to a clinic that may or may not have the blood-pressure medication she was promised last month. Forty minutes each way, on a road that washes out in the wet season. She goes anyway, because the alternative has a consequence she can measure in how she feels by the afternoon.

A construction foreman in Lagos arrives at a site where a residential tower is going up, poured concrete and rebar, fourteen stories, financed by a pension fund in London. He manages a crew of sixty and has no engineering degree. Everything he knows he learned on earlier sites, from men who learned the same way. Whether the building stands for fifty years depends partly on decisions he'll make before lunch. It will house families who are currently paying more than they can afford for less space than they need.

In rural Uttar Pradesh, a teenager is awake before dawn, studying by phone light for an exam that could move her from a village school to a district college. The lessons come from a free app, the phone belongs to her father, and the charge on its battery decides how long her morning lasts.

None of these people chose the same profession, speak the same language, or hold the same politics. What they share is narrower than any ideology and wider than any culture. Each woke into a condition, judged it against something better, and started acting on the difference. The farmer wants the crop to live. The surgeon wants the valve to hold. The child wants to understand the page, the foreman

needs the pour to set before lunch, the teenager needs enough battery to finish the chapter. Preferences acted on, eight billion times, every morning. Repetition makes the pattern, not any one of them.

FROM BURDEN TO CAPACITY

Everyone's instinct is to simplify. Eight billion becomes a crowd, a mass, a problem of resources divided by mouths. Population debates land on one of two positions, too many, straining the planet, or enough, if only distribution were fair, and both treat the number as a quantity to be managed. Neither camp has moved the other in fifty years.

I'd rather treat it as a count of active agents. Each person can learn, work, invent, care, trade, and choose, and the network those people make is the largest such network there has ever been, by an order of magnitude. Athens, Rome, the Song dynasty, the British Empire: every earlier civilization ran with a fraction of the participants now alive and awake on any given Tuesday.

Eight billion people can produce crises of coordination, resource competition, and mutual harm at a scale no earlier civilization had to manage, and can also produce redundancy, specialization, and problem-solving at a scale none of them could reach. Which we get depends on how far each person is allowed to grow.

Scarcity thinking treats each additional person mainly as a cost. Another mouth, another carbon emitter, another competitor for a job, a seat, a liter of water. Nobody in that accounting is ever anybody's colleague. And the pressures are real. Systems get overwhelmed when housing, water, education, health, and energy don't scale with the people they serve, and the grandmother walking forty minutes to a clinic that may have nothing for her isn't a hypothetical.

But the same person who needs housing, water, and a school is also a mind that can build housing, purify water, or teach. A billion people struggling to survive produce fewer solutions than a billion

people secure enough to think past next week.

Follow what stability does, one item at a time. Housing lowers emergency costs and gives you a fixed point from which work, learning, family life, and commerce can operate; the grandmother in Bolivia can't hold a steady job while she's walking eighty minutes a day for medication she may not get. Health preserves productive years; a farmer too ill to plant contributes nothing to the harvest, however much land there is. Education makes for better decisions, and better decisions are the thing that compounds. Reliable energy lets nearly every other system run, which is why the engineer in Chennai keeps walking the plant at dawn even though she knows it falls short. Partial is already better than none.

Policy debates frame basic stability as a cost that successful people pay for unsuccessful people. Run the numbers and it goes the other way. Stability does not merely consume wealth; it creates the conditions in which wealth can be produced by more people. Every adult whose blood pressure is managed is a worker, a parent, a customer, and a voter who stays in the system instead of leaving it. None of this requires generosity, though generosity isn't refused. It requires noticing that the platform produces more than it costs, once you count what the people standing on it go on to do.

A world of eight billion struggling people is unstable and economically small relative to its potential. A world of eight billion happily housed people, safe enough to plan, educated enough to choose, connected enough to participate, and supplied with clean energy, would be the largest productive and consumer market in human history.

Be exact about how little that promises. Such a world would still contain disagreement, failure, competition, grief, cruelty, and terrible decisions. Happily housed describes a platform of dignity, not permanent happiness. What changes is the denominator. Instead of a small fraction of the species generating the ideas, the products, the

art, the governance, and the markets that everyone else lives inside, a larger fraction takes part, and the returns to taking part compound in a way the returns to keeping people out never have.

A MARKET IS A NETWORK OF AGENCY

When economists describe a market they usually mean a place where buyers meet sellers. At the scale of a civilization a market is something more particular: a network of expressed preferences. More people able to form a preference, act on it, and signal what they need means more information in the network. Better information gets things where they're needed, shows unmet needs sooner, and rewards the answers that actually fit.

Which is why domination, over time, limits the dominator. A ruler can extract wealth from a silenced population, but silenced people generate weak signals. Their needs stay invisible until they become emergencies. Their talents go undeveloped. Their purchasing power is suppressed, and their fear has to be policed, policing being one of the few services that gets more expensive the worse things get. Extraction captures value that already exists. Agency creates value that did not exist before, and the difference between the two widens as the network grows.

Go back to the mornings. The surgeon in São Paulo draws on three continents because the knowledge was allowed to travel. The developer in Bangalore earns a living from a company eight time zones away because the code was allowed to travel. The teenager in Uttar Pradesh studies on a free app because someone decided the lessons were worth distributing rather than hoarding. In each case the system that let the signal move got more out of the person than a system that blocked it would have.

Across borders the logic is the same. A secure nation needs lawful systems and functional boundaries, and it also benefits when conditions improve on every side of those boundaries. Neighbors with

opportunity produce trade, stability, knowledge, and quieter borders. Neighbors without it produce pressure, and pressure is the one import no wall has ever managed to tax.

The better next is not borderlessness. It is security understood as a shared system outcome rather than a lonely fortress. A fortress is a bet that the other side stays poor, and it's a bad bet, not because it's unkind but because it has to be renewed every year, at rising cost, against people who are trying exactly as hard as you are.

THE EQUATION

Eight billion people x greater capability x voluntary exchange = the largest possible engine of sustainable prosperity.

CHAPTER 5

Good Is More Profitable Than Bad

Morality and economics increasingly point in the same direction

Bad pays. It pays well, for a long time, and often to the people who did it rather than punishing their grandchildren on their behalf.

In 1860 the roughly four million people held as property in the United States carried an aggregate market value somewhere near three billion dollars, which made them the largest single class of privately held asset in the country — worth more, by most reckonings, than every American railroad and factory combined.[13] Cotton grown by those people accounted for something like three-fifths of the nation's exports. Nobody was running this at the disreputable edge of a respectable economy. It was the respectable economy's principal holding, and the men who owned that holding were, per head, the wealthiest men in North America. An institution does not survive two and a half centuries inside a competitive market by losing money.

Fraud makes the same point with a lot less machinery. Bernard Madoff's fund reported, at the end, something close to sixty-five billion dollars of client assets that did not exist. The Securities and Exchange Commission examined him repeatedly across sixteen years and closed the file each time. What finally ended the scheme wasn't a regulator or an auditor. It was a market crash that produced redemptions he couldn't cover. He was seventy years old when the phone stopped working. A man fifteen years older, running the identical scheme, would have died in his own bed with his reputation intact and his name on a wing of a hospital.

Pollution is where the objection stops being theoretical and starts being profitable. A plant that installs a scrubber pays for the scrubber. Its competitor across the river installs nothing, sells the same product for less, and takes the customer. Where the discharge carries no price, the clean firm isn't being virtuous at a cost but outcompeted, for a reason the market has correctly identified as inefficiency. Virtue under those conditions is a business plan for going out of business, and calling the resulting industry morally defective mistakes an equilibrium for a character flaw.

By 1930 the medical literature described asbestos as a killer of the men who worked it; Britain wrote regulations for the trade in 1931; the American liability wave broke in the late 1970s.[14] Forty-five years is not a market signal. Whatever the mathematics eventually does, it does to a different set of shareholders, a different board, and in most cases a different century. No correction that arrives after the decision-makers are dead can work as an incentive, and an incentive nobody can anticipate is not an incentive at all.

On the evidence as usually presented, the objection wins.

THE UNRECORDED COST

Good is not only a moral category. Given enough time, and counted widely enough, it often describes what simply works better. Read as a general statement about the world, inside a fixed accounting boundary, over the horizon of a career, "good is more profitable than bad" is simply false, and no number of encouraging examples will rescue it. The truth is stranger. Bad, almost without exception, depends on a cost that is real and unrecorded. Change what gets recorded and the profitability changes with it. Nobody has to repent for this to happen. The mathematics moves and behavior follows it, which is the only kind of moral progress that has ever operated at scale.

So it comes down to bookkeeping.

Start with the hardest of the three cases. Slavery made individual men rich, and it did so across generations. What it conspicuously failed to do was build. By 1860 the states holding a third of the country's population were producing something under a tenth of its manufactured goods. Patents issued per head in the free states ran at several times the southern rate through the decades before the war. White illiteracy in the South stood at a large multiple of the New England figure, and the enslaved population was forbidden by law from learning to read at all, which meant that the single largest labor force in the country had been legally prevented from contributing an idea.

None of that appears as a cost in any planter's ledger, because in the ordinary sense it is not one. It amounts to the absence of a return nobody ever generated: the machine tools not designed, the firms not started, four million people spending their lives producing at the level a whip can enforce rather than at the level a mind can reach. Coercion has a ceiling built into it and the ceiling is low. You can compel a body to pick cotton. Nobody has ever compelled anyone to have an idea, and by 1860 ideas were fast becoming the thing that made a region rich. So the men who owned those assets were wealthy, and the region that ran on cotton was poor, and stayed poor for a century after the asset was liquidated.

HOW A COST GOES MISSING

Although double-entry bookkeeping is roughly 500 years old in its codified form and older still in practice, it may be the most successful information technology in commercial history.[15] Its central discipline is that no entry stands alone. Every debit has a matching credit, every transaction has two sides, and a set of books that will not balance announces that something has been left out. The whole system exists to make omission visible.

It has one blind spot, and the blind spot is structural, not moral. A transaction gets recorded because there's a counterparty, somebody on the other side who can present an invoice, sue, withhold delivery, or simply be owed. Where the party bearing a cost has no standing to present anything, that cost is unbookable rather than concealed. The farm downwind of the smelter, the fishery below the outfall pipe, the child who will be born in 2065 into whatever atmosphere we are currently assembling: each is a genuine counterparty to a genuine transaction, and not one of them can issue an invoice. So the books balance perfectly, and as a description of what actually happened they're false.

Which is why the fix has never happened, and I've stopped waiting for one. The creation of a claimant creates the missing entry. A tort doctrine that lets the downwind farmer sue. A regulator with the power to fine, or a disclosure rule that forces the number onto a page where an investor will read it. A satellite that photographs a methane plume from four hundred miles up and publishes the operator's name beside it. Each of those is a piece of accounting infrastructure wearing other clothes, and each converts an unbookable cost into a booked one.

Booking a future cost requires deciding what a future cost is worth now, and that decision rests on a single number nobody can derive from evidence. Nicholas Stern and William Nordhaus disagreed for years about the social discount rate, which determines how economists value future economic damages from climate policy. And a very large part of that disagreement came down to neither physics nor economics but to that one parameter. I've never seen it argued to a conclusion. Accounting can tell you a cost exists but not what it will be on a stranger's grandchild, and it turns evasive when asked.

THE WIDER LEDGER

The mathematics of better uses a wider ledger. Who benefits and who pays? And it insists those be two lists rather than one. How long does the benefit last? Nobody should mistake a gain that expires within a decade for a gain that compounds across a century wearing a different size. What new capability was created, if any, since a transfer moves value from one pocket to another and an investment makes more of it than there was? And can the structure continue without escalating force or deception? The fourth question matters most: an arrangement requiring more coercion every year to hold its position is publishing its own accounts whether it means to or not.

This does not eliminate difficult tradeoffs. It prevents a local gain from being mislabeled as a universal success.

Stated as a test you can carry: widen the frame, more people, more consequences, more time, and see whether the benefit survives the widening. Anything that only looks like a gain at one boundary was never a gain. It was a boundary.

Aimed at a policy, it asks who is downwind. Aimed at a company, it asks which of that company's costs are carried by people with no way to invoice it. Aimed at an empire, the question becomes how much of the return is simply a transfer from the governed. And aimed at a single person's single act, it asks the same thing in the same words, with the frame widened by exactly one human being. The mathematics doesn't change size when the subject does.

FOUR TIMES THE BILL ARRIVED

Since insurance exists to put a price on a cost that has not happened yet, it works by observing how often a thing goes wrong across a large population, and how badly. One requirement, though. Whatever the loss turns out to be, it has to appear inside a window the underwriter can observe.

Because mesothelioma takes many years to emerge and be diagnosed, that requirement was seldom met. For the whole period in which decisions were being made, which mine to open, which product to sell, which warning to leave off the bag, the actuarial machinery reported no loss, because no loss had yet occurred. An instrument functioning correctly, returning an answer that was catastrophically wrong.

When Johns-Manville filed for bankruptcy protection in August 1982, it was the largest asbestos producer in the United States and a profitable company, carrying something like sixteen thousand pending claims and its own correspondence from the 1930s in the plaintiffs' hands.[16] The trust established to settle those claims was funded with roughly two and a half billion dollars, and it has spent decades paying claimants a small fraction of assessed value, in recent years around five cents on the dollar. RAND's accounting of American asbestos litigation found that claimants received about forty-two cents of every dollar spent. Defense costs took thirty-one, plaintiffs' lawyers and related costs the remaining twenty-seven; by 2002 the total spent had reached something near seventy billion dollars.[17]

So the bill arrived, at the right industry, four decades late, and roughly three-fifths of it was consumed in the act of delivery.

Insurance didn't find this. Courts did. Actuarial practice detected nothing until the claims themselves arrived, at which point the exercise was no longer detection but subtraction. What the industry did, on reflection, was write asbestos exclusions into its policies, which is to say the instrument's contribution was to stop covering the risk at the precise moment the risk became knowable. Any argument that markets correct themselves has to sit with the fact that here the correcting was done by a judge.

Fraud raises verification costs. Short enough to sound like a slogan; the number attached to it isn't. Enron entered bankruptcy in December 2001; shareholder losses were later claimed in the order of

$74 billion; its auditor, Arthur Andersen, with 85,000 employees worldwide, some 25,000 of them in the United States, had ceased to exist within the year.[18] After Congress passed the Sarbanes-Oxley Act in 2002, Section 404 of that act required every public company in the country to document and test its internal financial controls. In the first year of compliance, large American public companies reported spending an average of something over four million dollars each on that section alone. Every one of them paid it. Almost none of them had done anything wrong.

So that's what a hidden cost looks like when it arrives. Enron no longer existed to receive it. What the cost became instead was a standing tax on everybody obliged to prove they are not Enron, and it is still being paid, two decades on, by firms whose executives were in secondary school at the time.

When the ledger runs the other way, the figures are occasionally almost embarrassing. The Environmental Protection Agency's retrospective study of the Clean Air Act put the benefits of the 1970-1990 program at something near $22 trillion against costs of about $520 billion, a ratio in the neighborhood of forty to one. The error bars matter: EPA's own range ran from $6 trillion to $50 trillion, which is an eight-fold spread and a reminder that the midpoint of a very wide estimate is still a very wide estimate. Even at the bottom of the range the ratio is better than ten to one.[19] A ratio that large should provoke suspicion well before it provokes applause, so here is what the number is made of. The overwhelming majority of the benefit comes from avoided premature death from fine particulate matter, and converting avoided deaths into dollars requires a value-of-statistical-life figure that is an assumption dressed as a measurement. Halve it, and halve it again, and you're still holding a ratio of ten to one. Read the study that way: the precise number is soft and the sign is not.

The northern cod fishery off Newfoundland is the one I find hardest to put down, because nobody in it was hiding anything.

Canada's fisheries scientists were counting. The index they counted with was catch per unit of effort, how much fish a standard vessel brought in for a standard amount of fishing, and it's a perfectly reasonable index in a world where the vessels stay the same. Through the 1970s and 1980s the vessels did not stay the same. Sonar, satellite navigation, larger nets, and freezer trawlers that could stay out until they were full each meant that the same nominal effort took more fish out of a thinner population. The index held steady while the stock fell out from underneath it. What the instrument was measuring was the technology, and it reported the result as abundance.

In July 1992 Canada closed the fishery. Around 30,000 people in Newfoundland lost their work in a single announcement, in an economy that had done that work for four hundred years, and federal aid programs eventually ran past two billion dollars. Spawning biomass had fallen to roughly one percent of its 1960s level. In 2024, thirty-two years later, Ottawa reopened a limited commercial fishery at a quota near 18,000 metric tons. In 1968 the catch had been around 800,000 metric tons.[20]

A firm's books miss a cost because no counterparty can invoice. Off Newfoundland there was one: a federal department with a legal mandate, a research fleet, and scientists whose entire job was to present that invoice. Standing was never the problem there. Measurement was, and the measurement failed in the direction that happened to be convenient for everyone being measured, which is the direction measurements usually fail.

WHAT HAS ACTUALLY CHANGED

Bad pays across the interval between the gain and the bill. Everything depends on the length of that interval, and the interval is a function of what can be measured. Asbestos ran forty-five years because the only instrument capable of seeing it was a cancer registry that nobody was reading against a customer list. Now a methane leak

runs for weeks, because there's a satellite above it and a reporter with access to the data. Water contamination that would once have taken a generation to attribute is attributed in a season. Across one domain after another, the lag is closing, unevenly and against resistance, and profitability follows the lag because profitability is only ever an estimate made inside a boundary.

Trust is the same mechanism run forward instead of backward. A merchant who people believe spends less proving every claim, meets less friction on the second sale, and survives a mistake that would finish a stranger. I've lived on that. In the early years of Crocs, before anyone outside Colorado had heard of them, a friend who ran a print shop had done the founders' business cards, and the galley, with eight or so cell phone numbers on it, was in his wastebasket. I asked if I could have it. When the shoe became one of the hottest products in the world and nobody could get stock, I could call those men, laugh with them about whatever was going on in sports, and get product. We sold a hundred pairs before lunch some days. None of that was a contract. It was a few people who believed each other, which is cheaper than a contract and, while it lasts, worth more. An institution people believe can move a population in a week. Trust is not softness. It is stored efficiency, and like anything stored it takes years to build and an afternoon to spend, because every interaction afterwards has to carry the memory of the betrayal.

Truth behaves the same way and for the same reason. A false story travels faster than a careful explanation and always will. What it can't do is stop the world from being what it is. No engineer has ever flattered a bridge into standing. No physician has voted an infection away. Confidence can be inflated for a considerable while, and then the cash flow arrives and declines to participate. Truth is profitable because it lowers the cost of correction, and correction is the largest line item in any long project.

Clean energy is the case where the two ledgers can be watched converging. Fossil fuels created immense prosperity because they concentrated portable energy when civilization needed it, and they were a better next than muscle, wood, and whale oil for most of what they were asked to do. Then the costs came into view, one instrument at a time: the smoke, the mines, the dependence on unfriendly governments, the carbon piling up in the air. Renewable systems are becoming morally attractive and economically dominant for the same reason: they reduce recurring fuel cost and external damage while expanding local resilience.

None of which requires a moral awakening, and I wouldn't advise waiting for one. Profit is available as the transition mechanism. Once the rules, the land, the finance and the demand make building more profitable than defending a shortage, a developer will build abundant housing. An energy company deploys storage when the return beats the return on defending a declining asset, and nothing in the boardroom has to change first. Even the outrage business is subject to this, though it's the slowest of the three: verified clarity gets sold when enough users will pay more for it than advertisers pay for their attention, and that day isn't close. Policy matters because it changes which costs the market can see. Technology matters because it changes what is physically possible, and culture matters because it changes what people reward. Durable change comes when the better option is simply the easier one to take.

Optimism does no engineering. What it gives the engineering is a direction. The idea still has to be tested by building things that get paid for the good they do, to many people over many years, instead of leaving that good lying about unpaid. Where that has been done, you get the result that matters: what is better for everybody becomes even better for those it is best for, because the market grows larger, safer, and more durable.

THE ADDRESS ON THE BILL

A cost coming back is not the same as a cost paid by whoever incurred it.

When the first asbestos verdicts came in, the men who had run Johns-Manville's medical department in the 1930s, who received the reports and decided what to do with them, were already long dead. That bill was settled by shareholders of the 1980s, by insurers who had priced their policies in good faith, and by dying men who collected five cents on the dollar.

The slave economy's account came due in a war that killed something like three-quarters of a million people, most of them young men who had never owned anything; the direct cost of the fighting has been estimated at roughly twice the market value of the human beings it was fought over, an arithmetic that nobody would have chosen and nobody did choose.

The Newfoundland fishing crews who lost everything in 1992 had not built the trawler fleet.

At Bhopal in December 1984 a plant released some forty tons of methyl isocyanate over a sleeping city; the official immediate death toll was under 4,000, the eventual figure is generally put many times higher, and the settlement reached five years later came to $470 million, paid by a corporation that has since been absorbed into a larger one.[21]

So the ledger widens, and it doesn't aim. The mathematics is real and the interval is closing, and I believe both. What I can't tell you is that the closing interval reaches back far enough to collect from the people who were responsible. It doesn't, and it collects instead from whoever happens to be standing there on the day of arrival, and that is generally the group with the least capacity to have prevented the thing and the least room to absorb it.

The cod in Newfoundland are the plainest version. Thirty-two years after the closure, the quota stands at about two percent of what those waters once landed. The men who ran the draggers are old or gone, the companies restructured or sold, the aid money spent, the outports emptied. Nobody who took the fish has paid for the fish. The balance is carried by a species that has not come back, a coastline that has not recovered, and a generation of people who moved away and did not return.

That is a real cost, correctly identified, and still outstanding. The mathematics is sound. The collection is not.

* * *

PART II

The Structure of Now

Agency, electrons, truth, and power are the variables already moving.

The second question is whether the pattern can be understood well enough to tend.

CHAPTER 6

Agency Is the Engine

The better next depends less on obedience than on people who can choose

"It is not from the benevolence of the butcher, the brewer, or the baker that we expect our dinner, but from their regard to their own interest."

Adam Smith, The Wealth of Nations (1776)

Agency is the capacity to perceive choices, evaluate them, act, learn from consequences, and revise. Civilization improves when more people possess more of it.

A man walks up his neighbor's path with a gun.

He has perceived a choice. He has evaluated it, probably for weeks, and probably in more detail than most of us bring to anything we do. He is about to act. He will certainly learn from what follows, if he survives it. Measure him against the definition at the head of the page, the capacity to perceive choices, evaluate them, act, learn from consequences, and revise, and he satisfies every clause. Every one. He is decisive where most people are merely opinionated. He accepts a cost the rest of us decline to accept. In one afternoon he raises his own capacity to act by more than schooling or income or a change in the law will raise it across a lifetime.

So if agency is the engine, why isn't he its finest expression? I'm not setting up a debating trick, and the case isn't a marginal one. Much of history has this shape. The men who most enlarged their own capacity to act generally did it by contracting somebody else's, and they did not think of themselves as thieves while doing it. A conqueror acquires the ability to move armies across a province by

removing the ability of the people living there to keep him out. When a landlord encloses a common, he acquires the ability to farm at scale by removing the ability of forty families to graze a cow. Monopoly performs the same operation on a price: one party gains the power to set it because everybody else has lost the power to offer a different one. Each of those is a spectacular expansion of somebody's agency achieved by subtraction, and each was defended in its own century in language uncomfortably close to the language I've just been using.

Which suggests that agency is a fixed quantity in the room, and the only real question is who walks out holding it. If that's true, a book urging more agency for more people isn't proposing a better arrangement, only a different winner.

That definition is mine, and it doesn't exclude the man on the path. A definition that admits the case it was written to rule out is only describing alertness, and alertness is available to anybody. The fix is to ask what job the word was doing.

WHAT THE WORD WAS SUPPOSED TO BE DOING

A command can produce order without alignment.

People may comply because they are frightened, because they are dependent, because they have been told something untrue, or because they cannot afford to leave. The system looks coordinated from above while storing conflict below. Nothing about the coordination is fake, the trains run, the quotas are met, the square is clean, but the order is being held in place by something other than agreement, and it is holding conflict rather than resolving it. When the pressure changes, the apparent alignment disappears, and everybody is surprised except the people who were being compelled.

Agency produces a different kind of order, by producing information. An independent person expresses a need. She tests an arrangement and finds it wanting. She leaves an institution that is failing her, or stays and argues with it, or builds the thing it refused

to build. Every one of those behaviors emits a signal about what is actually true for her. Noisy signals, contradictory, and frequently foolish. They are also more truthful than obedience, because obedience is the one behavior that looks identical whether the person means it or not.

Agency, in other words, is how a society finds out what's going on inside itself.

The connection seems paradoxical. More agency produces more visible disagreement. A controlled society may appear unified; a free society broadcasts conflict. Yet visible disagreement is the data from which durable alignment can be found.

Consider the difficulty in the form an engineer would recognize. Imagine a bridge designed after the engineers have quietly removed from the model every force pointing in an inconvenient direction. The drawing is beautiful and the calculations balance. On a day with a crosswind the actual bridge comes down, and the inquiry finds that nothing was miscalculated. The wrong thing was calculated correctly. Any resilient design has to carry the loads its designer would rather not think about: the gust, the resonance, the thermal expansion in August, the fatigue in the tenth year. Suppressed vectors do not vanish. They become unmeasured forces. Diversity is not noise to be deleted; it is reality to be modeled.

Now take the next step, what a system does with the information once it has it.

James Watt didn't invent the flyball governor; it was already running in windmills and water mills when he adapted it to the steam engine around 1788.[22] The device is two weighted balls on hinged arms, spun by the engine itself. As the engine speeds up, the balls swing outward and lift a collar; the collar closes the steam valve; the engine slows. As it slows, the balls drop and the valve

opens again. Nobody is watching. Its own output is fed back into its own input with the sign reversed, and the result is a machine that holds a speed it was never told, against a load nobody predicted.

Every control system in use today descends from it, and in 1868 James Clerk Maxwell wrote the paper that turned it into mathematics, because the governors had a habit of hunting, swinging above and below the speed they were supposed to hold, sometimes violently enough to wreck the engine. Maxwell's answer identified the conditions under which the loop stays stable, and the conditions are unforgiving in a way worth stating. A control loop requires three things, and fails if any one is missing: a measurement that reflects the actual state, a path that carries the measurement back, and enough time for the correction to take effect before the next one is applied. Cut the sensor and the system is running open-loop: still working, still confident, entirely blind. Reverse the sign and you have positive feedback, which does not merely fail to correct but amplifies, and amplifies faster the worse things get. Introduce a long enough delay and even a correctly wired loop begins to oscillate, overcorrecting for a condition that has already passed.

Set the man on the path against that description and the problem with him becomes something other than a moral intuition. He isn't merely harming his neighbor. He's removing from the system the one party best placed to report what he's done, and the report is the whole point. A society made of such acts has less agency in it, not more. What it has is an open loop, in which the sensors have been taken out one at a time by the people they were measuring.

There's a hard limit to the analogy. A governor works because somebody chose a setpoint: the engine is meant to run at a particular speed, and the whole apparatus exists to hold it there. No society has a setpoint. There is no number that a human community is supposed to hold, and the choice of what to stabilize around is precisely the political question that the mechanism cannot answer and was never

built to answer. Feedback tells you that you are drifting but nothing about where you ought to be. Anyone who claims otherwise is smuggling in their preferred setpoint behind a diagram.

ORDER WITHOUT A SIGNAL

Free people produce better order than commanded people: that sounds like something a person believes, not something anybody has measured. But it has been measured, in a small way, on roads.

Hans Monderman was a Dutch traffic engineer who spent the last part of his career removing things. His argument was mechanical, not sentimental. He held that a traffic signal does not make a driver careful; it makes a driver obedient, and obedience is a poor substitute because a green light carries no information about what the other road users intend. The driver stops reading the situation and reads the light instead, which means that at the moment something unexpected happens, the person behind the wheel has been trained out of the only skill that would help. Strip out the signals, the curbs, the markings and the priority rules, and drivers are forced back onto negotiation: slowing, looking, catching an eye, working out who goes first with the person who's actually there.

In the early 2000s the Laweiplein junction in Drachten was rebuilt on that principle. Something like 20,000 vehicles a day pass through it at the meeting of four roads, with no signals and no formal priority. Traffic moves more slowly and, by most of the accounts published afterwards, it moves through the junction faster overall, because the delay signals impose is the waiting rather than the driving. Injury collisions fell. Monderman's own summary was that a road which treats people as incapable will get behavior to match.

The design has a victim.

Shared space removes the fixed, predictable crossing point that a blind or partially sighted pedestrian depends on entirely. Negotiation by eye contact is not available to somebody who cannot make eye

contact. A design that returns judgment to the median user takes it away, completely and by construction, from a minority who had it before, and it's the minority least able to absorb the loss. Campaigning by British disability organizations made this argument through the 2010s, and in July 2018 the Department for Transport asked local authorities to pause new level-surface shared space schemes while the guidance was reconsidered.[23]

So here is a design that gives most people more say and takes it away entirely from a few, built by a man whose whole purpose was to give people more say. Producing agreeable answers was never what the test was for.

THE BOUND, AT THE SIZE OF ONE PERSON

Widen the frame, more people, more consequences, more time, and see whether the benefit survives the widening. Anything that looks like a gain at only one boundary was never a gain; it was a boundary.

The man on the path raises his own capacity at the price of ending his neighbor's altogether. Widen the frame by one person and the account is already negative; there is no further widening required and no appeal to sentiment involved. Widen it to the neighbor's children and it stays negative for a generation, which is the time dimension doing its ordinary work. Nothing about the men's characters, their grievance, or which of them the reader would rather have a drink with needs to enter the ledger. All it has to include is the person on the other side of the door.

So agency, as I'm using the word, means the capacity to choose among options that don't run on removing somebody else's capacity to choose. Raw ability to do as one wishes does not qualify. A slaveholder possessed enormous ability to act, inside a structure that manufactured its power by destroying agency in others, and we have already seen what that structure failed to build. By this reckoning,

that isn't agency. It's extraction wearing agency's clothes, and the wider ledger returns the bill eventually, late, and to the wrong address.

It's the same test applied to markets and constitutions and empires, at the smallest scale it has.

FOUR PLACES THE BOUNDARY GIVES AN UNCOMFORTABLE ANSWER

Any test that returns a clean verdict on every case is a test somebody has trimmed.

Self-defense. Here the verdict is clear and it isn't close. A woman who disables an attacker has reduced his capacity to act, dramatically and perhaps permanently, and the ledger widened by exactly one person still comes out positive, because the capacity she reduced was at that moment being spent entirely on removing hers. She did not create the subtraction. She interrupted one already in progress.

What it doesn't handle is proportion. It tells her she may act and says nothing about where the line falls between stopping him and punishing him, between the blow that ends the attack and the third blow after he is down. Every legal system that has taken the question seriously spends its effort exactly there, on imminence, on the duty to retreat, on whether the force was proportionate to the threat, and the wider ledger contributes nothing to any of it. The test can tell you which direction the account runs and not when to stop adding to it.

Coercion to prevent a greater harm. Quarantine. Involuntary commitment. The conscription of a generation. Such acts can pass. That's the answer I trust least. Comparing magnitudes is how the test works, and in every one of these cases the party doing the comparing is the party doing the coercing. A public health authority weighs the liberty of the confined against the infections prevented, and the authority also writes both numbers. There is no mass detention in the historical record that was not justified in precisely this form, using precisely this arithmetic, sincerely, by people who had done the sum

and found it favorable. The wider ledger offers no protection against its own misuse, because misuse of it looks identical to use of it from the inside.

Whatever protection exists has to come from somewhere other than the arithmetic: from the counting being done by a body that is not the one doing the confining, from a fixed expiry that requires a fresh argument to renew, from a route of appeal that a confined person can actually reach. Those are procedural, unglamorous, and constantly eroded. I have no better account of it than that, and I'd rather leave the gap showing than paper it over with a principle.

A parent overriding a child. A four-year-old moving toward a road is picked up and carried back, and her capacity to choose is removed entirely, by force, without consultation. Passes, for a reason that generalizes badly.

She isn't trading the child's agency against somebody else's. She is trading the child's agency today against the same child's agency for the following eighty years, which makes this the one case where reducing a person's capacity and expanding it are the same act. One ledger, and the widening happens in time rather than across people.

Trouble is, the identical structure is the oldest justification ever offered for empire, for guardianship over the poor, and for every arrangement in which one group has decided that another is not yet ready. Same words. What separates them has to be something checkable from outside, and I think there is one thing: a child's incapacity is temporary, observable, and shrinking, and the arrangement is built to end. Every reasonable parent is running an expiring authority and knows the expiry date within a few years. Colonial administrations were not built to end, made no measurement of readiness they intended to act on, and produced no schedule. A guardianship with no expiry isn't a guardianship. So the question to put to any paternalism is whether it is building toward its

own dissolution, and whether anybody outside it can check the progress. Which is answerable, and the answer is usually available quickly.

The state's monopoly on force. Here the test returns its least satisfying answer.

The monopoly passes. Where one accountable body holds the legitimate use of force there is vastly less total coercion than where the capacity is spread among whoever can raise a band of men, and the historical record on that is not seriously in dispute. Widen the frame across a population and across a century and the account is strongly positive. Trouble is what happens next: the entity that passes the test is also the entity that defines the frame, appoints the auditors, and holds the records. A police force reduces the aggregate coercion in a society while holding an unreviewable share of what remains, and asking it to widen its own ledger is asking a party to audit itself.

The test supplies one usable handle, the fourth question of the wider ledger: whether the structure can continue without escalating force or deception. A state that requires more coercion every year to hold the position it held last year is publishing its accounts whether it means to or not, and the publication is legible from outside without anybody's cooperation. Slow instrument. Historically, a reliable one.

There is a fifth case, and it is easier. Competition destroys capacity. A better store opens across the street and a family business that stood for thirty years is finished. Nobody was defrauded, nothing was coerced, and a man's ability to act in the world has been cut down hard by someone else's choice. Does the bound forbid it?

It doesn't, and the reason is that what was destroyed was a position, not a person's capacity to choose. The proprietor retains his skills, his relationships, his ability to open somewhere else or do something different, and the ledger widened by one includes the customers who now pay less for the same goods. I believe it. I've also

closed a store, and I can tell you that to the man turning the key for the last time this analysis is a piece of paper. A book that pretends the ledger's verdict is also his verdict has stopped being about people. The ledger widens. It doesn't console.

WHAT CONSENT HAS TO BE MADE OF

Agency without grounding becomes manipulation with a choice architecture around it. A person cannot choose well among options built to deceive, and the appearance of choice is exactly what a well-built deception supplies. Recommendation systems narrow attention while calling it personalization. Propaganda manufactures the preference it then satisfies. Financial desperation turns consent into a legal fiction, and it does so without anybody lying about anything.

Contract law has been working on this problem for several centuries, empirically, one dispute at a time, and it has arrived at a set of conditions that map almost exactly onto what is needed here.

A signature isn't, by itself, consent. The law treats an agreement as unenforceable where it was obtained by duress, where one party lacked capacity, where the terms are unconscionable, or where the assent was procured by a misrepresentation the other side relied on. Each doctrine is really a statement about information and alternatives. Duress says the promise carries no information because refusal was not available. Incapacity says the promise carries no information because the promisor could not evaluate it. Unconscionability, in the standard American formulation, has two halves that have to be present together, a defect in the bargaining process and terms that are grossly one-sided, which is the law insisting that a bad deal isn't enough on its own, and that something must have been wrong with how the choosing happened.

One case fixed the doctrine in American law. Ora Lee Williams bought a stereo set from a Washington furniture company in 1962 for a little over five hundred dollars, on installments, while supporting seven children on about two hundred and eighteen dollars a month

in public assistance.[24] Buried in the printed form was a cross-collateral clause: every payment she made was pro-rated across every item she had ever bought from the store, so that until the last one was paid off, none of them were. She defaulted, and the company moved to repossess everything she had purchased since 1957. In 1965 the Court of Appeals for the District of Columbia, in an opinion by Judge J. Skelly Wright, sent the case back to be examined for unconscionability, and in doing so gave the modern doctrine its shape. What the court identified wasn't cruelty but a term that no reader of that contract could have priced, sold to a buyer with no realistic alternative source of credit.

The law can void a bargain but can't supply the alternative that would have made the bargain unnecessary. Ora Lee Williams got the clause struck down and didn't thereby get a better furniture store, a lower interest rate, or two hundred more dollars a month. And the doctrines are drawn far tighter than the moral case would draw them, deliberately. Economic duress almost never succeeds; courts require both an improper threat and the absence of any reasonable alternative, and hard bargaining between unequal parties is not enough, because a doctrine that voided every agreement made under pressure would void most of commerce and could not be administered. So the law stops well short of where this test would stop, and it stops there on purpose, for reasons of workability rather than principle.

I take that as a warning about my own instrument, not a complaint about the courts. Any test that finds more defective consent than an institution can act on is describing a problem, not remedying it, and the gap between the two is where most of the actual work lives.

Which is why the mathematics of agency has to include the material conditions and not only the rule. Access to evidence, and the ability to inspect where it came from. An alternative that genuinely exists, close enough to reach. Memory of how the last several choices

turned out, since a person with no record cannot learn from consequences and the definition requires her to. Education improves the ability to name a need. Health provides the physical capacity to act on it. Income and secure housing create a time horizon long enough to plan inside. Legal rights lower the price of dissent, which is the difference between having an opinion and being able to afford one.

And it includes the ground itself. A child should know the feel of soil and the scale of a tree, not only the image of each on a screen, and should have some idea where water goes and where food starts, at first hand, before being handed an interface to it. Technology can then do what it is good at: identify the tree, describe its habitat, name the insects that depend on it. The screen should point back to reality rather than away from it.

This is why the future of education is neither a retreat from technology nor surrender to it. It is grounded augmentation. Handwriting, unhurried reading, building something with tools, walking somewhere, tending a growing thing, sitting with a person who needs attention: these develop forms of attention that no application should be allowed to replace, and the case for them isn't nostalgia: attention of that kind is the substratum everything else in the definition runs on. What machines add is real and should be taken: translation across a language barrier, a memory that does not decay, a comparison run in a second that would have taken a library and a week, a simulation of a consequence before it is incurred, an answer to a question at the moment the question occurs. The combination raises agency further than either extreme, and the extremes are both well represented and both wrong.

FROM RECIPIENTS TO PARTICIPANTS

Most institutions were built to deliver something to people, and the delivery model is written into their architecture. The news is broadcast, the policy announced, the curriculum set, the diagnosis

given; the aid arrives. In each case the person on the receiving end is modeled as a destination, which is a reasonable design when information moves in one direction and expensively, and an obsolete one when it doesn't.

The next structure treats them as participants. A patient now arrives with three years of her own continuous data and a list of questions, which is a different consultation from the one medicine was organized around. Students can shape a path through material rather than following the single sequence a syllabus can afford to offer. Any citizen can now read the competing constitutional interpretations in the original rather than the summary of a summary. Tools that required a corporation's budget a decade ago sit on a laptop.

Participation does not end expertise; it ends expertise held as unchallengeable possession. An expert whose reasoning can be questioned, translated into plain language, tested against a case, and combined with what somebody local already knows is more valuable than one whose authority depends on nobody being able to check. Greater agency does not flatten competence. It makes competence accountable, and accountable competence is the only kind that compounds.

What the bound doesn't do is settle proportion, or protect against a party that counts its own ledger, or tell a paternalism when to end without an independent check, or console anybody whose capacity was reduced by a process the ledger scored as a gain. Some conflicts survive any widening of the frame, because two people's capacities genuinely cannot both be satisfied. Those do not get easier. What the bound does is separate them from the much larger category of harms that look irreducible only until somebody widens the ledger.

THE EQUATION

Grounded knowledge + meaningful choice + feedback + memory = individual agency. More agency = more honest vectors = more durable alignment. Capacity purchased by removing another's capacity does not enter the sum.

CHAPTER 7

Electrons Arranged as Energy, Knowledge, and Memory

The same medium is beginning to carry both abundance and intelligence

An electron arranged in a grid can keep a baby warm. Arranged in a computer, it can preserve a library. Arranged in a network, it can carry a question around the planet.

One physical thing does three jobs that civilization spent 5,000 years routing through separate systems, fire, paper, and voice, and it happened so recently that most of us living through it haven't noticed.

For most of history, energy and information moved by different means through different institutions. Energy came from the bodies of animals, from wood and coal, from water falling and wind blowing. Information came through voice, through marks pressed into clay and scratched onto vellum, through teachers and libraries and the slow postal systems that let ideas cross an ocean in six weeks if the ship didn't go down. Both energy and information were scarce. Both rewarded whoever controlled the source. And the systems that delivered them had nothing in common except that they depended on civilization existing at all.

Electricity changed that, and it took a century for the change to be fully visible. An electron is a charged particle that can be arranged to do either job, depending on the medium it moves through and the instructions given to the circuit. In a heating element it dissipates as heat. In a logic gate it encodes a bit. In a transmission line it carries

both to wherever the infrastructure reaches. The energy transition and the intelligence transition are not separate revolutions. They are different arrangements of the same underlying capacity.

* * *

A THURSDAY IN JULY, PHOENIX, ARIZONA

A single-family house sits on a street in the East Valley. By 7 am the air-conditioning has been running for two hours, drawing about 3.5 kilowatts, sustained and invisible behind a thermostat. In the garage a plug-in vehicle sits at the end of an overnight charge; its 75-kilowatt-hour battery holds about as much energy as eight gallons of gasoline, except that it arrived silently over a cable from a grid fed that night mostly by wind turbines turning in the dark outside Tucson, eighty miles south.

By noon the solar panels on the roof are producing more electricity than the house is using. The surplus flows out onto the local grid, a trickle nobody is watching, but similar trickles are leaving several hundred thousand rooftops across the state at once. On the grid operator's dashboard, net demand drops, and the wholesale price for electricity in the midday hour falls below two cents per kilowatt-hour.

At 4:45 pm, something changes. The sun is still high enough that panels are producing, but a wave of demand is approaching: air-conditioners cycling back on as people return from work, electric vehicles beginning their evening charge, commercial buildings drawing their peak load of the day. Somewhere a grid operator has fifteen minutes to find eight hundred additional megawatts before the price spikes and, in the worst case, before reliability is threatened.

What happens next is invisible to the homeowner on that street, and it involves her house in a negotiation she never knowingly entered. Her smart thermostat, enrolled in a demand-response

program by the utility, allows the grid to raise the house temperature setpoint by two degrees for forty-five minutes. Her vehicle, connected to a bidirectional charger and set to charge automatically, receives a signal to pause. Together with several hundred thousand other enrolled devices across the state, these adjustments stand in for a medium-sized power plant, not by building anything new, but by timing what already exists.

The electrons don't know any of this. What they do is follow the path of least resistance through a grid that has been arranged, in real time, by software running on behalf of people who set their preferences once and then forgot about it. Orchestration, at the level of a single afternoon, means this: the instrument is the electron, the music is the schedule, the conductor is the algorithm, and the composer, somewhere upstream, is the policy about whose demand gets shifted and by how much and in exchange for what. That last question isn't a technical one, and nobody on that street was asked it.

* * *

THE SHAPE OF THE PROBLEM: THE DUCK AND THE SPIKE

California's grid operators have a graph they call the duck curve. It just looks like one. Plot net electricity demand across a spring day in California, and the line falls through the morning as solar generation rises, reaches a trough in the early afternoon when panels are producing abundantly, then climbs sharply as the sun drops and demand does not. Charted, the profile looks like a duck: a curved body in the middle of the day and a steep neck rising in the evening.

Its neck is the problem. In the space of two or three hours, the grid must bring online the equivalent of several large power plants' worth of capacity to fill the gap left as solar generation falls. Before storage and demand response at scale, this meant keeping fast-ramping gas

turbines idling through the day, ready to respond within minutes but running badly: burning fuel to be ready to burn fuel. Solar abundance created its own constraint.

A barrel of oil carries stored energy and must be replaced after use. A solar panel is a conversion device that continues receiving fuel from the sky. A battery moves energy through time. A transmission line moves it through space. Software moves demand. An electric vehicle can be transportation, storage, and a flexible grid asset. The value lies increasingly in orchestration. I want a BYD, and not because I like Chinese culture better than my own. I think it's a better value: better engineering for less money, and one of the better batteries on wheels.

Storage changes the mathematics. Charged during the midday trough, when solar is abundant and prices are low, it discharges into the evening ramp, flattening the neck of the duck. On the grid, a large storage installation does this for hundreds of thousands of customers at once; in a garage, a car battery wired both ways does a smaller version for a single street. Everything economic about it depends on the price difference between the charging hour and the discharging hour, and those prices, in a well-designed market, reflect what electrons really cost at different times rather than a flat average that hides it. Buying the cheap hour and selling back in the dear one is called time-of-use arbitrage, and it turns out to be a kind of patriotism: it keeps the neighbors' beer cold and the grid steady.

A conductor can tell the strings to hold while the brass comes in, and neither section is damaged by the instruction. A battery that charges and discharges through many deep cycles degrades: its chemistry changes, its capacity diminishes, and at some point it must be replaced. Electrons are patient. The thing that holds them isn't. In a grid with plenty of storage, the design question is partly engineering, which chemistries last longest at which duty cycles, and partly social: whose batteries take the wear, in exchange for what,

and who pays for the replacement? Software will optimize for the grid. Whoever owns the car is doing a different sum, and without a fair deal she'll simply unplug.

This is why batteries still matter, even in a future of abundant generation. Overbuilding solar and wind can make electricity extraordinarily cheap during many hours, but clouds, seasons, geography, and evening demand remain. The better system is not one miraculous device. It is a network that makes more use of assets already present. Curtailment, clean electricity available when the grid cannot use it, is usually treated as waste. It can also be read as an invitation: move computing, water heating, industrial processes, vehicle charging, or storage into the abundant hours, and turn surplus electrons into future mobility, heat, cooling, materials, or data. The transition is from extracting scarce fuel to scheduling plentiful flow.

* * *

A SECOND ABUNDANCE

Knowledge has the same shape as energy: scarce for most of history, controlled at the source, delivered through institutions that decided what arrived and who got it.

A printed book is not consumed when read, but for most of its history reaching one meant being near it, being let in, and finding your way through a system nobody had designed to be found through. A university library in 1970 held perhaps a million volumes. A researcher with a question that touched three disciplines had to know which building held which catalog, which journals were on paper and which only on microfilm, and which reference librarians had memorized connections no index contained. The knowledge was there. Reaching it was a second profession on top of the first.

Digital transmission made copying cost nearly nothing. A paper published in Geneva can be read in Nairobi within seconds of its release, by anyone with a connection, at no marginal cost to either party. Lectures from universities that once required admission and tuition are available in a browser, along with papers, demonstrations, and courses that would have been unreachable to most people in my childhood. Medical literature that doctors used to receive months after publication is now searchable the same day.

Access is not understanding. When information is everywhere, what runs short is attention you can trust. A person can reach a million documents on any contested question and still be unable to determine which interpretation is grounded, which study has been replicated, which institution has an interest in the answer it is publishing, or what matters to the decision in front of them. That is why memory and cross-correlation become the next infrastructure.

Here the electron comes back. At its most basic, an AI system is a very large arrangement of electrons performing arithmetic on representations of language at extraordinary speed. What it can do is hold more of the existing record in working relation than any single person can, and apply the comparisons that relation makes possible: this claim appears in forty-three sources, twenty of which share a funder; this conclusion was revised in a paper three years later; this study's methodology has been questioned on these specific grounds. None of it is understood by the electron, which is doing what it always does, following instructions through a circuit. But the circuit, at sufficient scale and with sufficient structure, begins to produce something that functions like memory: a record that persists, that can be queried, that connects a question asked today to an answer recorded before the person asking was born.

A grid manages the flow of energy among sources, storage, and loads, smoothing the mismatch between when power is generated and when it is needed, and a well-designed AI interface manages the flow of information among sources, memory, and questions in exactly

the same way. Both require infrastructure to be useful. Both require someone to decide whose needs get prioritized when the demand exceeds what the system can serve simultaneously. A grid without a pricing signal to allocate scarce electrons in peak hours becomes unreliable; a knowledge system without a truth structure to allocate trust among competing claims becomes noise. Both problems have the same shape because they're running on the same substrate, and the solutions, storage, routing, prioritization, and a way for the user to see what is being delivered and why, are the same solutions. Energy needs a grid. Knowledge needs a truth and value structure.

A grid can be physically checked: the electrons either arrive at the socket or they do not, and the meter is not subject to interpretation. An AI's answer arrives in confident prose regardless of whether it is accurate, and the user has no meter to put at the socket. The checking has to be built into the system itself, as part of the structure, because what it is supposed to deliver is exactly the ability to tell a reliable answer from a plausible one. A grid that delivers the wrong voltage damages the appliance. A knowledge system that delivers the wrong answer with fluency damages something harder to repair.

* * *

THE ELECTRON BANK

Namoruputh Health Centre, in Turkana County in northern Kenya, sits in a region where the national grid is connected but outages caused by heatwaves and flash flooding are routine. When a modular solar plant with battery storage was installed there in 2022, it was set to take over the moment grid power drops, so the pharmacy refrigerators stay between two and eight degrees, where vaccines have to be kept.[25] The refrigerators come before everything else in the building. Vaccines need cold that never stops, and getting the

power back an hour later doesn't save a dose that has already warmed. We have the technology to keep medicine cold, and keeping people well means using it.

That clinic isn't unusual anymore. What's unusual is saying plainly what it is. Money is stored permission to claim future value. Electricity is immediate physical capacity. A stored kilowatt-hour at the right moment may have more practical value than a speculative token, because it can perform work: keep medicine cold, run a factory, or stabilize a neighborhood. The phrase electron bank is not a proposal to replace all currency. It is a reminder that the energy system is becoming an asset network, in which homes, vehicles, buildings, batteries, and appliances can store or shift value, and in which the difference between a surgical light staying on and going out is a battery somebody owns.

Follow that in the direction it wants to go. Describing the energy transition as a shift from fossil fuels to renewable generation is accurate and incomplete. What is also changing is who holds the power. Under the old system, energy resided in deposits: coal seams, oil fields, gas reserves. Whoever controlled the deposits controlled the energy, and a household, a clinic, a small business had no reserves of their own; they were purely consumers, dependent on the supply chain from extraction to delivery holding together. Storage changes this. A solar panel on a roof and a battery in a garage are a small but real reserve that belongs to whoever lives there, a position in the energy system instead of plain dependence on it. Coordinated at scale, those positions are something genuinely new: an energy system whose value is spread among millions of participants instead of drawn from deposits held by a few. A car that sells its spare electrons back at the evening peak makes its owner a small factor in the price, where before she was only the one paying it. Put a billion of those cars on the road and together they can push more power into the grid than most nations generate. Oneness stops being a philosophy and becomes an engineering fact. Whether that lowers

the price for everyone, the neighbor without a car included, is a fair question, and the answer depends on who pays for the wires. Panels and batteries are capital assets, and capital assets take capital to acquire. A household that can't afford the upfront cost takes part in the energy transition as a customer, not an owner, and gets the benefit as lower prices rather than as something it owns. Ownership is spreading faster than in any earlier energy transition, and unevenly, and the choices made now about policy and markets will decide whether it spreads fairness or repeats the old unfairness. Coordinated abundance can reward millions of participants rather than a few owners of deposits, but only if the participants can afford to hold a position in it.

* * *

TWO ABUNDANCES

For the first time in the long history of civilized life, the same medium delivers both.

Twice before, the machinery of energy and the machinery of information grew up side by side and fed each other. The first was the printing press and the water-powered mill, both from the same mechanical tradition, both spreading along the same trade routes in the fifteenth and sixteenth centuries, but serving different masters and never sharing a supply chain. The second was the steam engine and the telegraph, which arrived within a generation of each other in the nineteenth century and together made the industrial city possible, the engine concentrating labor and the telegraph coordinating it across distances the engine had suddenly made irrelevant. But steam and signal still ran on different fuels through different wires. The electron, alone among energy carriers, can do both jobs in the same medium, through the same infrastructure, at the same time.

The electrons that keep the clinic's vaccines cold tonight will flow through a different circuit tomorrow and carry the query that a nurse enters into a tablet connected to a medical AI, which will compare her patient's symptoms against a database assembled from a hundred countries' diagnostic records, and return an answer that the clinic's single physician will evaluate against what she knows from standing in that room with that person. Nothing separates the electron carrying the energy in kind from the electron carrying the information. They move through different architectures, obey different protocols, and serve different purposes, but they are the same particle, obeying the same physics, drawn from the same abundance.

Nobody designed this convergence. What produced it was a century of separate changes, factories wired for electricity, messages turned into digits, computers shrunk to fit a pocket, that happened to run on the same thing. That one buildout can now serve both transitions at once is an opportunity, and it won't organize itself.

An electron arranged correctly can keep a baby warm, preserve a library, carry a question around the planet. What arranges it correctly is the structure of human choices about who gets to participate in the abundance, and on what terms, and with what obligations to the people downstream of the decision. The electrons are ready. That part has been solved.

THE EQUATION

Renewable electrons + storage + networks + intelligent timing = coordinated abundance. Digital knowledge + memory + truth structure = cognitive abundance. Both abundances run on the same medium. The constraint now is not the medium. It is the arrangement.

CHAPTER 8

Truth Is Infrastructure

A civilization that cannot distinguish signal from confidence cannot steer

"Facts are stubborn things; and whatever may be our wishes, our inclinations, or the dictates of our passion, they cannot alter the state of facts and evidence."

John Adams, closing argument for the defense, Boston, 4 December 1770

Here is a claim you may have met: taking the lead out of gasoline caused the great American crime decline of the 1990s.

Is that true? We have machinery for true and machinery for false, and an enormous amount of machinery for whose side are you on. For a verdict that is broadly right in its core, overstated in its headline number, unsettled in its mechanism, and aimed at a situation that has since changed, which is the verdict most real claims deserve, there's almost nothing.

Dirty information is like dirty water. A person may survive one glass. A civilization cannot build health on a contaminated supply.

Truth gets discussed as a moral virtue or a philosophical puzzle, and it's both, and it's also infrastructure in the dull sense that a water main is infrastructure. Every bridge, medicine, contract, election, market, forecast, and personal relationship depends on claims corresponding well enough to reality that action can succeed. When the correspondence fails, the failure rarely announces itself as a philosophical error. A footing settles. A drug gets approved and then withdrawn. A family spends on a treatment what it needed for rent.

The internet expanded access to information faster than it expanded the ability to evaluate information, which was a predictable stage, not a surprise. Printing created pamphlet wars before it created libraries, and broadcasting created propaganda at scale before it settled into the evening news. Social networks then lowered publishing costs to nearly zero and optimized distribution for reaction, because reaction was the thing they could measure. Artificial intelligence has now lowered the cost of producing persuasive text, images, audio and video to roughly the cost of asking for it.

Whatever the answer is, it can't be a central authority declaring one official truth, because authority can be wrong, captured, slow, or self-protective, and usually gets there in that order. Nor can the answer be that every claim is merely perspective. Reality imposes consequences. The bridge falls down on the people who were right about it and the people who were wrong about it alike. The better next is a distributed truth infrastructure that makes provenance, corroboration, uncertainty, contradiction, and correction visible to the individual.

TWO TRADES THAT ALREADY DO THIS

Here's a question nobody asks at the counter. When the butcher's scale reads one kilogram, why should you believe it?

Not out of any confidence in the butcher. That scale was checked against a working standard held by the company that services it, which was checked against a reference standard held by a calibration laboratory, which was checked against a standard held by the national measurement institute, which is tied to the international definition of the kilogram itself. Every link in that chain is written down: who compared what against what, on which date, with which instrument, to what tolerance. It is the temple clerk's receipt, kept for a kilogram. Metrologists call the property traceability, and the word carries a precise meaning: an unbroken chain of comparisons, each

with its own stated uncertainty, terminating in a definition rather than in a person. Two features of that arrangement are worth stealing. First, the chain is auditable, not authoritative. Nobody inside it is asking to be believed. Each link produces a certificate, and any link can be checked afterwards by somebody who was not present when it was made.

Second, and this has never made it into public argument at all: no accredited laboratory reports a measurement without reporting an uncertainty alongside it. Bare, 1.0002 kilograms is not a result; 1.0002 kilograms plus or minus a tenth of a gram is a result. Uncertainty belongs to the number instead of apologizing for it, and a measurement offered without one gets treated as unusable. Not confident, not bold. Unusable.

Metrology also knows what happens when a chain terminates in an object instead of a definition. For 130 years the kilogram was a cylinder of platinum and iridium kept under three glass bells in a vault outside Paris, with several dozen official copies distributed to national laboratories. Whenever the copies came home, every few decades, they were compared against the original. Over the twentieth century they drifted apart, by something like fifty micrograms, a fingerprint's worth of mass. Nobody could establish which side had moved, because the cylinder in the vault was the definition, and a definition cannot be wrong. It can only become inconvenient. In 2019 the world's measurement bodies gave up on the object and redefined the kilogram in terms of the Planck constant, a value fixed by agreement and realizable in any sufficiently equipped laboratory on earth.[26] They replaced an authority with a procedure, and the procedure doesn't drift.

Mass has a physical referent, and a laboratory can repeat the measurement on a Tuesday and get the same answer. Most contested public claims have nothing of the kind underneath them; whether the removal of lead from gasoline reduced assaults in Chicago is not a quantity anybody can put on a balance, and no amount of

institutional design will make it one. Such a chain is also expensive, and it exists only because commerce, medicine and engineering pay for it, through funded institutions, across a century. Nobody is funding a traceability chain for claims about a school board.

Provenance can be recorded. Uncertainty can be stated as part of the result instead of buried in a caveat. Both can be made cheaply enough to sit beside an ordinary claim rather than only inside a laboratory report. Far less than metrology has, and a great deal more than we currently do.

A single source is a vector. It has direction, strength, context, and possible interest.

When several independent sources point toward the same conclusion, confidence in it rises. A source that conflicts with the others may be revealing an error, a change in conditions, or a distinction that nobody had thought to draw. What matters is the structure of the support rather than the number of links, and confusing the two is where most public reasoning goes wrong. A claim can be evaluated along several axes at once: how good the sources are, whether they are genuinely independent of one another, how recent and how direct the observation was, what each source stood to gain, and whether the whole assembly squares with something a person could go and look at.

In a working trade this method has an older name. A surveyor standing at an unknown point fixes their position by resection: sights three landmarks whose coordinates are already established, takes the bearing to each, and draws the three lines back across the map. If everything were perfect, those lines would meet at a single point, and that point is where you are standing.

They never do meet. The three lines enclose a small triangle, surveyors call it the triangle of error and navigators the cocked hat, and the size of that triangle is the most useful thing on the sheet. A

large one says the bearings disagree and the fix should not be trusted with anything valuable. A tight one says they agree, and roughly how well.

Now take away the third landmark. Two lines always cross, exactly, at precisely one point, however wrong either bearing happens to be. Perfect fix, apparently. It arrives with no triangle, which is to say with no warning, and the surveyor has no way to learn from the instrument that anything has gone wrong at all. A third source earns its place by making error visible, which is a different service from making the surveyor more confident and is frequently the opposite of it.

Resection has a failure mode of its own, and no amount of care at the instrument will catch it. Resection assumes the three landmarks were established independently of one another. Where two of them were fixed from the same control monument, and that monument was in the wrong place, both bearings carry the same error in the same direction, the triangle closes tight, and it closes tight around a spot where the surveyor is not standing. Precision arrives; accuracy does not. The tighter the triangle, the more confident the wrong answer, which is why a traverse is closed back onto its starting point and the misclosure measured. The check has to come from outside the network being checked.

THE CLAIM, RUN THROUGH

Back to the gasoline.

What holds. In 1953 a Caltech geochemist named Clair Patterson worked out the age of the earth — about 4.55 billion years — by measuring lead isotopes in fragments of a meteorite.[27] He only got there after discovering that his laboratory, his glassware, the air of Pasadena and his own hands were so saturated with industrial lead that his samples were useless, and after building one of the first ultra-clean laboratories to get around it. Contamination that had nearly ruined his measurement then became his subject for the rest of his

working life. Greenland ice cores showed atmospheric lead deposition climbing sharply from the 1920s, which is when tetraethyl lead began going into motor fuel. Industry-funded researchers disputed him for two decades and he lost contracts and committee seats over it. He was right.

In its core, the claim stands about as solid as environmental science gets. Lead is a neurotoxin with no established safe level of childhood exposure, and it damages exactly the faculties, impulse control, attention, the executive functions, whose failure turns up in crime statistics twenty years later. Between the late 1970s and the early 1990s, as the additive came out of the fuel, American children's mean blood-lead levels fell by roughly four-fifths. Neither of those findings is in serious dispute among people who have looked at them.

The number. Then there is the figure that travels alongside the claim. A 2007 economics paper estimated that the phase-out accounted for something like fifty-six percent of the fall in violent crime during the 1990s, and that number escaped into general circulation, where it has been repeated ever since as though it were a measurement.[28] One specification of one model on one country's data produced it. Ask what happened when other people tried. A 2022 meta-analysis of roughly two dozen lead-and-crime studies found strong evidence of publication bias, the pattern in which findings too small to be interesting never reach print, and reported that correcting for it shrank the average estimated effect to a fraction of the headline claim. Not nothing, though: on their reckoning, lead abatement might still account for somewhere between seven and twenty-eight percent of the fall in homicide. The direction of the number survives. The famous magnitude does not.

So the headline figure is carrying far more weight than the evidence beneath it can bear, while the thing it points at is probably real.

The cause. American violent crime fell by roughly half between 1991 and the middle of the 2010s.[29] Crime fell across most of the wealthy world in the same period, on different lead timetables and under very different policing regimes, which is the first thing that should give anybody pause. At least eight explanations are seriously in the field: mass incarceration, changes in policing method, the collapse of the crack cocaine market, an aging population, legalized abortion, economic growth, better vehicle and household security, and lead. Several have good evidence behind them. None can account for the whole decline on its own, and any three of them stacked together over-explain it, which is a reliable sign that something is wrong with how the question has been posed.

On causation, then, lead is one contributing vector among several, that its share is genuinely unknown, and that anybody quoting a percentage is quoting a model rather than reporting a finding.

And a fourth thing, which public argument almost never reaches. What the claim describes is a world that has since been altered. American children's blood-lead levels today are a small fraction of what they were in 1976. Leaded petrol ended worldwide in 2021, when the last country still selling it ran out of stock. Whatever the coefficient was in the 1990s, it cannot be used to argue about the value of the next dollar spent on remediation, because the exposure it measured has largely been removed already. What's left, service lines, old paint, contaminated soil, piston-aircraft fuel, is a different problem with different numbers, and the argument for fixing it doesn't need the crime statistics at all. It's poisoning children now. So, is it true?

Mostly, in its core. Overstated in its headline figure. Unresolved in its causal share. And aimed at a condition that has largely passed. Given an afternoon inside the literature, anybody can reach that verdict. In the ten seconds during which the claim is actually encountered, almost nobody can.

Nobody wins: not the environmentalist who wanted lead to explain the crime decline, and not the skeptic who wanted the whole hypothesis to be an academic fashion. Both walk away with less than they came for. A conclusion that disappoints everybody who cares about the question has no constituency, no distribution, and no reason for anyone to repeat it. That's the operating problem with truth infrastructure: the correct output of the process is the least shareable thing the process can produce.

The result should not always be yes or no. People still need an answer, which is why the output cannot simply be a shrug. A truth tool should give a clear bottom line while preserving the path behind it: what is high-confidence, what is plausible but unconfirmed, what is disputed, what evidence is missing, and what the better next question would be. Refusing false certainty is not evasion when the uncertainty is explained.

WHY TRUTH BECOMES MORE PROFITABLE

Every one of these methods rests on a single assumption, and it is the assumption now under direct attack.

Corroboration means something only when the corroborating sources are independent. Two people who watched the same fire from opposite streets are two witnesses. Two people who read the same wire report are one witness counted twice, and the count is the only quantity that has actually increased.

People have always been poor at telling those two situations apart, and networks make it worse in a specific, mechanical way. Six outlets carry a story that originated with one agency. Forty accounts repeat a figure that came from one post that misread one table. A survey measures what people believe their neighbors think, while the neighbors, asked separately, report back what they believe the first group thinks. None of this requires anybody to lie. All it takes is that repeating cost less than checking, which it always does, everywhere, for everyone.

Old mechanism. What has collapsed is its price. Manufacturing an additional apparently-independent source used to demand a person, a publication, a plausible history, and time. It now demands a prompt. So the one input every verification method depends on being expensive, independent origination, has become the cheapest thing in the system to counterfeit.

I don't have a solution to this. Partial measures exist, and they are worth taking. Weight for independence instead of counting agreement. Of any two sources that agree, ask whether they had access to one another before they agreed, and treat unanimity as a warning rather than as a result. A source that had to do work, take a measurement, visit a place, file a document, is worth more than one that only had to hold an opinion. Each of those helps. None survives an adversary who has read the specification, and the specification has to be public for the test to be worth trusting in the first place. I can't see the way through it.

Attention markets have rewarded outrage because reaction is easy to measure. Trust is harder to measure but more valuable over time. A service that repeatedly helps a user avoid fraud, correct a mistaken belief, understand a medical claim, or see manipulation can become utility-like infrastructure. The business model shifts from selling attention to selling decision quality.

That shift will not happen automatically. Platforms may remain more profitable when users are agitated. Governments may prefer narratives that preserve power. Individuals often prefer confirming stories. Truth infrastructure must therefore be personally accessible and structurally independent. It should let a user compare models and sources, inspect the value assumptions behind an answer, and retain control of personal context.

THE LAYER NEXT TO THE WRITING

For what a truth layer would feel like, the nearest working model I know of is already sitting under your cursor.

Because it costs so little that nobody rations it, spell-check runs on everything; no writer has ever decided that a document was important enough to spell-check. It flags and doesn't block: the red line stops nothing, and the writer who wants to keep the word keeps it, which leaves the decision where it belongs. Spell-check also ships with a way of being told it was wrong, and "add to dictionary" is a more radical piece of design than it appears, because the authority defers permanently to the user on a case the user knows better. Best of all, it changed behavior without persuading anybody of anything. Nobody ever resolved to become a better speller.

There's the specification, sitting in plain sight for thirty years: a cross-correlated truth layer beside ordinary browsing, the way a language layer sits beside ordinary writing: ambient, cheap, correctable, advisory. Experiments in that direction are worth running. IsItTrue.xyz is one of them, and I'd rather they were run by people who expect to be corrected than by people who expect to be believed. An oracle is the failure state of this idea, not its grown-up form.

Work on IsItTrue.xyz has pushed the spell-check analogy one step further. Whatever else changes, the simplest user question should stay exactly what it is: Is it true? But the useful answer is rarely a binary stamp. First a person needs a short reliability call, high, mixed, low, or not enough evidence, followed by the reasons it earned that level.

Under that first signal, the tool can show where independent sources align, where they diverge, which assumptions differ, what evidence is direct rather than repeated, and which part of the disagreement is factual rather than definitional. Alignment in this sense is descriptive, not moral. Two sources can align around a harmful objective; a minority source can be right. What's valuable is seeing how the evidence is arranged, not being told that the majority is right.

Persistence changes the architecture again. One check is a verdict on one claim. Repeated checks can become a memory of relationships: which sources repeatedly depend on the same origin, which disagreements recur, which claims change as new evidence arrives, which model is confident outside the area where its grounding is strong. Those patterns are not automatically true either. They are hypotheses produced by the history of checking and offered back for another check.

In that direction a three-model council becomes useful without becoming mystical. Agreement among models is not truth; shared training data can produce shared error. Disagreement among them is not truth either. What several models can do is show, cheaply, where they frame the question differently, pick different evidence or draw different conclusions, so the reader knows where to look hard.

A truth layer should not autocorrect belief. It should make the reasons for confidence and doubt easier to inspect.

English has a dictionary. There exists a finite, enumerable, maintained list of correctly spelled words, and any given string is on it or is not on it. Truth has no such list, nobody is going to compile the set of true sentences, and the claims that matter are precisely the ones whose status is unsettled or contested or dependent on a definition that nobody has agreed. When a spell checker with an incomplete dictionary underlines an unusual surname, the writer shrugs and moves on. When a truth layer with an incomplete dictionary underlines a true claim, it has manufactured doubt where none was warranted, and to anybody in the business of manufacturing doubt, that's worth real money.

Then there's autocorrect. Spell-check's more confident descendant doesn't flag the word but changes it, silently, while the writer is looking somewhere else, and people have been sending messages

they didn't write for fifteen years. One product decision, made by somebody optimizing for engagement, separates a truth layer from an autocorrect for beliefs, and nothing in the technology stands in the way of that decision being made. A system that tells you what is true, that you have stopped checking, and that answers to somebody else would be the most efficient instrument of control ever built. From the inside it would feel exactly like being well informed.

Which is why structural independence matters more than accuracy, and why it is the property that gets designed out first. Accuracy can be advertised. Independence can only be paid for, over and over, out of money somebody would rather book as profit.

WHEN VERIFICATION BECOMES PEDESTRIAN

Of everything our small lab is trying, I keep coming back to IsItTrue.xyz, not because it knows what is true but because the ambition is to make checking ordinary. Pedestrian. Spell-check changed writing partly because checking a word stopped requiring a dictionary and became something quietly available beside the sentence. Why shouldn't evidence work the same way? A claim appears, and beside it a person can see where it came from, what supports it, what contradicts it, how independent the sources are, what is missing, and what to ask next. Not an authority telling you what to believe. A better-lit stream of information.

Something like it is already happening in a narrower trade, on the floor above my shoe store. Valid8 Financial now has the second-floor space NVIDIA once used. Its software helps forensic accountants rebuild a trail of transactions back to the original records, work that used to take trained people a long time before the real investigation could even start. The accountant still decides what the evidence means. The software makes it easier to see. That is what I mean by pedestrianization: not making expertise worthless, but making some of its tools ordinary enough that many more people can use them.

The public stream of information needs the same change. Anyone can now receive a beautifully written falsehood from anywhere on earth in a fraction of a second. Sending a claim around the world costs almost nothing. Checking it still costs an afternoon. If checking ever becomes nearly as easy as receiving, lies and mistakes won't disappear, and neither will propaganda, but their head start begins to shrink. That could be one of the very big NEXTs.

WHERE THIS STANDS

Facts alone do not choose goals.

A weapon's range can be measured to the meter. Whether building it increases anybody's security is a differently structured question, involving what the adversary does in response, what the money would otherwise have bought, whether deterrence holds in this particular case, and whose fear is admitted as evidence. Every one of those is arguable. None of them is settled by the range. A policy can reduce one price while raising a cost that sits outside the frame in which the price was measured, and the measurement will be perfectly correct.

So a system can be accurate throughout and still leave its user worse off, because a true answer to a badly framed question is a fully functional instrument of error. From inside the system that failure is invisible, since every component reports success.

So a truth layer has to interrogate the question and not only the claim. Having answered what appears to have been asked, the useful next move is to ask what the person is actually trying to accomplish, which assumptions are riding inside the phrasing, who is affected who has not been mentioned, and what evidence would change the conclusion. The highest-utility answer is sometimes a better question.

And the objection to that arrives immediately, as it should: a system that reframes the question has taken something from you. Somebody chose which reframings it offers.

One asymmetry sits under all of this. Producing a persuasive false claim now costs a few seconds of compute. Checking one costs a person's attention, which hasn't got cheaper, won't get cheaper, and is the one input in this whole book whose supply is fixed. An afternoon of the work behind these pages covered a single claim. An ordinary person meets several dozen before lunch. For twenty years that ratio has been getting worse, and every technology arriving now makes it worse faster, because generation scales and attention does not.

At this moment, truth is losing. Everything described here, provenance recorded, uncertainty stated, independence weighted, the decision left with the person, is a specification, not a thing that exists. Fragments of it exist, run by people with mixed motives and short funding. Nobody is paying for the chain.

I'd like to end with a sentence about how the specification eventually gets built. I don't have the evidence for it. What I have instead is a calibration chain that took a century and a statutory budget, a surveyor's third landmark that helps only if it was established independently, and the observation that the correct answer to a contested question is the least shareable thing anybody can produce.

A civilization cannot coordinate on a reality it cannot reliably see. The instrument is not being built. And the first faculty to go, in a population that cannot see, is the faculty for noticing that it cannot.

* * *

CHAPTER 9

The Old Mathematics of Power

Why protection, fear, and control can become self-reinforcing systems

After he became President of the United States following his heroism in World War II, former General Dwight D. Eisenhower understood military necessity better than almost any other president. After that war, he came to conclusions voiced in his 1953 Cross of Iron speech, where he said that every gun, warship, and rocket ultimately represented resources not used to feed, clothe, educate, or shelter people. And later, when he left office, he warned in his 1961 farewell address that a permanent armaments industry and a vast military establishment created a new concentration of influence.[30]

He knew about weapons because he had used them. He had spent three decades in the military. Sometimes weapons were necessary.

However, the danger begins when a device defines the problem. More weapons create perceived threat, which justifies more weapons.

THE ARITHMETIC HE DID IN PUBLIC

Shortly after Soviet leader Joseph Stalin died on 5 March 1953, six weeks into Eisenhower's first term as President, Eisenhower spoke to the American Society of Newspaper Editors in Washington, categorizing the defense budget aloud as a list of things it would prohibit from being built:

A heavy bomber cost more than thirty brick schoolhouses; two electric power plants, each serving a town of 60,000; two fully equipped hospitals; or fifty miles of concrete highway. A single fighter plane cost half a million bushels of wheat. A single destroyer cost new homes for more than 8,000 people.

Then he uttered the words that are still remembered: that every gun, warship, and rocket made signifies "a theft from those who hunger and are not fed." And the image that closed it: "Under the cloud of threatening war, it is humanity hanging from a cross of iron."

What makes that speech unusual is its accounting. Eisenhower compared military spending to the alternative foregone, and he did it item by item so that it could be checked by anyone who cared to.

But what followed took things in a different direction. Four months later, in August 1953, the Soviet Union tested a massive bomb, and an even bigger one in 1955.

Consequently, during Eisenhower's two terms the American nuclear stockpile grew from roughly 1,500 warheads to more than 20,000, despite the Cross of Iron talk, one of the most admired speeches in American political history, also known as the "Chance for Peace" speech.

The arms race is not simply caused by evil people. It is an old mathematics of uncertainty.

WHY THE ARITHMETIC LOST

Weapons persist because trust remains local. Nations can't assume that rivals share their values or will honor an agreement, verification is genuinely hard, and the consequences of being wrong aren't symmetric. Underestimating a threat can be catastrophic and permanent. Overspending on defense is embarrassing, expensive, and politically survivable. If you are the official who guesses low and turns out to be wrong, nothing else you ever did will be remembered; if you guess high and are wrong, you retire on schedule. Every individual actor is behaving sensibly, and the aggregate is a machine nobody chose.

Eisenhower lived inside that asymmetry more painfully than anyone. From 1956 his U-2 aircraft were photographing the Soviet interior, and the photographs told him there was no missile gap, that the American arsenal was overwhelming. But he couldn't say so. Saying so would have revealed the source, and the source was the only thing verifying his own estimates. So the general who had commanded the largest military operation in history stood silent while a senator and then a presidential candidate accused him of leaving the country defenseless, and he absorbed it, and the budgets went up anyway.

Eighteen months before Eisenhower left office, in July 1959, an American exhibition opened in Moscow's Sokolniki Park with a model kitchen in it, and Richard Nixon and Nikita Khrushchev spent an afternoon arguing in front of the appliances. Nixon's case was that competing over washing machines was preferable to competing over rockets. Khrushchev didn't really dispute it. He claimed the Soviet Union had washing machines, too, and the two men went on being pleasant to each other.

It made no difference to anything. Both governments agreed in public that the kitchen was the real prize and went home and funded the rockets, which tells you how much a good argument weighs against an incentive.

WHAT HE SAID ON THE WAY OUT

On 17 January 1961, five days from the end of his presidency, Eisenhower spoke to the country on television for about a quarter of an hour. He was seventy. He had nothing further to gain and nothing left to protect.

He gave numbers first, as he often did. Three and a half million men and women were then directly employed in the defense establishment. The United States was spending more each year on military security than the net income of all American corporations combined. This was new, he said. No such permanent armaments

industry had existed before, and the country had improvised it out of necessity within his own lifetime. Then the phrase: "in the councils of government, guard against the acquisition of unwarranted influence, whether sought or unsought, by the military-industrial complex."

The line survived, and I think its survival did some damage, because a phrase that portable can be aimed at almost anybody by almost anybody. Most of the rest of his speech didn't survive, and the rest is where the trouble is. He warned in the same breath that public policy might become the captive of a scientific and technological elite, that a solitary inventor tinkering in a shop had been displaced by task forces and federal contracts, and that free universities could find the substance of their research shaped by whoever wrote the checks. He warned, in a passage almost nobody quotes, against plundering the resources of tomorrow for the ease and convenience of today.

Eisenhower saw the military-industrial complex and where it would take us.

He thanked the country and left. He had eight years of unmatched authority, the credibility of the uniform, and the largest audience available to any human being alive. He used all of it to describe a mechanism, and he offered no instrument for switching it off, because he didn't have one. I've looked for the instrument he was missing. I don't have one either.

THE BETTER QUESTION

The old question is: how do we make our side strong enough to defeat the other side? A better question is: what conditions would make attack less useful, less likely, more visible, and easier to resist without making everyone permanently poorer? Nothing about that reframing abandons defense. It places defense inside a larger system, and it changes what counts as a defense expenditure.

Energy independence is the clearest case. A country that cannot be squeezed on fuel cannot be coerced by the party holding the valve, which means a domestic generating asset does work that no additional squadron can do. Europe learned this the hard way and expensively. Distributed infrastructure does something similar by arithmetic alone: a grid with 10,000 generating points has no single target worth attacking, and the cost of protecting it falls out of the design rather than being added to it afterward. Reliable information reduces panic. Prosperous neighbors can reduce migration pressure and illicit markets. Schools and jobs can shrink the pool of people that violent groups recruit from. Transparent verification reduces uncertainty. None of these is a substitute for a soldier. All are security assets that old budgets misclassify as unrelated social spending.

Collective AI can push in either direction here. It can translate signals, detect coordinated deception, model supply chains, and reveal where populations share interests even when their governments do not; it can equally accelerate targeting, surveillance, cyberattack, and autonomous violence. The technology does not choose the mathematics. The value structure does.

It's satisfying to explain war through villains, and there isn't one available here. Some leaders do commit crimes and have to be resisted. But a system designed to work only when good people are running it isn't a system, and lasting improvement has to come from changing what pays, what people know, how institutions work and who gets a say, so that fewer people can turn private fear or ambition into public catastrophe.

Global military spending reached about $2.9 trillion in 2025, according to SIPRI, and it has gone up rather than down in each of the recent years.[31] Run Eisenhower's own conversion on it and the figure is roughly $350 for every person alive. Some of that money answers real aggression and protects open societies, and defense

cannot be abolished by declaration. What remains open is whether civilization can build forms of security that require a smaller share of its intelligence, materials, talent, and fear.

Nothing in the $2.9 trillion is signed for. What it is instead is the sum of many individually defensible judgments, each made by someone who would be ruined for guessing low and forgiven for guessing high, and there is no one accepting the bill. Eisenhower described the machine from inside it, with every advantage a warning can have, and it went on running. That's where he left it.

I've changed my mind about one thing since I first wrote this. I still don't have an instrument that can switch Eisenhower's machine off, and I'd distrust anyone who claimed to. But there may be one for reducing what the machine has to do: define the objective more narrowly and measure every asset against it.

Defense remains necessary while threats remain. How much destructive capacity a society can manufacture was never the question. What matters is the least destructive power it takes to make an attack fail, so that as much as possible is left for everything else. In that frame, a military is a security instrument with a job to do, not an identity and not a permanent market.

The eight billion people who can say yes to trade, learning, migration, investment, friendship, and exchange can also say no. No, a border does not become more legitimate because enough people have been killed to move it. No, destroying a market and its people is not a durable way to create wealth. No, one population's fear cannot be the only variable admitted to the security calculation.

None of that requires eight billion people to agree about governments, history, religion, or the exact location of every disputed line. What it requires is a sufficiently broad norm that territorial conquest should not be rewarded by the larger system. If invading a neighbor reliably gets a country isolated, resisted, cut off from its markets and beaten back, instead of enriched, the sums change.

A wall can make crossing harder. A military can make invasion fail. But the deeper border includes the lives on both sides. Imagine families on both sides with homes, jobs, schools, healthcare, friends, businesses, and plans for next summer. The line remains. Jurisdiction remains. Defense remains. But both sides possess more to lose by destroying what is across it, and more to gain by keeping the line peaceful.

Which gives the earlier phrase shared prosperity a harder meaning. Your neighbor's floor contributes to your security. Their market can be your market. Their stable institutions reduce the pressure on yours. Their family having something to lose isn't a sentimental advantage to you. It's part of the mathematics of deterrence. If you believe in border walls, the best thing you can do is make both sides of the wall nice. You can have a shorter wall.

Then consider the engineering capacity accumulated inside modern defense: aerospace, robotics, materials science, artificial intelligence, precision manufacturing, logistics, medicine, communications, and capital. Where it's safe to, moving some of that skill toward energy, water, housing, transport and medicine is simply good use of an asset, not a rejection of defense. The skill that made destruction difficult can also make continuation easier.

> *The best weapon is not the largest weapon. It is a world in which fewer people calculate that firing one improves their future.*

Eisenhower was the best-credentialed man alive to say it. He'd commanded the largest amphibious invasion in history, he'd run the machine from the inside for thirty years, and he said it twice, six weeks into the job and again on his way out the door, on television, to the whole country.

The number has gone up in every one of the recent years.

Being right, at the highest available volume, from the most credible available position, isn't the mechanism. Something else has to be.

* * *

PART III

The Collective AI Horizon

When memory, translation, comparison, and question refinement become personal tools.

The third question is what to do with the tools that are arriving faster than the institutions meant to govern them.

CHAPTER 10

AI Turns the Lights On

Expanded memory is not wisdom, but it changes what wisdom can inspect

AI is not a new species arriving from elsewhere. It is a new arrangement of human language, mathematics, machines, energy, data, and questions.

The popular imagination swings between worship and fear. One story says the machine will solve everything. Another says it will take everything. Both stories make the same mistake: they treat AI as an independent object separate from the civilization that trained, powers, prompts, deploys, owns, regulates, and corrects it. Trained on what human beings have written and said, such a system is in that sense a particular kind of mirror, one that reflects a very large sample of the record rather than the person standing in front of it.

A language model is not a perfect copy of human knowledge, and it does not contain humanity in any mystical sense. It is a statistical system, trained on patterns in human-created data, adjusted through human-designed processes, and operated through physical infrastructure. Yet the practical effect can feel like access to a compressed planetary memory. A person can ask across disciplines, translate across languages, compare eras, and receive explanations at a speed no library or human generalist could previously match.

That is super-sentience as I use the term: not machine consciousness, but a human's nervous system connected to memory and cross-correlation far larger than the unaided brain can carry. The roughly three-pound brain remains the source of lived value, responsibility, embodied experience, and choice. The system extends reach. The real question is what decides where the reach goes.

NATURE BECOMING SENTIENT

To me, super-sentience means nature becoming sentient as a description of a civilization, not a claim about the planet. Nothing has awakened. No rock has developed opinions. What has happened is that through human beings, nature produced a form of life that can perceive a galaxy it will never visit, remember a species it drove extinct, model a future that has not occurred, and then change its own behavior on the strength of what it found.

Our instruments extend perception. Our libraries extend memory. Our institutions extend coordination. Artificial intelligence now extends cross-correlation, at a speed and across a scope that no committee, no library staff, and no human generalist has ever matched. In that limited but profound sense, humanity is nature increasing its capacity to notice itself.

Noticing it isn't deciding what to do next, and that gap is the whole remaining problem, which is why the phrase describes the present without celebrating it. Perceive more of yourself clearly and you can make better choices or enforce bad ones with unprecedented efficiency. Lights illuminate whatever is in the room. They do not choose the room.

THE QUESTION DEFINES THE FIELD

Arguments about AI tend to become arguments about the machine. More useful is an argument about the question AI is asked, because a question creates a gradient through an almost infinite answer space and determines what evidence matters, which comparisons are relevant, whose costs count, and what success means, and it does all that before the model has processed a single token.

Questions are the compression algorithm of intelligence. "How do I win?" retrieves a different worldview than "What outcome leaves both sides capable of creating more value tomorrow?" "How do we stop migrants?" retrieves a different worldview from "What

combination of border integrity, lawful movement, and opportunity reduces forced migration?" The model may be equally intelligent in both cases. The question defines the field.

Take a third case, of a kind I've sat through on a city council. A large employer is negotiating to move its regional headquarters to town, bringing several thousand jobs, in exchange for a serious package of tax breaks and infrastructure from the city. Somebody has to decide. But the question can be asked in at least two ways, and the framing isn't neutral.

The first way: Should we approve the package? Ask it like that and a particular kind of evidence shows up. Yes or no is the answer, and the deal is the unit. What arrives is comparative: what other cities offered, the company's revenue projections, the immediate job count, what similar deals looked like. The horizon: the deal's term. The parties: the city and the employer. The costs: the ones attached to the package. Costs that don't appear, the existing businesses competing for the same workers, what happens to rents, the obligations that outlast the tax break, the road built for one use, don't appear because nobody asked.

The second way: What would make this a good deal for the people who will live here in twenty years? Ask it like that and a different world arrives, with a horizon that runs past the incentive period, a set of people who count that takes in the residents, the residents to come, the workers the employer will draw in, and the workers whose wages will move when the labor market does, and evidence that includes cities that took similar deals and where they will be a decade later. Now the city is a place with a long future, not a party to a contract with a term.

Answer the first question and you get a recommendation: approve, reject, or renegotiate on these terms. Bounded, actionable, addressed to the people at the table. Answer the second and you get a set of conditions: this is a good deal if these provisions are in it, these risks

are allocated, and somebody is checking. Harder to execute, more demanding of the institution, and more useful to the people who will live with the result and were never at the table. Nothing differs about the intelligence of whatever produced the analysis. What differs is how each question would lead to a complete answer. That was decided before the model was asked anything.

None of that says the second question is always right and the first always wrong. Some decisions need a firm recommendation, and a government that can't eventually say yes or no isn't governing. Most consequential decisions get answered the first way by default, not by choice, because the first way is shorter, fits the approval process, and produces something the institution already knows how to process.

Someone must ask the second way, and asking it becomes a political act before it's an analytical one; it widens who counts and what costs must be reconciled. The question for every public use of these tools is whether AI makes the second question easier to ask and answer honestly or makes the first faster and more authoritative.

Both questions are about the same decision. The same model answers both. The change is what counts as relevant, who counts as affected, and how far forward the accounting runs. A question is not a search term but is a value structure working on the evidence before any answer is produced. That's what compression means here: the question decides which dimensions of a vast and undifferentiated record come into the foreground, and which stay as background noise. What results is a real map with true things on it, but it is a projection, and every projection chooses what to flatten. Whoever writes the question makes that choice, usually without noticing. Making the choice visible matters more, for most purposes, than improving the model.

QUESTIONS ARE SACRED

Consider what was involved years ago in asking a question of the record, one that nobody nearby could settle. A real question, whose answer is somewhere in the accumulated work of people now dead. Until very recently that involved using a building, a reading language, a card, a librarian's patience, and sometimes weeks of time. Often it also needed a credential, because the rooms where the answer sat were not open to everybody. And sometimes it was a journey. Even in the end, the record did not just deliver an answer. It provided a shelf of material to look through, and nobody else could turn the shelf into an answer.

For most of the time since human beings have been writing things down, few people had that ability. The result turns on the questions that were never asked rather than on the ones that went unanswered. When the answer can't be reached, the question stops forming, and it stops forming early enough that the loss never registers as a loss. Whole populations have gone through life with the settled sense that certain things are known elsewhere, by other people, in other kinds of rooms.

Now, a person with a connection can put a question to nearly the whole record at 11 pm and get back an answer in their own language. They can ask what the doctor's diagnosis means and then ask what the doctor didn't say. I did exactly that after a procedure on my back, at an hour when no doctor was taking calls, and the answer told me what to expect at three, five, and ten days, and that I might be less comfortable before I was more comfortable. Being informed made it bearable.

If the first answer turns out not to be what they meant, they can ask again, differently and instantly. Asking again used to cost another week, and a great many questions were abandoned at exactly that point, one revision short of being useful.

It's hard to describe what that's like, so here's the word. It's joy. Not just satisfaction or convenience, which undersell the thing by a wide margin. The nearest comparison: the first time a language you've been studying turns into a conversation, when you discover that you've not just acquired an ability, but that a wall you'd stopped seeing is a door.

An ordinary person, holding no standing and no credential, can interrogate centuries of accumulated human memory and expect it to answer. That hasn't been true before. From the inside, that is super-sentience, and it's worth being exhilarated about.

Questions are sacred. My first good one was to a librarian.

Answers are the commodity here. They are getting cheaper by the month and they will go on getting cheaper. What stays irreducibly human is the question, because the question is where the values live: what somebody decided was worth knowing, whose interests they counted, how far forward they looked. A machine can generate questions all day. It cannot want an answer.

An uncomfortable edge on that. A system that makes questions cheaper to ask, easier to revise, and safer to ask badly does more for human agency than one that makes answers more accurate.

And a system so smooth that the next question never forms, that closes the matter, satisfies, and sends the user away content, has quietly taken something back from the person it just helped. One kind shows you the door. The other answers so well that you stop noticing the wall.

None of which is evidence. Joy isn't a guarantee, and an exhilarating instrument is just as capable of being an exhilarating instrument of error.

MEMORY IS NOT JUDGMENT

A thousand-times-larger memory can produce thousand-times-larger confusion. Models can synthesize incompatible claims into confident prose. They can inherit prejudice, follow a manipulative prompt, and optimize a metric that has little to do with human flourishing. More information does not automatically produce better direction.

Confusion of this kind has a particular texture that makes it harder to detect than ordinary error. Trained on the full diversity of human opinion on any contested question, a model can access articulate, internally consistent arguments for incompatible positions, and the confidence in its prose comes from the format, not the facts. That is true whether the underlying record supports the conclusion strongly, weakly, or not at all. Always, anyone reading confident prose has had to judge independently whether the confidence is earned, and a system that generates fluent, structured, well-cited prose at high speed makes that judgment more necessary and, for many, harder to exercise. Fluency, the signal of quality people learned to use, is now available to any output regardless of its grounding.

It's an old problem, distinguishing knowledge from persuasion; only the scale and speed are new. Every technology that has lowered the cost of publishing has faced a version of it: the press made more falsehood available alongside more truth, broadcast put fabricated content in the same format as verified reporting, and the internet handed self-publishing to everyone, including the fraudulent, at no cost. AI lowers the cost of fluent, structured, apparently sourced text to near zero, in every language and at any volume.

Genuinely new is the texture of the confidence. If a false story was published, the falsity could have been in the selection of facts, the framing, or an invented source, and the argument was often inconsistent in ways a careful reader could catch. Persuasion had a

human structure, subject to human failure modes: the bad argument that contradicts itself, the implausible leap, the vagueness where a fact should be.

AI-generated text does not have those failure modes in the same places. A model can build a consistent argument for a false conclusion out of true premises correctly assembled, and the seams, where the argument leaves the evidence, don't show in the prose. So the verification problem has moved upstream, from the logic of the argument to the grounding of the record the argument draws on, and that upstream problem is exactly the one a reader cannot solve without the system's own help.

Distrusting the output is the wrong answer, for the same reason nobody stops taking X-rays because the image needs a trained reader. Instead, build the reading into the delivery: show how the question was framed, let people open the sources, say how sure the answer can be given what's behind it, and let the person asking trace the path from question to answer. Any knowledge system that delivers an answer without showing its architecture is asking to be trusted rather than verified.

AI systems are commonly said to hallucinate when they produce a fluent statement that is not supported by the record. A vivid word, and it hides the mechanism. In the human sense, a hallucination is a conscious perception without the external stimulus it appears to represent.

Why the distinction matters: the same generative freedom that produces an error can also produce a novel conjunction. Two separate ideas from the retrieved record can be placed beside one another by the model. Most such conjunctions will be useless. Some may be interesting. Label the output a hypothesis and send it to be checked.

Science already has a name for one version of this move: abduction, the proposal of a possible explanation that must then survive evidence that could be capable of killing it. A machine can participate in that cycle without being conscious. Its novelty becomes valuable only when provenance and testing prevent novelty from being mistaken for knowledge.

A bad idea should be allowed to die cheaply. A surprising idea should be allowed to earn its life.

THE LIGHTS CAN BE AIMED

When I say AI turns the lights on, I don't mean that illumination guarantees virtue. A spotlight can guide a rescue or enable a weapon. AI increases visibility, prediction, persuasion, and coordination. Those capacities amplify the values of the system around them.

So individual control matters, and it's the structural question the rest depends on. A personal AI — one that remembers what the user chooses to preserve, shows the sources behind important claims, distinguishes knowledge from inference, exposes the value assumptions built into its own framing, and allows the user to correct the model — is doing something qualitatively different from one that does not. Intelligence isn't where the difference lies; it's who the intelligence is working for. A tool that expands the user's capacity to leave, compare, refuse, and ask again is an instrument of agency. A tool that narrows those options while appearing to expand them is a very sophisticated version of the delivery model Eisenhower understood and resisted in a different medium.

What the means of verification requires in practice is more specific than the phrase suggests. A patient asking about her own condition can see not only the answer, but which studies it draws on, when they were done, how big they were, and whether there's a minority view in the literature that the summary has smoothed over.

Students using AI on a research question can trace an argument back to its sources and check whether the summary is faithful to them, instead of trusting that it is. Citizens reading an AI summary of a bill can see which provisions were emphasized, which were left out, and what a different framing of the same text would produce. None of these are technically impossible. All of them are design choices that cost something to build and cost something to leave out. The costs of leaving out hit users while the costs of building hit developers. That's why they won't reliably get made right without outside pressure.

Collective benefit arises when millions, and eventually billions, of such empowered individuals can perceive more of the same reality and bring their perceptions into comparison with one another. In 2025, the International Telecommunication Union estimated that approximately 6 billion people, about 74 percent of the global population, were using the internet.[32] The next expansion will not only connect people to information. It connects them to tools that help interpret, interrogate, and test that information against their own circumstances and their own questions.

The remaining divide is not merely access to a network. It includes language, affordability, literacy, disability, local relevance, data rights, and the nerve to question a system instead of taking its answers as final.

Most of the world's knowledge is in a handful of languages, and a translation tool is not knowledge produced in a community's own terms. A person whose data trains a model, and who receives no share of what the model captures, has contributed to someone else's abundance. An AI age that leaves billions as passive data sources will repeat the old mathematics of extraction. An AI age that gives them tools of memory, translation, creation, and verification can enlarge the field of agency rather than concentrate it.

Applied to AI, the mathematics of better returns the same test it applies to anything else: does this leave the people inside it more able to run their own lives, year after year, or does it take more from them, year after year? Nothing about the model itself determines the answer. The question does, and the ownership structure, and the design of the interface, and the policy about who holds the data. Those are human choices, being made by institutions that are being built right now, in the same years these words are being written, and they'll be far harder to revise later than to make right the first time.

CHAPTER 11

History Written by Vectors

The archive gets larger; the weighting becomes the argument

"The life of the law has not been logic: it has been experience."

Oliver Wendell Holmes Jr., The Common Law (1881)

Two forces of the same size pulling in opposite directions add up to nothing. Anyone who took a physics class knows it, and almost nobody applies it to people.

In 1893, in a church basement in Oberlin, Ohio, a Congregational minister named Howard Hyde Russell founded an organization on that arithmetic, though he would not have described it in those terms.[33] The Anti-Saloon League wasn't the first temperance body in America and it was far from the largest. The Woman's Christian Temperance Union, a national organization, had a quarter of a century's head start, and under Frances Willard a program that took in suffrage, labor conditions, prison reform and public education as well as drink. That breadth was its virtue as a movement and its weakness as a force. A politician who disappointed the WCTU on liquor could recover with it on suffrage.

Russell's League took no position on tariffs, on currency, on the Philippines, on anything. It endorsed or rejected Republicans and Democrats, Protestants and Catholics, honest men and scoundrels, on the single question of alcohol. Wayne Bidwell Wheeler, an Oberlin graduate who became the League's general counsel and then its strategist, worked out the rest. His method has been called pressure politics ever since, and it centered on a method of counting that any officeholder could do. Wheeler didn't try to convert the country. He

identified the districts where the last election had been decided by a thin margin, and he demonstrated to the officeholder that the League could deliver, or withhold, more votes than that margin contained.

A Congressman facing a 500-vote margin doesn't have to be persuaded of anything. He has to be able to count.

What followed came at a speed that still seems improbable. State by state, the League took on legislatures, and legislatures were what ratification required. The Eighteenth Amendment, prohibition of alcohol in the United States, cleared Congress in December 1917 and was ratified on January 16, 1919, by forty-six of the forty-eight states; only Connecticut and Rhode Island declined.[34] The enforcement statute that gave the amendment teeth was drafted in substantial part by Wheeler himself, though it carries the name of the congressman who introduced it. A minister's single-issue league in Ohio had amended the Constitution of the United States inside a working lifetime.

Nobody knows what a majority of Americans wanted in 1919. Scientific opinion polling didn't exist yet. The League's bloc was the only measurement of national direction anybody had.

The rest of the country answered back with speakeasies, home stills, juries that would not convict, an enforcement budget that never came within sight of the coastline it was meant to patrol. Finally, in December 1933, the Twenty-first Amendment repealed Prohibition, making the Eighteenth Amendment the only article of the Constitution ever to be repealed.

The bloc Wheeler assembled was never a majority. Alignment did the work numbers couldn't.

History was once written by victors because victors controlled the archive. Wheeler controlled something smaller and, for a while, more useful: the only device in the room that created a number.

THE ARCHIVE THAT NOBODY OWNS

Digital civilization is beginning to record vectors the victor cannot fully own. A victor chooses monuments, schoolbooks, maps, names, and official memory, and the resulting history may contain a great deal of truth while remaining shaped by the power to select what survives. The experiences of ordinary people stayed scattered in letters, artifacts, oral traditions, graves, songs, and silences.

Something has changed about the raw material, the amount of data out there. The archive is vastly larger and more distributed. Sensors measure atmospheric temperature without regard to national rhetoric. Scientific databases preserve methods as well as results, which means a finding arrives with the means of its own refutation attached. Satellites observe fires, crops, ice, construction, and conflict. Social networks capture emotion, manipulation, solidarity, and rumor. A container ship writes its position into a public transponder record every few seconds, so a government's assurance that its largest port is operating normally can be checked against the queue of vessels sitting outside it, by anyone, from anywhere. Markets register millions of choices. Phones record events from several angles at once, including angles nobody planned.

More data does not automatically create a more truthful history. Digital records can be fabricated, censored, selectively amplified, or interpreted without context, and the tools for each of those have improved at least as fast as the tools for correction.

What changes is the possibility of cross-correlation. No single narrator has to carry the entire burden of truth when many independent vectors can be compared against one another.

Everything there rests on the word independent. What happens when it fails?

WHAT IS A HUMAN VECTOR?

A vector has direction and magnitude. Direction is what a person is moving toward. Magnitude is how hard, how consistently, and at what cost they are moving. That is why a stated opinion and a change of address aren't the same size.

In human systems a vector may be a stated value, a repeated choice, a need, a fear, a purchase, a migration, a vote, a question, a refusal, an invention, a measurement, or a sacrifice. None is the whole person. Together they reveal patterns that labels conceal.

Take a man who votes every two years for the candidate promising to hold the school levy down, and who also drives his daughter forty minutes each way to a district he doesn't live in, because the school there is better. His vote and his mileage point in opposite directions. Both are real, both are his. What he says he wants and what he does disagree, and he isn't lying in either. Human beings are internally plural. Groups are even more so.

Collective intelligence should preserve that complexity rather than flatten people into demographic predictions. Vector analysis exists to let people see the field they already create, and using it to manipulate a population more precisely is available as an option, cheaper, and already the main thing anyone does with it. Where millions of vectors point toward safety, dignity, creativity, family, belonging, and a livable future, those overlaps become evidence of possible alignment. Where they conflict, the conflict becomes a design problem rather than an excuse for a ruler to declare what everyone wants.

THE RESULT IS NOT THE SIMPLE SUM

Take a hundred vectors of equal length and point them in random directions, and their sum doesn't come to a hundred. On average it comes to about ten. That's the mathematics of a random walk, it can be checked with a ruler and a protractor in an afternoon, and it holds wherever vectors add: molecules in a gas, footfalls on a deck,

preferences in a population. A hundred people pulling with equal force in scattered directions move a thing about as far as ten of them pulling together.

A resultant vector is therefore not the loudest opinion or the largest vote. Sum with a direction, and magnitude in that sum can stand for intensity, for durability, for how much of a life is riding on the outcome, or for how many people the outcome reaches. Change what magnitude means and the arrow swings. Nothing about the count changes at all.

Financial markets do this sum in public every day, which makes them worth watching even for people with no interest in them. A large listed company might have a billion shares outstanding. On an ordinary trading day perhaps three million of them change hands, under a third of one percent, and the closing price, the number reported everywhere as what the market thinks, is set by the last transactions at the margin of that small flow. The other 997 million shares held an opinion and expressed nothing. Index funds make the point sharper still: several of the largest holders of most large companies buy and sell according to a rule that has nothing to do with the company at all, and are therefore, on any given day, silent by construction. What is called the market's view is the resultant of the people who acted, weighted by how much they moved, and the overwhelming majority of the owners are nowhere in it.

As a model of a public, a market fails twice. A market can add intensity because intensity is denominated: you say how much you care by how much you pay, and the unit is the same unit for everybody. Preferences in a population have no such unit. Anybody who weighs a child's asthma against a landowner's inconvenience has invented an exchange rate between them, and the invention isn't a technical step on the way to the politics. It is the politics. Worse is the second break. Markets weight by wealth, so a market resultant is a

vote in which the rich have more votes, which is exactly why it makes a poor model for a public one and a very convenient one for anybody who would like the public sort to work that way.

So the weighting has to be declared. Whose safety is being counted, and how far into the future does the count extend? What level of uncertainty is acceptable? Which claims are constraints on the answer rather than preferences inside it? A small minority may carry a decisive moral claim when the majority's convenience destroys the minority's basic agency, and no headcount will ever return that result, because a headcount is not the instrument for it. Children, future generations, ecosystems and the people who will live downstream of a decision do not appear in ordinary polling and still belong in the calculation.

Hidden weights produce hidden power. Explicit weights create an argument that can be inspected and revised.

Wheeler's weights were explicit. He published them. One question, absolute priority, everything else at zero: a scheme so transparent any voter in Ohio could recite it, and being transparent made it more effective, not more correct. Declared weighting buys an argument and not an answer, and a reader who has watched a well-organized minority carry a legislature will not need telling that the argument is frequently won by the side that showed up.

THE MEASUREMENT CAN BE CAPTURED

There's a harder problem than weighting, and the most dangerous version of it leaves the instrument looking perfectly healthy. Off Newfoundland the catch index looked healthy almost until the fishery closed.

A resultant is only as honest as the independence of the vectors inside it. Two people who witnessed the same event separately are evidence. Two who agreed on the story beforehand are one witness counted twice, and a crowd manufactures the second case

continuously, without anybody intending it. People report the position they believe their neighbors hold, and their neighbors turn out to have been reporting back to them.

In the summer of 1989 the German Democratic Republic had, by every measurement it or anyone else possessed, the compliance of its population. Elections returned figures near unanimity. Western analysts, sovietologists, intelligence services with large budgets and long experience didn't forecast what came next. On the fourth of September a few hundred demonstrators gathered outside the Nikolaikirche in Leipzig. Five weeks later, on the ninth of October, more than 70,000 walked the ring road; a week later it was 120,000; by the end of the month it was in the hundreds of thousands, and on the ninth of November the Berlin Wall was open.[35] Nothing new had been learned in those nine weeks that wasn't known in June. What had changed was each person's estimate of what everybody else was prepared to do.

Timur Kuran, an economist, gave this its name. In *Private Truths, Public Lies* he called it preference falsification: people state the position they judge to be safe, and because everybody is doing so, everybody's estimate of public opinion is assembled out of everybody else's estimate of public opinion. A society in that condition holds its shape until a small shock moves a few people, whose movement lowers the cost for the next few, and the whole distribution can relocate inside a week. Kuran's uncomfortable conclusion was that such events aren't merely unpredicted but unpredictable in principle, because the information required to predict them is the information the system is built to suppress.

Structural engineers met the same mechanism in a form you can watch move. The Millennium Bridge across the Thames opened on the tenth of June 2000 and closed two days later.[36] Under crowd loading, it swayed sideways, alarmingly, and the cause turned out to be a feedback loop nobody's design code had checked for. Every walking person applies a small sideways force with each step. In a

crowd on a still deck those forces are uncorrelated, they point every which way, and the same mathematics as before, they cancel down to almost nothing.

When a bridge sways, walkers naturally adjust their steps to stay balanced. If their timing matches the movement, a dangerous loop begins; more swaying forces more people to walk in sync, making the bridge sway even more. Once a crowd hits about 150 people, this effect spins out of control.

While the military has banned marching in step on bridges since a deadly 1850 disaster, that rule only stops up and down bouncing. The bridge was structurally strong – engineers just prepared for the wrong direction of movement, vertical instead of side to side.

Engineers still argue over the details. A 2021 paper in *Nature Communications* makes a good case that the instability can emerge without the walkers synchronizing at all. Inserting ninety-odd dampers has kept the bridge steady since.

Damping worked because the coupling between the walkers lived in the steel, where an engineer could reach it. Among a population, the coupling is the communication, and a damper on it is a restriction on speech, installed by whoever owns the deck, quieter and worse. Secondly, not one person on that bridge changed their mind about how they wanted to walk; the synchrony was in their feet, and when the sway stopped they kept on.

A preference cascade alters what people say, then what they are willing to hear, and eventually what they believe, so there is no original gait waiting underneath to be recovered when the motion stops.

Erroneous published results make it worse because the measurement then becomes an input to the thing being measured. Show a population that 70 percent hold a mistaken or just wrong view, and the population moves toward it. So the next reading confirms the first without a single new piece of information having

entered the system. Which is how a measured society ended up more confidently wrong and why wider measurement has never meant a larger poll.

GETTING FURTHER, AND NOT OUT

So an honest vector system has to weight for independence instead of counting agreement, and has to keep the minority reading rather than rounding it away. Where two sources agree, the question worth asking is whether they had access to one another. A system unable to answer that question is measuring an echo and calling it a population.

HISTORY BECOMES PROSPECTIVE

When vectors can be compared as they arrive, history stops being only retrospective. Patterns of misalignment can be seen before they become collapse. Supply stress, public fear, ecological change, disinformation, and institutional distrust are all measured somewhere by somebody; correlating them is mathematics, and the mathematics is not the hard part. Like an X-ray, the image does not cure the fracture, but it makes denial harder.

There's another kind of information that becomes possible when the archive is machine-readable at scale: information about the stream itself. What share of the claims people make is grounded in direct evidence? What share is an exaggeration of something true? What are billions of people actually asking about? Which questions are rising, disappearing, or changing shape?

WHAT THE INSTRUMENT CANNOT SEE

The archive of information is larger and it is not neutral about who appears in it.

A sensor sits where somebody paid to install a sensor. A phone records where phones are carried, charged and kept. Markets register the choices of people with money to move, in the amounts they

moved it. Emission is not evenly distributed across a population, and it is least available to exactly the people whose reading would matter most. Children do not emit. Ecosystems do not. Neither do the unborn, nor the incarcerated, nor anybody who has learned that answering a question about their opinions is a kind of risk, and East Germany is the standing demonstration that a population can hold its most consequential information in exactly that condition, invisibly, for forty years.

So somebody has to put them in. The weighting for the absent parties, the intensity correction, the constraint that stops a majority's convenience from consuming a minority's agency: none of it is discovered by the instrument. All of it is supplied by a person, in advance, as a choice. And whoever chooses which absent parties are represented, and at what weight, is doing the thing the victor did to the archive, with better tools, a cleaner conscience, and a defensible-looking number at the end of it.

Wayne Wheeler died in 1927, six years before the repeal of Prohibition, and by every account he died believing the question had been settled. He had assembled a resultant that was genuine, that was measured, and that pointed where the country would not follow. A larger archive would not have corrected him. His instrument could not see what no instrument sees: the size of what people decline to say out loud.

The archive got larger. The selection moved into mathematics where it is harder to notice and a great deal easier to defend. And the people it leaves out are the same people, in every generation, who were never going to be in the archive either.

THE EQUATION

Many independent vectors + transparent weighting + cross-correlation = a more truthful picture of shared direction and hidden misalignment.

CHAPTER 12

Value-Structured Intelligence

A reasoning system is stronger when its values and tradeoffs can be inspected

In 1569 Gerardus Mercator published a map that solved a real problem. Navigators could lay a ruler across it, read off one compass bearing, and hold that bearing until land appeared. On a globe that path is a spiral; on Mercator's chart it is a straight line. Ships found ports. The map has lasted four centuries, and it bought that convenience by lying about size. Greenland covers a little over two million square kilometers and Africa about thirty million, and on a Mercator wall map the two look about the same size.

Three centuries later James Gall proposed the opposite trade, and in 1973 Arno Peters pushed it hard and in public: keep the areas true, let the shapes stretch, and call the result the honest map. Both men were selling a trade-off as a virtue. A round surface can't be laid flat without giving something up. Gauss proved that in 1827, and nobody has found a way around it since, so the quarrel between the two maps was never really about honesty.[37] Each map is a choice about which loss the mapmaker could live with.

Now think about the room where one of those maps ends up. A child standing in front of it doesn't see a choice. She sees the world, the way things are. Somebody decided before she walked in, and nothing on the wall says so. For most of the twentieth century the map in most American classrooms was a Mercator, so the countries nearest the poles, several of which had done the mapmaking, looked a good deal bigger than they are. Good mapmakers have a fix for this. They print the projection's name in the corner. They add a scale bar that changes with latitude, or a small inset showing the same circle at the equator and again far to the north. None of that removes

the distortion, because it can't be removed. What it does is tell the person looking which distortion they're being shown, so that an argument about Greenland can become an argument about the map instead of an argument about the world. You need a globe.

WHERE IT CAME FROM

VSI didn't begin as a theory. It came out of the experiments described at the front of the book, one question at a time. OptimismAI.app taught us to improve a question before answering it. AskRichard.xyz asked whether what sits underneath a person's writing, what he keeps coming back to and what he weighs against what, could be put to work on a question he had never been asked. AskTheFounders.us made that harder, because the founders disagreed with each other and with themselves, and a set of values had to be able to hold the disagreement.

THE PLUMBING

Start with machinery that already exists. Most AI systems that answer questions about the real world now use a method called retrieval-augmented generation, or RAG. Before the model answers, the system goes and fetches something relevant: a current document, a private archive, a person's own earlier work. It puts that material in front of the model along with the question. That's the whole trick. The model stops answering from memory alone, which matters, because its memory is a blurred average of an enormous amount of text.

VSI uses the same trick, but it fetches something more. Along with the facts, it brings what is needed to interpret them: the principles, priorities, tradeoffs and definitions of success of the person or the community asking. Nobody can collect a town's values the way you'd collect its rainfall. But a community writes down more than you'd think. Its commitments leave a paper trail: minutes of the meeting where they voted; a charter; an environmental review with the

objections still attached; a letter from a tribal council; the budget item a school board has protected through four rounds of cuts. Those values are declared, and they can be quoted.

Other values aren't written down anywhere. They show up as patterns, in what someone keeps choosing over years of decisions or comments, and a machine can sometimes spot them. That kind of value is a guess, and it should be labeled as one. It can be wrong. It needs a source, a date and a way to be corrected, and it should never arrive looking like the values somebody actually wrote down.

Put either kind beside the engineering report and the model answers inside values someone can see, instead of inside whatever values it soaked up in training. And it did soak some up. A system trained on the writing of the last several decades carries a projection of its own, a sense of what is obviously important and what is background noise, and nobody printed its name in the corner. VSI doesn't claim to discover anyone's hidden moral truth. It asks that the projection get a label in the corner. The test is practical: does showing the values improve the question, expose a contradiction that was hidden, or help a person ask a better next one? If not, the values were decoration.

FOUR DAMS ON THE KLAMATH

A fact can be stable while its significance changes by worldview, and the Klamath River shows it at full scale. Four dams stood on that river for most of a century, and everything about them could be measured: how much power they made, what they cost the utility's customers, the salmon runs they blocked. None of those measurements settled the question. To the company that owned the dams and the households they supplied, they were infrastructure. A community that prioritizes energy independence may judge such a thing differently from one that prioritizes a sacred landscape. To the Yurok,

the Karuk, the Hoopa Valley, and the Klamath Tribes, who had been on that river far longer than the dams, they stood between a people and its own food.

The last of the four dams was removed in 2024.[38] Nothing in the engineering had changed. What changed was whose values the numbers were finally read through, and that took decades of testimony, lawsuits, hearings and a licensing deadline.

Suppose the question had gone to an AI in 1998: should these dams come out? A system holding only the engineering report gives a reasonable answer that leans toward what can be measured, because that's all it was given. Add the tribal council's submissions and the same question, over the same engineering, comes back different. VSI should not hide the difference inside an opaque model response. It should surface the value structure and show how the conclusion changes when priorities change. No reverence for a river has entered the machine. What has entered is an idea it didn't have before: that a fishery can be a people's food supply and not a pleasant view, and that a run of salmon can matter to a people the way a power line matters to a grid.

Better information did not remove those dams. Tribal governments with legal standing removed them, helped by a federal licensing deadline and by a utility that eventually worked out that removal was cheaper than compliance. Making someone's values visible gives them no power at all. A system that lays out the Yurok or Karuk position beautifully, and changes nothing about who decides, has improved the record and left the outcome alone.

THE LABEL IN THE CORNER

People who aren't engineers usually ask the harder question. Does the model understand a community's commitment to energy independence the way a neighbor would understand it across a fence? The years of it? The humiliation of a bill nobody can argue with?

I don't believe it does. I see no evidence that it does and no theory by which it could.

It doesn't need to. What it needs is simpler, and anyone can check it: the system has to show which values are in play, and its conclusion has to move when those values change. Swap the documents and watch whether the recommendation shifts. If it shifts for reasons that hold up, it's working. If it doesn't shift at all, the values were decoration. Two runs and a comparison is the whole test, and a school board, a planning commission or a reporter on deadline can run it in an afternoon without any access to the model's insides. Whether a machine understands would take a laboratory to settle. Whether the answer moved takes a printer.

Anyone can work out the Mercator distortion at any latitude, because the formula is fixed, published, and the same on Tuesday as it was on Monday. A model's sense of what matters has none of those properties. It drifts between versions, moves when a question is reworded, and can't be inspected even by the people who built it. That's why the values have to be declared up front: you can't read them off the answer the way you can read the distortion off a map. Mercator, whatever else he was up to, showed his work.

Courts have spent centuries on this exact problem, and they left their reasoning in writing.

A trial doesn't let every true thing into the room. In the United States a judge may keep out relevant evidence when it would mislead or inflame a jury more than it would help. That's Rule 403, a few lines long. Rule 404 goes further: a jury generally can't hear about a defendant's character to show that he acted true to type, which is why it may never learn about a prior conviction. Neither rule says anything about what the jury understands. Both are about what is put in front of the jury, and what goes on the record.

VSI is the same kind of rule. Rules of evidence don't make a judge wise. They make the choosing open to challenge. A lawyer can argue that a document belongs in; the judge rules in open court; the ruling is written down; and a higher court that was never in the room can review it. The jury never sees what was kept out, but the appeals court sees all of it. Choosing is as unavoidable in a trial as it is in a model, and the courts' answer was to keep a lasting record of the choice.

A courtroom also comes with lawyers on both sides and an appeals court above them both. Nobody has to pay a judge in money for a verdict to matter to the judge. A ruling can follow a judge for a whole career, speed it up or stall it, and win or lose the good opinion of the people the judge has dinner with. That is why the lawyers, the written ruling and the court above are there. Those are what make the rules mean anything, and an AI system has none of them. Nobody argues that your document should be included. Nobody rules on it in public. There's no appeal, and usually there isn't even a list of what was considered and left out. So the courtroom describes no system that exists today. It describes what one would need first.

A map is easy to label, because it's a fixed object and the label can be printed once. A system that answers questions is harder, and most of the difficulty is in what the person sees on the screen, not in the fetching.

Four things must be visible, and none of them is a breakthrough.

1. The documents used, by name, with the passage that was used, so a reader can go to the original and see whether it says what the answer suggests it says.

2. Where each document came from and who put it there, because a district's own board minutes and a vendor's sales pitch should never arrive looking alike.

3. The disagreements the system found and didn't settle, left as disagreements.

4. And a control, a switch or a menu, something a person can work without taking a course, that loads a different set of values and shows what moved.

Mercator's chart can't show you which areas it has drawn too big; you need the other map, and somebody has to hand it to you. A system can hold both sets of values and switch between them at no cost, and that is the genuinely new thing. It comes from the answer being made at the moment you ask for it, instead of printed a year earlier. Compare the default: one confident answer, built from values the reader never declared and can't inspect, delivered in the voice of mathematics. Nothing hypothetical about that failure. It's what most of these systems do now, competently and at very large scale, and the competence is what makes it hard to notice.

All four cost the same thing. Answers get longer, slower and less satisfying, and they stop sounding like verdicts. There's a reason products aren't built this way. A label in the corner asks the reader to do something with it, and most people, most of the time, would rather be handed the number.

A FOUR-DAY WEEK

Projects such as AskRichard.xyz and AskTheFounders.us are experiments in this direction. One asks whether a person's long-lived values can become something others reason with, without pretending to bring the person back. The other asks whether the recorded disagreements among the American founders can inform today's constitutional questions without turning history into an oracle. The intelligence there lies partly in the disagreement. Neither is far enough along to produce a session worth printing. So here is the method, laid out so the parts can be checked and disputed.

Take a real and thoroughly unglamorous question. A rural school district with 1,100 students is weighing a move to a four-day school week. The superintendent asks the system directly: Should we do this? With good fetching and no declared values, a system does

something perfectly reasonable. It finds studies: several showing flat or slightly lower test scores, several showing that teachers stay longer, one or two showing the effects fall hardest on the students who were already behind. It sums them up fairly, notes that the evidence is mixed, and stops. Accurate, and also what the superintendent could have gotten from anyone, and it leaves the decision exactly where it was.

Now declare the values, which in this district means adding six or seven documents: the board's mission statement; the minutes of the meeting eighteen months ago where the board voted down combining two bus routes, and the reasons it gave; the memo the principal wrote after losing four teachers in a year; a letter from the parents' association; the number of children on free lunch; the transcript of the public comment session.

The answer changes shape, not direction. Instead of a balanced summary, the system shows which commitments are in play and where they collide. Keeping teachers pulls toward the four-day week. The free-lunch figures pull hard against it, because two hundred children eat their most reliable meal at school. And the bus-route minutes say, in the board's own words, that this district treats travel time for families as a cost it will carry rather than pass on. The conflict is shown, not settled. Nobody's position gets averaged into a recommendation that satisfies neither side.

Then the check, and if you ever want to test one of these systems, this is the one to run. Swap the parents' letter for one from a neighboring district where most families work fixed shifts at a single plant, and watch whether the childcare problem moves from a paragraph near the bottom to the center of the answer.

And then the part I find most interesting: the original question turns out to be the wrong one. Once the fifth day is looked at instead of assumed, what it provides now, to whom, and whether any of that is the calendar's job, the decision in front of the board turns out to be

about meals, supervision and a bus route, with the schedule as one tool among several. That's a better question than the one the superintendent brought. He didn't ask it because it wasn't the question he had.

I made up that example, the documents, the question, and the tidy reversal at the end, which is exactly the condition under which any method works. What's still missing is a session somebody else ran, on documents they chose, producing a result they didn't expect. Until that exists, read the example as a description of how the mechanism would work, not as evidence that it does.

PERSISTENT QUESTION EVOLUTION

A set of values isn't one lump, and the school district shows why that matters. Somewhere in those documents is what the district takes as settled fact: enrollment, budget, what the state requires of it. Elsewhere are things it believes but could be talked out of, such as the idea that its teacher problem is really about pay. Then come the rules it uses to judge a tradeoff, visible in the bus-route vote and applied again whenever a cost is about to land on families. Under those sit the things it treats as simply important, dignity, safety, the idea that a child's day should have an adult in it, which nobody writes down because everybody assumes them. Deeper still is what the district thinks counts at all: whether a family is a unit, whether the town is one, whether a child ten years from now has a say in a decision made tonight. And running through all of it is memory. The four teachers who left. The year the levy failed. The superintendent before this one.

Six layers, listed in the order they're easiest to see rather than the order they matter, and even that order is a projection.

People grow. Institutions reinterpret their missions, sometimes cynically and sometimes because they have genuinely learned something. New evidence changes belief, or ought to. So the values have to be editable, open to inspection, and dated: if the district's

commitments shifted after the levy failed, somebody should be able to see the shift, put a date on it, and ask who made it. A worldview that cannot evolve becomes ideology encoded as software.

Persistent question refinement began before VSI, with OptimismAI.app. What VSI added was one more thing for the questions to test: the values themselves. A normal system treats the user's first question as the target. Treat it instead as the first draft, and the purpose of the system changes underneath you.

In plain terms, the cycle goes like this:

1. Answer the question as clearly as the evidence allows.

2. Look at what the answer had to assume, who it left out, and what would have to be true for it to be wrong.

3. Hold the answer up against the declared values and against what the person actually wants to happen.

4. Ask the next question, which ought to be a better one than the first.

5. Look again, carrying forward what was learned and what is still uncertain.

6. Stop when a new question no longer improves the decision.

Everyone skips the step of stopping, and it isn't a small omission. A system that always comes up with another question isn't helping you think. It's keeping you on the app, and from the outside the two look the same for about six rounds. A system that can't say why it stopped has quietly changed what it's working for.

VSI MUST PROTECT AGENCY

Value-structured systems become dangerous the moment the values are hidden, imposed, or tuned for whoever owns the platform. A government that controlled the values could build in obedience and file the result under civic education. A company would build in its own profit, and it wouldn't need to be sinister about it: a default that

quietly ranks its own products above the alternatives would do the job, and would pass any review that only checked whether the answers were true. One charismatic person with enough documents, and an institution willing to host them after he's gone, could preserve prejudice as legacy. Nothing about this design is automatically good. I'd go further. A declared set of values is also a set of values somebody can seize, and I suspect it is easier to capture than the hidden systems it's meant to improve on, because a hidden system has nothing so convenient to grab.

The safeguards follow from the mathematics of agency rather than from anybody's good intentions:

* Keep the user's own values separate from the institution's defaults and from whatever the model guessed on its own.

* Allow correction, deletion, and refusal, including refusal of the whole thing.

* Measure success by what the person can decide afterward without the system, not by how often they come back to it.

Every item on that list costs money to build and cuts the time people spend inside the product, which is why none of them is standard.

None of this stays private between one person and one product. A hundred million people reasoning inside values nobody declared add up to a civilization looking at a map with no label in the corner, and the second effect is worse than the first. Any single wrong answer can be fixed; somebody checks it. What can't be fixed is losing track of where people disagree. Two people who ask the same system and get the same shape of answer have no way to find where they actually differ, and an argument nobody can find is an argument nobody can settle. Declared values make disagreement visible again. Two neighbors can look at the same question, see that one is reasoning

from the board minutes and the other from the principal's memo, and then have the argument they were always going to have, with the material in front of them instead of hidden underneath.

VSI is a research program, not a settled field. People are already working close by, on constitutional AI, value retrieval, cultural alignment, moral graphs, explicit memory systems and graph-based retrieval. The question our work raises is whether a written set of values, declared where possible, guessed where useful, always open to correction and allowed to change, makes the answers better when it sits beside the evidence and the questions that improve as they go. It needs testing against alternatives. If it has been done better elsewhere, I'd like to know. The label in the corner matters more than who prints it.

Underneath all of it is a question. Alignment, as AI researchers use the word, has mostly been asked in one direction: how to keep the machine in line with what people intend and value. The reverse question gets far less attention. How does the machine help keep a person's beliefs in line with the evidence, and their decisions in line with values they've actually declared, while the last judgment stays with the person? VSI leans toward the second question, and I don't think it's been well asked yet, let alone answered. If it turns out to matter more, a good deal of today's alignment work is measuring the wrong end of the loop. I can't settle that. It's the next better question.

VSI is therefore not a machine claiming sentience, and it is not presented here as a demonstrated solution to AI alignment. It is an architecture for making some of the values that already direct intelligence more explicit, retrievable, comparable and correctable. Wall maps are going up in a great many rooms right now, faster than anybody is governing, and all this asks for is the label in the corner. It doesn't tell the child what to think about Greenland. It does tell her that a decision was made before she arrived.

THE EQUATION

RAG retrieves relevant knowledge. VSI makes relevant value structure visible and revisable as another input to reasoning. Persistent Question Evolution improves the inquiry until the decision has a wider field of view.

CHAPTER 13

Alignment Without Uniformity

A stable system can preserve disagreement while improving coordination

The strongest objection to all of this runs as follows.

You want expanded agency, wider ledgers, shared visibility, and difference turned into resilience, and you assume that more people with better tools and fuller information will tend, over time, to produce more durable collective goods. But that assumption has a name, it's a form of liberal universalism, and it has a history. Genuine progress is in that history, and so is a pattern that repeats. A reformer knows what the better arrangement looks like. She builds the mechanism that will produce it. And a generation later the mechanism has been captured by people who look a great deal like her, and has been delivering its benefits mainly to them. A wider ledger was promised and not delivered. What was delivered was a set of credentials, a set of offices, a compliance apparatus, and a generation of people who spent their careers administering the promise rather than keeping it, and who, when the promised outcomes failed to arrive, concluded that more administration was required.

* * *

WHAT UNIFORMITY WOULD ACTUALLY REQUIRE

The case for uniformity, for enforcing one set of answers, one culture, one definition of the good, is an engineering case before it's a moral one, and it deserves to be heard as one. Coordination is expensive. The more divergent the values of the people in a system, the higher the transaction costs of reaching any agreement. A community that

shares religion, language, kinship structure, and historical memory can coordinate at low cost because the shared background does most of the work before anyone opens their mouth. No such community is imaginary. For most of human history, most human beings lived inside it. The institutions that made large-scale cooperation possible, the empire, the church, the nation-state, were each, in their way, a bet that enough shared ground could be imposed or inherited to let people work together.

Whoever didn't fit paid for that bet. A language that is not the official one becomes a private matter rather than a public resource. A family structure that is not the legal one becomes unrecognized. A religion that is not the protected one becomes a liability. What forced uniformity costs has a long history, and it is primarily paid by people with less power to resist the imposition.

Uniformity is also fragile, in ways that make it a poor engineering solution even on its own terms. A field planted with one genetic variety can be efficient until one disease arrives. A company that rewards one type of thinker can move quickly until the environment changes. A political movement that purges internal tension loses the information needed to correct itself. The Irish Potato Famine is the case usually reached for, and it's worth handling carefully, because the easy version of it is wrong. Its genetic half is real. Phytophthora infestans arrived in 1845 and found in the Lumper potato, a population with no resistance at all, which is why a staple crop failed across an entire island rather than in patches.[39] Monoculture, behaving exactly as the argument predicts.

But the blight destroyed the potato. It didn't decide who starved.

Ireland remained a net exporter of food through most of the famine years. Grain, cattle, and butter left for British ports, some of it under armed guard, out of districts where people were dying. Roughly a million people died and roughly a million emigrated over

the following seven years, and the size of those numbers is a fact about land tenure, rents, relief policy and a government's choices, not a fact about a fungus.

Two monocultures were operating in Ireland in 1845. One was genetic and one was political, and the second had removed the alternatives just as thoroughly as the first. A system with no variety left in its crop and no variety left in who was permitted to decide anything about it failed in both dimensions at once.

A civilization that enforces uniformity of belief, method, or approach is making the same bet as the Lumper potato: that the environment will remain stable enough for the dominant form to keep winning. When the environment changes, and the present one is changing faster than any since the industrial revolution, the system with diversity in reserve has the material to adapt. Without it, a system doesn't. Diversity is a biological and informational hedge against a future that nobody has correctly specified in advance.

* * *

WHERE THE ARGUMENT BREAKS

Admitting that uniformity is fragile doesn't automatically rescue the case for managed diversity. Here the reformer's critique lands, and it lands hard.

Whatever mechanism gets built to preserve and value difference will reflect the priorities of the people who designed it. Those people are usually drawn from a narrow cultural, economic and educational range. In wealthy countries they reliably are. So the mechanism treats the diversity it already recognizes and has words for as the real kind, and simply does not see the rest. What you get is a system that produces visible markers of inclusion while concentrating its actual benefits among those already positioned to navigate it.

One case from American housing policy shows this exactly.

In 1977, the Community Reinvestment Act directed federally insured banks to serve the communities in which they took deposits, including lower-income neighborhoods historically excluded from mortgage credit. The goal was genuine. Redlining, the systematic refusal of mortgage credit to predominantly Black neighborhoods, had produced a two-tier housing market in which homeownership, the primary vehicle of wealth accumulation for the American middle class, was structurally unavailable to millions of families. The CRA was followed in 1992 by the Federal Housing Enterprises Financial Safety and Soundness Act, which required HUD to set Fannie Mae and Freddie Mac explicit affordable-housing targets: a stated share of their mortgage purchases had to serve low- and moderate-income borrowers. The share began around thirty percent and was ratcheted upward for the next fifteen years: fifty percent by 2001, fifty-six by 2008.[40]

What happened between the intent and the outcome isn't a simple story, and the people who tell it simply are usually serving a political point rather than an analysis. On the political left the story is that the policy worked until deregulation undermined it. On the right it's that the affordable-housing mandates caused the financial crisis. Both versions are wrong, and wrong in ways that matter.

What actually happened is that the mandate was one signal among many in a market that had stopped caring whether a loan got repaid. An originator could package a mortgage, sell it on, and be out of the deal before the first payment came due, and the buyers of the riskiest paper were mostly private issuers, not Fannie and Freddie, with rating agencies grading the bundles and banks holding them on borrowed money against a house price that everybody assumed would keep rising. When the Federal Reserve's own economists went looking afterwards, they found little evidence that either the Community Reinvestment Act or the housing goals played a significant part, and the crisis commission's majority concluded that the goals contributed only marginally to what the two agencies

bought.[41] Nobody needed a mandate to make a bad loan in 2005. What the mandate did was smaller and more particular. It gave a name and a virtue to loans that were being made anyway. The incentive to make affordable loans became, at the margin, an incentive to make loans that were classed as affordable, and teaser rates, balloon payments, and prepayment penalties that looked affordable at origination while moving the risk onto the borrower counted toward the targets like any other. Some of the credit that was supposed to build wealth for excluded households was structured to extract it from them instead.

When the market turned in 2006 and 2007, the households who had received that credit lost it in foreclosure. A 2011 Pew Research Center analysis of Census Bureau data put the numbers on it. Between 2005 and 2009, median wealth fell 53 percent among Black American families and 66 percent among Hispanic families. Among white families it fell 16 percent.[42] Read those three figures side by side and then read the phrase widening access again. Those who were supposed to have been included lost the most. Those who ran the inclusion, and whose pay was tied to volume rather than outcomes, didn't.

Captured, and the capture was structural before it was ideological. Anyone paid for the number of loans made, rather than for how the borrowers fared, had every reason to make more loans.

* * *

WHAT FOLLOWS FROM THE CAPTURE

The failure of a mechanism doesn't refute the goal the mechanism was trying to serve. Redlining was real. Its effects on the distribution of American wealth are measurable and large and have not been repaired by the failure of the policy designed to address them. The

families who lost wealth in 2008 lost it on top of wealth that had already been denied them for two generations before the CRA existed.

What the capture raises is whether the mechanism was designed in a way that made capture predictable, and whether a differently designed one would have resisted it.

A program holds up better when the people running it are paid for results, not activity. A mortgage program that makes the originating institution keep a share of the credit risk for the life of the loan gives it a reason to ask whether the borrower can actually repay, because it will share in the consequence if they can't. Originate-to-distribute, in which the originator's exposure ended at sale, eliminated that incentive systematically, and the reform that would have made the mechanism work was about who bore the consequences, not about intent.

A program holds up better when anyone can see how it's doing. Suppose the affordable-housing targets had been measured differently. Not by the volume of qualifying loans written, but by homeownership rates and household wealth among the targeted population five and ten years later. Any gap between the metric and the goal would have opened in plain view, years before the crisis, instead of being discovered in the wreckage afterward. Count the activity instead of the result, and the capture stays invisible.

A program designed by the people it is meant to help holds up better than one designed for them. Their design will often be worse. More robust anyway, because the designers will personally bear some of the consequences of getting it wrong, and that's still the only reliable incentive anyone has found. Also the harder reform, because it means changing who has access to the design process, which is the access problem the mechanism was trying to solve in the first place.

None of this says such programs shouldn't be built. A program that isn't answerable for its results will, over time, come to serve the people who run it. The results have to be counted among the people it was meant to serve, not the people running it, because those two numbers drift apart quietly and for years. Nobody's values have to fail for that to happen. Structure does it on its own.

Expanding access and mandating outcomes are separate operations, and only the first follows from the mathematics. A market gains when a person who was locked out can finally participate: the transaction that was impossible because of the exclusion becomes possible, and both sides can benefit from it. Nothing is gained from a rule that assigns a position regardless of whether the person can hold it. Worse, the rule damages the person it was written for. She now carries a suspicion that the position was a grant rather than a demonstration, and she carries it whether or not it's true, which is the part that can't be argued away. The case was always economic, not moral: an educated customer and a stable neighbor each add capacity that did not exist before. Opening the door is what that argument calls for, and it doesn't say who must be seated at the table once inside, and the distinction matters because saying so is where capture begins.

Harder, and usually left out, is the second point. Any mechanism designed to widen access creates a position worth holding, for the designer, the administrator, the credentialing body, the compliance office, and positions worth holding attract people who want to hold them rather than serve the purpose they were created for. That's a design requirement, not a counsel of despair. Hold the gatekeepers accountable for outcomes rather than for activity. Not how many doors were opened. Who got through, and whether it did them any good.

* * *

WHAT ALIGNMENT ACTUALLY REQUIRES

Humanity does not need eight billion identical answers. It needs systems in which eight billion different lives can discover enough shared direction to cooperate.

Diversity is not automatically wise. Difference can produce paralysis, mistrust, and conflict, and anyone who has watched a genuinely diverse group try to decide without a shared framework has seen this. Its value appears when a system can translate and integrate perspectives without erasing them. Alignment is the coordination capacity that turns difference into resilience.

The word oneness can sound spiritual or authoritarian. Here it means something simpler: the recognition that human lives are embedded in one physical, ecological, informational, and economic system. Atmospheric carbon does not stop at a border. A virus does not read a constitution. A supply chain joins workers and customers who may never know one another. A false financial claim can cross the world before breakfast. Shared system does not imply shared culture. People may continue to speak different languages, worship differently, paint different pictures, organize families differently, and sing in the shower with radically different talent. Diversity is part of the abundance. A shared system does mean sharing some responsibility: consequences that no technology and no border can wall you off from, and an agreed way of measuring what they touch. The better next protects the individual because the collective becomes more intelligent when more genuine perspectives survive.

Democracy is often defended as a moral inheritance. It is also a correction architecture. It increases the number of people who can signal failure, replace leadership, organize alternatives, publish criticism, and revise law. Its procedures are slow and vulnerable to money, propaganda, tribalism, and apathy. Its structural advantage is

that it does not require one ruler to perceive the whole field: a system with more correction points fails less catastrophically than one with fewer, even when the additional points are imperfect.

The U.S. Constitution embodies this incompleteness. Its framers disagreed profoundly. The document contains mechanisms for tension, amendment, federalism, separation of powers, and recurring elections because no generation could finish the argument. The contradictions of the founding, freedom beside slavery above all, do not erase the structure's value. They show why amendment and expanding participation were required from the start.

AI can strengthen democratic correction or destroy it. It can translate legislation, compare claims, model consequences, and let citizens question a civic corpus. It can also personalize propaganda, flood discourse, and centralize surveillance. The test is whether the technology increases the citizen's agency and the visibility of institutional power.

When more people can understand choices, express values, inspect consequences, and remember outcomes, alignment becomes more likely by the nature of the mathematics. Not guaranteed. More likely.

Why? Because ignorance creates false conflict. People argue over labels while sharing goals. Hidden costs allow one group's local better to damage another without either party being able to see the connection. Poor translation turns difference into threat by making the difference look larger than the common ground beneath it. Short memory permits the same manipulation to be run again on a population that cannot recognize it. AI in people's own hands can reduce each of these by making the background, the consequences and the common interest easier to see, and so easier to bargain over.

And that same visibility is the danger. A system good enough to show people where their interests overlap is good enough to manufacture the appearance of overlap. So the alignment is conditional: more likely where the tools expand what each person

can independently check, and less likely where they expand what everybody sees at once while nobody controls what they see. Design decides which.

Greater agency can also reveal genuine conflicts that cannot be dissolved. Land cannot always serve incompatible uses. Rights can collide. Resources can remain limited. Alignment then means agreeing on a fair process, legitimate constraints, and a method of revision that keeps the agreement open to challenge, not pretending that everyone receives a preferred outcome. The mathematics of better doesn't dissolve genuine conflict. It separates the conflicts that are real from the much larger set that come from hidden costs, poor translation, and short memory, and clearing those away is where most of the work is.

What the housing example adds is specific. Alignment without uniformity doesn't administer itself. Mechanisms are required, mechanisms can be captured, and good intentions on the part of the designers do not prevent it. What prevents it is a program that lives with its own consequences: one that counts results among the people it is meant to serve, lets anyone see how it's doing, and is designed partly by people who have something real to lose if it fails.

A gyroscope remains upright not by freezing, but by continuously responding to forces. A healthy democracy is similar. Balance is dynamic. Agency supplies motion and correction. Shared truth supplies orientation. Values supply the definition of falling. Correction is built into the structure, and the building is the work.

The Union is strengthened through understanding, not uniformity.

THE EQUATION

Diversity + agency + translation + shared evidence + fair revision = alignment without uniformity.

CHAPTER 14

Two AI-Literate Generations

A planning horizon for collective memory, agency, and correction

> *"For this discovery of yours will create forgetfulness in the learners' souls, because they will not use their memories; they will trust to the external written characters and not remember of themselves."*
>
> Plato, Phaedrus, trans. Benjamin Jowett — King Thamus, on the invention of writing

Civilizational transitions feel slow from inside a day and fast from inside a lifetime.

The wheel and the washing machine are separated by millennia. The television and the smartphone fit inside one lifetime. The smartphone and generative AI fit inside a fraction of it. The acceleration is not simply that inventions arrive faster. Each new layer becomes infrastructure for inventing the next layer.

A child who grows up with AI will not experience it as a visitor. It will be part of reading, creating, diagnosing, designing, translating, governing, and learning, just as electricity became an invisible assumption of modern life. The first AI-literate generation will build institutions and businesses around capabilities that today's leaders still treat as demonstrations. The second will inherit those institutions as ordinary.

Two generations is not a prophecy tied to a calendar. It is a planning horizon, the time it would take to finish, and nothing has to wait for it to start. Some of it happened yesterday. In less than fifty years, people educated from childhood to use AI critically could hold most positions of responsibility. They could normalize source

inspection, model comparison, personal memory control, value transparency, and persistent question evolution. They could also normalize surveillance, synthetic persuasion, and dependency. Literacy determines which future compounds.

WHAT THE WORD LITERACY HAS ALWAYS MEANT

Literacy, in the original sense, meant more than recognizing letters. Meeting a written claim and bringing independent judgment to bear on whether it was true: that was literacy. Anyone who could decode text but couldn't question it was functionally illiterate in a way that mattered, with access to the record and no means of evaluating it. What worried people about early print wasn't that ordinary people would read but that they'd read without the critical training that had been kept, before print, for people with a scholastic education. King Thamus worried that writing would make people forget. It did, a little, and in exchange a temple clerk's receipt went on saying the same thing after he was dead.

AI literacy is the same problem at a new scale, and writing a prompt is the least of it. Anyone who can write a prompt but can't evaluate the answer is in the position of the 16th-century reader who could decode text but had no way to compare the account to others, inspect the source, or recognize the argument that was not being made. The tools are powerful enough to require the second kind of literacy: the capacity to interrogate, not merely to receive.

It includes knowing that fluent output is not the same as truth. Asking for sources, uncertainty, alternatives, and missing evidence. Understanding how data, incentives, ownership, and values shape an answer. Protecting personal memory and deciding what a system may retain. Using AI to expand creation rather than avoid learning. Recognizing manipulation, synthetic media, and automated consensus. Comparing models and preserving the right to work without one. Grounding digital knowledge in embodied experience, community, and nature.

In practice, that looks like this. A student writing a research paper with AI help who can't trace the argument back to its sources, and check whether the summary is faithful to them, has a credential that outruns their understanding, and the mismatch is worse than plain ignorance because it's invisible.

Calculators brought the same problem into classrooms, and the schools mostly got the question wrong for a decade. It was never whether calculators should be used. It was whether a student using one could still reconstruct the arithmetic underneath. Arithmetic wasn't the point. Being able to tell when the machine was wrong was the point. With AI, the question is harder because the output is in prose rather than numbers, and prose is harder to spot-check than a calculation.

A medical AI that summarizes a patient's risk factors needs a physician who can ask which studies does this draw from, how large were the samples, is there a contested view in the literature that this synthesis has smoothed over? Without that capacity, the physician has delegated the clinical judgment to a system they cannot audit. With it, they have an instrument that covers more of the record than they could cover alone and can be checked. Nothing about their access to the tool separates those two physicians. What they learned before they met it does.

THE EDUCATIONAL TASK: WHAT TO PRESERVE, WHAT TO ADD

Finland's 2016 national school curriculum revision is worth examining because it made two moves simultaneously that most educational debates treat as incompatible.

First, phenomenon-based learning: periods in which students work across subject boundaries on real questions, the local watershed, the history and economics of their own town, the physics and policy of energy, rather than receiving each discipline as a separate delivery. The idea was to put students inside a genuine question before they

are given the conceptual tools for answering it, on the theory that understanding what a tool is for is a prerequisite for using it well, rather than an optional supplement.

Second, and this got less attention abroad, was a deliberate expansion of unstructured outdoor time and manual skill instruction: building things, cooking, tending growing plants, navigating without a device. Not nostalgia. Unusually, the argument was made explicitly in the curriculum documents. Embodied activity builds forms of attention. Abstract reasoning and digital tool-use run on top of those forms. A student who skips the embodied part doesn't arrive as more efficient. He arrives with less to stand on. Nobody can yet credit the reform with anything, or blame it for anything either. Finland's PISA (Programme for International Student Assessment) scores have been falling since around 2006, a decade before the curriculum changed, and the 2016 reform was phased in gradually enough that it had barely reached the fifteen-year-olds sitting the 2018 test. The Finnish National Agency for Education's own account of the decline names widening socio-economic inequality between families, education funding, collapsing reading among the least motivated, social media, and adolescent mental health. Curriculum is not at the top of its list. Most Western systems have been sliding over the same period, so Finland isn't an outlier. It also isn't a proof, and a reform whose first properly exposed cohort has not finished school is not evidence of anything yet. It is a hypothesis with a country attached.[43]

What matters is the combination. Adding AI literacy to a curriculum that is already purely abstract and screen-based produces students who are more capable with one more abstract tool. Add it instead to a curriculum that also contains embodied learning, real questions, and the experience of working with incomplete information in the physical world, and something different happens. Students arrive at the tool carrying something: a reference to compare the output against, a sense of what a claim feels like when it isn't true, a body of experience the model can't fabricate away.

Children already understand persistent inquiry. They ask why until adults run out of answers. The educational task is to preserve that instinct while adding evidence, patience, moral imagination, and technical skill.

THE TRANSITION MUST INCLUDE THE PEOPLE IT DISPLACES

Every better next creates losers through transition. Appliances changed domestic labor. Automobiles changed horses, streets, and cities. Renewable energy changes regions built around extraction. AI changes tasks performed by clerks, analysts, artists, programmers, teachers, lawyers, and executives.

Telling displaced people that progress is inevitable is not a policy. The mathematics of better must count them. A transition that concentrates gains and socializes disruption may be technologically advanced and structurally backward. People who helped build the old system should have pathways to ownership, training, income, dignity, and roles in the new one. A person whose skills were made obsolete by a transition they didn't design and didn't consent to has had their capacity reduced by a process the ledger scored as a gain. No harder case exists.

Denmark is the most closely studied working example of what it takes to manage that case at scale.

In the 1980s and 1990s, Denmark's traditional manufacturing industries, textiles, shipbuilding, furniture, faced the same competitive pressure from lower-wage economies that was devastating comparable industries across Western Europe. Politically the options were the familiar ones: protect the industries through tariffs and subsidies, or accept the displacement and manage its consequences. Denmark chose neither, or rather chose a third path with three components and required all three to function.

First, flexicurity. Employers could dismiss workers with minimal restriction, which removed the need to avoid hiring. In exchange, a dismissed worker received income replacement at roughly ninety percent of previous wages up to a ceiling. That matters because it makes the scheme far more generous to a cleaner than to an engineer, for up to four years. The duration was cut to two years in 2010, the last step in a series of reductions that had already reduced it from eight.[44]

Second, mandatory worker participation in active labor market programs, retraining, relocation assistance, job placement, as a condition of continued income support.

Third, and it gets left out of the abbreviated version, which is why the abbreviated version keeps failing to transplant: a tax structure capable of funding the first two. Denmark's total tax burden as a share of GDP is among the highest in the world. Decades of trust-building between employers, unions and government went into the political consensus behind that, and those three could only sit in a room together because of institutions built in the nineteenth century.

None of it was frictionless. Workers whose skills were made obsolete didn't always find equivalent work. Some regions changed permanently. In the 2008 financial crisis the model came under serious strain when the replacement-rate funding came under pressure. The model produced a labor market with low long-term unemployment relative to comparable European economies, high labor force participation, and a workforce that moved between industries at a rate that kept skill mismatches from compounding. People who lost their jobs in the old industries had pathways, although imperfect, demanding, sometimes humiliating, to the new ones.

For an AI transition the lesson is structural, not Danish: income replacement without retraining produces dependency; retraining without income support produces impossible choices between

training and eating; both without the trust between employers, workers, and government that allows the first two to function produces neither. All three legs are required. Welfare should not only catch people after they fall. It should help move skills, capital, and identity toward the next productive structure.

A named program shows the shape of an answer without being one. In one respect the AI transition differs from the manufacturing transitions Denmark managed, and it makes the structural requirements harder, not easier. When a textile factory closes, the displaced skill is specific and locatable: the worker ran a particular machine, and the question of what retraining looks like is at least well-defined, even if the answer is difficult.

When AI changes the tasks performed by an analyst, a paralegal, a radiologist, or a junior programmer, the displaced skill is partly cognitive, and cognitive retraining is harder to define, harder to sequence, and harder to verify as complete than manual retraining. A worker who spent a career building pattern recognition in one field can't be retrained into another field's pattern recognition in six months, and a system that claims otherwise is measuring activity again rather than outcomes.

What the AI transition requires in addition to the three Danish legs is a fourth: honest accounting of what AI can do well, what it cannot, and which human judgments can't be replaced, for practical reasons rather than sentimental ones. A system that takes over routine office work but still needs judgments that depend on being in the room, knowing the people, knowing the place and carrying the blame creates demand for exactly those things, even as it removes the clerical jobs around them. The nurse who cannot be replaced by the diagnostic AI is the one who can read what is happening in the room and bear the consequence of being wrong about it. Training for that capacity, training that starts in schools, continues in institutions, and

is recognized and paid for in labor markets, is the retraining the AI transition requires, but it isn't what any country's training system is built to supply.

WHAT THE TOOLS ACTUALLY MAKE POSSIBLE

The better future is not one in which machines do everything and humans receive an allowance. It is one in which tools lower the cost of being capable. A person with a small business can access strategy, design, translation, accounting, and research. A village can model a microgrid. A teacher can adapt material to each learner. A citizen can question a policy. A family can preserve a worldview across generations while allowing descendants to argue with it.

Three examples. A teacher in a rural school with thirty students and one subject preparation period per day can use AI to generate and adapt materials for students reading at five different levels, in two different languages, across a subject she did not major in. She doesn't become a worse teacher by doing this. She becomes a teacher who can serve her actual classroom rather than the averaged classroom that the single textbook was written for. The AI is doing what a well-funded school district does by hiring specialists, and doing it for a school that couldn't hire them. Capability that used to belong only to the wealthy institution now belongs to the one-room school. The institution that gains the most from the change is the one that had the least to begin with.

Or consider the small business. A contractor in a mid-sized city knows her craft. She doesn't know her contracts, and has been signing them anyway for eleven years. She can now have a complex subcontractor agreement read back to her, have the clauses that quietly transfer risk onto her pointed out, and keep asking questions until she understands what she is about to sign. Before that tool existed, she either signed contracts she didn't fully understand, which is where most small contractors' business problems begin, or she paid a lawyer, which costs more than most small contracts are worth. For

the complex transaction the tool doesn't replace the lawyer. What it gives her is access to the lawyer's function for the transactions that were previously too small to be worth protecting. What it expands isn't legal sophistication. It's legal access.

Or consider a smallholder farmer in India cultivating two acres of cotton under a rainfall pattern that has become less predictable in each of the last fifteen years. Historically, planting decisions in that context were made by reading the sky, consulting older farmers, and following the regional averages that extension services published once per season. AI weather and soil modeling can now give that farmer variety-specific planting advice based on forecast rainfall, historical soil moisture, and the price of water at different times in the irrigation calendar, tools that used to be available only to large commercial operations with staff to commission and interpret them.

Even the crop insurance she can now buy is priced differently when the insurer has the same data, because the information gap that used to make insurers price against their own uncertainty, premiums too high and payouts too narrow to be useful, is smaller. She doesn't become an agronomist by using these tools. She becomes a farmer who can bring her knowledge of her specific land into dialogue, using a record she could not previously access, and the combination is more reliable than either alone. That is the extension that matters: her knowledge, connected to a record it couldn't reach before.

This is the largest possible expansion of agency. It is also the largest possible market expansion. Eight billion people with better tools can create needs, solutions, art, relationships, and enterprises that no central planner or corporation can anticipate. The economic value isn't in replacing all those people. It's in giving them the means to help more people, and the gain is largest for the people who currently have the least.

DIRECTION, NOT ARRIVAL

The future will still contain catastrophe, injustice, and grief. The direction has not always held. Civilizations have collapsed, not only declined. Technologies have enlarged cruelty as readily as kindness. Progress has been reversed, and the reversals have been long. What the mathematics of better has going for it is duration. It has been running long enough, in enough places, to constitute a pattern, and the pattern has structural reasons behind it. Systems that help more people survive, learn, cooperate, create, exchange, and correct error tend to generate more durable value than systems that depend on keeping most people weak or excluded.

The claim is not paradise. Nobody gets to Eden. The road just gets nicer, and so do the hotels along it. The claim is direction: the same mathematics that made survival more reliable, literacy more common, energy cheaper, and knowledge more accessible can now make judgment and coordination more widely available. Not guaranteed, not evenly, and not without people who are harmed on the way and whose harm the ledger has to count.

But in the direction of more people with better tools, more honest accounting, more accessible knowledge, and more visible consequences, the mathematics of better has been running for longer than any of the forces arrayed against it, and it is still running. Who it runs for is another question.

THE EQUATION

Two AI-literate generations x agency-preserving tools x inclusive transition = a plausible path from collective memory to collective better judgment.

PART IV

THE FLOOR AND NEXT

The ceiling shows what is possible. The floor shows whether civilization can carry it.

The fourth question is who it is for.

CHAPTER 15

The Floor Is the Measure

A civilization is durable when ordinary life is worth protecting

There is a line painted on the side of every merchant ship in the world, and the law says the sea must not close over it.

Life at the top has always been good. Harder to ask is what life feels like below, and for most of history nobody with the power to ask had reason to.

Civilizations are described by their ceilings. The tallest building, the fastest computer, the wealthiest citizen, the most powerful weapon, the greatest university, the largest company, the rarest achievement. Those numbers aren't vanity. They record what a system can produce when it is genuinely trying, and a society that produces nothing remarkable has usually stopped producing anything at all. They are also close to useless for predicting whether the thing will still be standing in a century.

A palace can rise above a hungry village. Repeatedly, in history it has, and for surprisingly long stretches. A dazzling technology sector can coexist with families who are one illness, one rent increase, or one missed paycheck away from losing everything. It can do so while shipping genuinely astonishing products the whole time. The ceiling represents possibility. The floor represents durability. Those are two different measurements, and they are confused constantly, usually by people standing near the ceiling.

For the mathematics here, the floor is the controlling variable. Not equality of outcome, which is a different argument with different evidence behind it, and not a promise that disappointment can be abolished.

The floor is the minimum practical platform from which a person can participate in tomorrow: food, shelter, safety, health, education, information, lawful agency, and enough stability to plan past the next emergency. A floor worth standing on creates time. Time creates choice. Choice creates vectors. Vectors are the only information this runs on.

THE LINE ON THE SIDE OF THE SHIP

In the early 1870s a British coal merchant and Member of Parliament named Samuel Plimsoll became interested in a question nobody wanted asked, which was why so many merchant ships were sinking.

A good number of them were being sunk, it turned out. Not directly, and not deliberately by anyone who thought of himself as a murderer. A vessel could be over-insured, loaded far past its limit, sent out with a rotten hull with too few crew aboard. If she failed to arrive, the owner collected. Sailors had a name for these ships. Coffin ships. Plimsoll published the evidence in 1872 in a book called *Our Seamen*, was sued by shipowners who regarded him as a menace to legitimate business, and kept going anyway.

What he finally extracted, in the Merchant Shipping Act of 1876, was a line. Yes, a line. A circle twelve inches across with a horizontal bar through it, painted on the hull where anyone standing on the dock could see it. It marked the depth past which that particular ship must not be loaded.

Ship owners don't become virtuous because of the line, and nobody has to care about a single sailor. It takes a hidden condition, how much cargo is too much for this hull, in these waters, in this season, and makes it a visible one, in public, visible at the moment the loading decision is being made. Then it makes the visible condition enforceable. It worked.

It also had a loophole. For the first eighteen years the law required a load line but didn't say where it went, so a certain number of owners painted theirs higher than a naval architect would have. Parliament did not regulate the line's position by law until 1894.[45] A floor the interested party gets to draw isn't a floor. It's just a decoration in the shape of one.

One number, for one hull, settled by hydrostatics: that's a ship's load line.

A civilization's floor is a dozen quantities at once, food, rent, a clinic within reach, a school that teaches something, a court that answers, and no equation settles where the mark goes. That's a real difference. The shape carries over: measure the minimum, publish it, and put the number where the people making the decision cannot avoid seeing it.

THE FEAR OF FALLING IS ITSELF A COST

A person can sit comfortably above any statistical poverty line and still live in continuous fear of going through the floor. Fear of that kind has a price, and the price shows up in behavior: shorter planning horizons, less risk taken, a coercive employer far harder to leave. Illness, housing, transport, and childcare can turn from problems into existential threats, and a person managing four existential threats is not going to start a business this year.

In 2013 four researchers, Anandi Mani, Sendhil Mullainathan, Eldar Shafir and Jiaying Zhao, published a study in *Science* that made this unusually concrete. They tested sugarcane farmers in Tamil Nadu twice: once shortly before the harvest, when money was tight, and once shortly after, when it was not. Same farmers. Same fields. Same tests.

The farmers performed measurably worse before the harvest than after it. The authors put the gap they measured at some thirteen IQ points, which they compared to the effect of losing a full night's

sleep.[46]

Since then the specific number has been argued over, as headline effect sizes in social science generally are. A test administered on two afternoons is not a life. What the study does show is that the deficit traveled with the *condition*, not with the person. The farmer didn't get cleverer in November. He got his attention back.

Scarcity is more than a shortage of money: it is a standing tax on the cognitive capacity of the person paying it, levied at the worst possible moment, and it produces exactly the short-sighted decisions that are later offered as proof that the tax was deserved.

Avoidable suffering is a human wrong, because suffering hurts. Reason enough, and it needs no economics attached.

But a wider ledger records a second cost, and refusing to look at it has never helped anybody. A sick person who can't get care isn't only suffering. They have been removed from working the field. A child sleeping somewhere unstable is not only frightened, but their future learning is being taxed in advance. And a family consumed by crisis has less attention for work, invention, care, community, or hope, and attention, as the sugarcane farmers show, is the scarcest input in the whole system.

Clean water, primary care, housing stability, basic education, public safety, reliable energy and usable knowledge are productive infrastructure in the same sense that a road is. Every solved floor problem releases human attention upward.

The ceiling is evidence of what is possible. The floor is evidence of what will last.

EIGHT BILLION REPEAT CUSTOMERS

What could possibly be more sustainably profitable than eight billion housed, healthy, literate repeat customers?

Not eight billion identical people. Not identical incomes, not equal boats, not a world sanded flat. A deeper harbor.

Somebody healthy, housed, educated and connected isn't merely someone who has stopped needing help. She can make something, teach something, buy something, fix something, start a company, solve a problem belonging to somebody else, or ask a question nobody had thought to ask. Raising the floor doesn't require lowering the ceiling. It increases the number of people with any chance of reaching it.

And the mathematics runs the other way too, which is the part the fortunate tend to skip. Where large numbers of people believe they have nothing much to lose, a society has to spend more guarding what the rest possess. More walls, more suspicion, more emergency response, more political fury, more of the destructive kind of competition. None of that is free. Security bought only at the ceiling is expensive because the floor underneath it is moving.

Peace is nice. Cooperation is nice. Helping people is nice.

Nice isn't required by the calculation. It's great to be righteous. But the reason most people are righteous is that it has stopped being a sacrifice. It feels better, and it's more profitable.

War destroys customers, producers, infrastructure, capital, trust and every transaction that would have happened afterward. Preventable disease removes capability. Poverty strands it. Unstable housing consumes it in survival. Bad information corrupts decisions made with it. On the other side of the ledger: education creates capability, health restores it, housing releases it, cheap energy multiplies it, truth improves what gets done with it, and peace preserves the lot long enough to compound.

None of those require anybody to be good. They require only that the accounting be complete, which is a lower bar and, inconveniently, a rarer achievement.

Technology changes the price of compassion.

A service that once needed a scarce expert can become widely available when software, automation, better logistics or cheaper energy cut its marginal cost. Moral arguments for providing it may be 4,000 years old. Economic ones can be four years old. What was charity in one decade becomes, in the next, simply the cheapest way to run the system, and the second argument moves institutions the first one had been failing to move since Deuteronomy.

A real raise is increased command over a decent life, whatever the pay stub says. If energy, health care, knowledge, transport, housing, communication and manufactured goods each require fewer hours of work to obtain, purchasing power rises. Owners gain as well, and not as a consolation prize: a healthier, wealthier, more capable population is a larger market, and it is a market that can keep transacting for another forty years.

THREE WAYS A FLOOR FAILS

First, a floor can be built that traps people in poverty. Support that vanishes the moment someone earns a little more creates a wall precisely where the ladder was supposed to start, and the person facing an effective marginal tax rate near a hundred percent isn't lazy for declining to climb it. They've done the arithmetic correctly. Badly designed floors have produced this outcome in many countries, repeatedly. One of the more durable stupidities in public policy is pointing at the resulting behavior as evidence about the poor, rather than as evidence about the design.

Second, a floor is an instrument, and instruments get picked up by whoever is nearest. A floor supplied by a single provider who also owns the housing, the store, and the means of leaving isn't a floor. It's a fence with amenities.

Third, money. Somebody pays. A floor isn't free. Neither is the current arrangement, and its costs are booked somewhere that doesn't appear on the invoice: in emergency rooms, in prisons, in guarding, in the compounding loss of everything the trapped participant would otherwise have made. Which ledger the bill lands on is the argument, not whether there is one.

We know how to measure a ceiling. We have never once agreed on where to paint the line.

CHAPTER 16

Something to Lose

Hope becomes durable when people possess a life they want to protect

"He that hath wife and children hath given hostages to fortune; for they are impediments to great enterprises, either of virtue or mischief."

Francis Bacon, "Of Marriage and Single Life," Essays

Hope is lovely. Something to lose is structure.

Nobody has to become virtuous for a society to be durable, which is lucky, because the virtue supply has been unreliable for 4,000 years and shows no sign of improving.

It depends on something far more ordinary: on millions of people, and eventually billions, waking up to lives they are invested in keeping. A home. A family. A trade. A friendship. A neighborhood, a business, a reputation, a pension, a school, a trip booked for August, a garden halfway planted, a child who is doing well. These are stakes in tomorrow, and they change the mathematics of destruction more reliably than any sermon ever has.

A person who believes tomorrow contains something valuable faces a genuine cost when destruction threatens it. So does a community. So does a nation. The best weapon is not the strongest security. The strongest security is to have fewer people who calculate that they gain more from firing a weapon than from waking up peacefully the next morning.

THE TWO SETTLEMENTS

The twentieth century ran this experiment twice, on the same country, within thirty years, and published the results.

After World War I in 1918, the victors settled with Germany by extracting. The reparations figure moved around for years, fixed, revised, rescheduled, argued over by commissions, but the shape of the arrangement never changed: a defeated population was to spend the next several decades transferring value outward, under terms it had no part in setting.

By the mid-1920s the payments were being made largely with borrowed American money, which is a way of saying they weren't really being made at all. When American lending stopped in 1929, the whole arrangement collapsed, and took the political settlement with it.

After 1945, the same victors did close to the opposite. The European Recovery Program ran from 1948 to 1952 and moved roughly $13 billion into Western Europe, about 90 percent of it as grants rather than loans.[47] Germany was a recipient. Seven years after a war it had started and lost, the defeated country was being handed capital.

The Marshall Plan wasn't charity, and nobody involved pretended it was. There was a Soviet Union nearby, and a strategic argument for a prosperous, non-Communist Western Europe that had nothing to do with generosity. Economists have argued for decades about how much of Europe's recovery the money actually caused; the more careful accounts give a good deal of the credit to institutional reform, to the removal of trade barriers, and to the simple fact that Europe still had its skills and its habits even after it had lost its buildings. Thirteen billion dollars wasn't, in the end, an enormous sum against the size of the economies receiving it.

And we don't get the counterfactual. Nobody could have run a control Europe. The first settlement left a large population with very little to protect and a permanent grievance to organize around. The second gave a large population a reason to want the arrangement to last. The first ended in the worst war in human history. The second produced the longest peace that part of the world has ever had. Not proof. Two data points and a mechanism, which is more than most of political philosophy runs on.

GRATITUDE IS MEMORY OF A WORKING FLOOR

Hope is usually described as facing forward. Durable hope also faces backward, and the backward-facing part is the one that gets neglected.

A person who wakes up expecting clean water, electricity, emergency medicine, enforceable contracts, food in shops and information on a screen is living inside achievements produced by generations whose names are almost entirely unknown. None of it was inevitable. Most of it was argued over bitterly by people now dead, and a fair amount of it was opposed at the time by serious men with serious reasons.

Gratitude for that is accurate accounting, not obedience, and accurate accounting is the only thing that reliably tells you what you'd be giving up.

Achievement produces something worth preserving. Preservation lengthens horizons. Longer horizons support investment. Investment creates capability. Capability can raise the floor again. The loop isn't automatic and nothing enforces it, but once it's running, a society has something it didn't have before: a constituency for continuation.

FOURTEEN HUNDRED YEARS OF THE FAMILY BUSINESS

In 578, a carpenter from a Buddhist kingdom on the Korean peninsula arrived in Japan at the invitation of Prince Shōtoku, who wanted a temple built and had no one locally who knew how to do it.

The man's name was Kongō Shigemitsu. The firm he founded to build Shitennō-ji was still operating, under his descendants, in the 21st century.

Kongō Gumi built and repaired temples for something over 1,400 years, through forty generations. It survived civil wars, the Meiji Restoration, the collapse of the shogunate, the firebombing of Japanese cities, and several centuries in which its principal product was in and out of official favor. When the family ran short of male heirs it adopted them, or promoted sons-in-law, or on at least one occasion put a widow in charge. The rule was never that the blood continue. The rule was that the building continue.

Long-lived family firms learn a peculiar kind of accounting, and Kongō Gumi is its extreme case. Whoever is running it is never the only stakeholder. Inventory, reputation, customer relationships, property, technique, debt, habits, mistakes and opportunity all arrive from people who worked before, and decisions made this morning will be lived with by people not yet in the room. Nobody ever gets to hold the closing meeting.

I know a small version of that accounting from the inside. For eighteen years I sat on the board of an arts center in Boulder that I'd helped design, build, and pay for. Now I send a check once a year. When I walk into the center, it's crowded and happening. It doesn't need me. That's the only kind of success I've ever fully trusted, and it's the kind you can't take with you.

Civilization behaves similarly, on a larger scale and with a longer list of assets: language, agriculture, medicine, law, science, infrastructure, art, markets, institutions, accumulated error, memory. We receive them without having earned the original capital. We modify them merely by using them.

Two things, though.

First, continuity isn't sameness. A healthy family enterprise retires products, changes suppliers, adopts machines, closes premises, enters markets and occasionally sells the building. Kongō Gumi did all of it. Freezing the arrangement was never the point of an inheritance. Preserving enough capacity to keep improving it is.

Second, the story doesn't end well. In January 2006, after fourteen centuries, Kongō Gumi's assets were bought by a larger construction group.[48] What finished it wasn't war, or the loss of the craft, or a failure of the sixth-century business model. It was debt taken on during Japan's property bubble, in the 1980s, by the fortieth generation. A firm that had survived everything the medieval world could throw at it was undone by a decade of easy credit and a bet on land.

It's still operating, as a subsidiary. Which is the right ending for the parable, because it's what actually happens: the enterprise continues, the independence doesn't, and the generation that lost it was in most respects the most sophisticated one the family ever produced.

WHAT STAKES DEFEND

There is a cost to giving people something to lose.

People who have stakes defend them. That's the whole point, and it's also the problem. The homeowner with equity is exactly the person who turns up to object to the building of more homes. The professional with a license is exactly the person best organized to make the license harder to get. And the worker in a mature industry is the person with most to lose from the transition that would raise everyone else's floor, and he isn't wrong about his own position. He has a stake in tomorrow, we told him to acquire one, and he is protecting it.

So "something to lose" doesn't produce a population of peaceable improvers. It produces a population that's peaceable about some things and immovable about others, and the immovable set is the set

with the most to lose.

I don't have a clean resolution. What can be said is that this is a far better problem than its opposite. A society of people defending what they have argues, lobbies, litigates, delays, and infuriates its reformers for a generation at a time. A society of people with nothing to defend does something else, faster, and the something else doesn't stop when the argument is over.

Both are expensive. Only one of them compounds.

Hope is the feeling. Something to lose is the architecture.

THE EQUATION

A life worth protecting x credible opportunity tomorrow + inherited capacity = a larger peaceful stake in continuation.

CHAPTER 17

Nature Put Everything in the Box; Man Found Good and Bad

Energy, water, and the physical floor

Four-fifths of every breath you take is nitrogen. For all but the last century of human history, not one atom of it was any use to you.

Civilization doesn't negotiate with chemistry. It learns to arrange it, and the arrangement is usually the whole of the difficulty.

Before any policy sits on it there is a physical floor. People need breathable air, usable water, food, shelter from temperature, and energy capable of moving matter and information. Every civilization there has ever been is an arrangement inside those limits, and none of them got a vote on the limits.

What's remarkable is how generous the box turned out to be.

Nature supplied an absurdly rich set of ingredients and a set of laws for combining them, and then declined to supply any instructions. Almost never has the constraint been that the necessary thing doesn't exist. Usually it's that the useful arrangement of it hasn't yet become cheap, scalable, reliable, or available to anyone who isn't already rich. Different kind of problem, and the kind that gives way.

THE AIR WAS ALWAYS 80 PERCENT FOOD

For the whole of human history until about a hundred years ago, the binding constraint on how many people could eat was nitrogen.

Plants need it. They can't take it from the air, where it sits in enormous quantity, roughly four-fifths of every breath, locked in a molecule so stable that almost nothing in nature can break it apart.

So agriculture depended on the small trickle that lightning and certain bacteria could fix, plus whatever manure could be carried to the field. That trickle set the ceiling on the population, and the ceiling was hard.

Which is why the nineteenth century went looking. Europe stripped the guano islands off Peru. It fought over Chilean nitrate deposits. A serious body of educated opinion held that the industrialized world was on a path to famine, and it wasn't obviously wrong. Malthus had said as much in 1798.

Then in 1909 Fritz Haber, a German chemist, working with a table-top apparatus, pulled ammonia out of the air under extreme pressure and temperature over a catalyst, and Carl Bosch spent the next four years turning the demonstration into a fertilizer industry. Estimates vary and the accounting is genuinely hard, but the range isn't really in dispute: something in the region of half the nitrogen in the food eaten today was fixed by that reaction. Canadian scientist Vaclav Smil's version of the calculation is that without it, roughly two-fifths of the people currently alive would not be.[49]

Nothing new entered the world. The nitrogen had been in the air the whole time, in front of everybody, for the entire history of hunger. What was missing was the arrangement.

Haber proposed chlorine as a weapon and stood on the ground outside Ypres, in Belgium, in April 1915 while it was released. His wife Clara Immerwahr, a chemist herself, killed herself shortly afterward. He was awarded the Nobel Prize in Chemistry for 1918, and a substantial portion of the chemical community objected. Later, having been driven out of Germany for being Jewish, he died in exile from a country whose army he had served with more enthusiasm than it ever repaid.

Then there's the bill the process itself carries. It consumes on the order of one to two percent of the world's total energy supply, most of it from fossil fuels, and it accounts for something near one percent

of global carbon dioxide emissions. The nitrogen that does not stay in the crop runs off into rivers and coastal water, where it produces the largest single human perturbation of a planetary nutrient cycle.[50]

So the box was full and nothing was labeled. Nature put everything in the box. Food, hydrogen, oxygen, sunlight, wind, gravity, heat, emissions, uranium, minerals, membranes, chemistry, and the laws that govern them were already here. Intelligence adds no new ingredients to nature; it discovers arrangements that let the ingredients do useful work. That's a statement about availability, not about safety, and it comes with no assurance that the arrangement that solves your problem won't create somebody else's.

WATER IS STORED AGENCY

Clean water does more than prevent disease, and the disease prevention is only the visible half.

It returns hours. Hours otherwise spent carrying, boiling, queueing, searching, and caring for people who got sick from the alternative. A reservoir, a pipe, a membrane, a pump, a treatment plant, a recycling system, a desalination plant: each is an agency machine. Its output is measured in liters. Its civilizational value is measured in what people do with the part of the day that the water problem used to occupy.

Water systems need energy. Energy systems need water, and land, and materials. Running them as separate policy silos hides the relationships a wider ledger considers, and it's the reason so many plans assume only one will happen. Which technology will win the argument is never the design question. The question is which setup gets people the water and power they need while asking least of everything else, for longest.

THE ELECTRON NOBODY IS USING

Renewable energy generation creates a novel and rather pleasant problem: at certain hours the system produces more electricity than the grid can immediately absorb.

Nothing about the scale is theoretical anymore. In 2024, California's grid operator curtailed about 3.4 million megawatt-hours of utility-scale solar and wind, a 29 percent increase on the previous year, of which roughly 93 percent was solar. April was the record month, at 1.46 million megawatt-hours, equal to about 18 percent of everything the state's grid-scale wind and solar produced that month.[51]

That's close to a fifth of one month's clean generation, produced and then thrown away, in the most aggressively renewable large economy in the United States.

Surplus has value if flexible demand can be moved toward it. Batteries store it. Water can be pumped, treated or desalinated. Heat and cooling can be shifted by hours. Vehicles charge. Computation can be scheduled. Industrial processes can be run when the electron is nearly free rather than when the shift pattern says. *Curtailment capture* is the habit of asking, every time, what useful work could be moved into the hours when the alternative is throwing the energy away.

I'm a curtailment capture guy, and part of the appeal is the irony. Carbon capture, in my judgment, mostly doesn't work. This does. Tesla now sells its grid batteries as Megablocks, four Megapacks and a transformer wired together, twenty megawatt-hours in one unit, and the trucks carrying them go off quietly in every direction. Elon Musk, the greatest showman in the business, who has never understated anything, says batteries like these could roughly double the electricity the American grid delivers in a year without building a single new plant, because the grid is built for a peak of about 1.1

terawatts and averages about half that. Cut his number in half and it is still enormous, and Tesla isn't alone; modular batteries are getting cheaper and smaller across the industry.

Smaller matters as much as bigger. At Wild Bear, a project I'm working on that is becoming a sustainability hub, there are 50 kilowatts of solar on the roof, the grid, and a set of electric vehicle chargers. One of them is never used. It is kept free for police and fire. With a battery behind it, a fire truck or anybody else can charge whatever the grid is doing.

The marginal electron isn't free, and pretending otherwise is how enthusiasts lose arguments they should win. Panels, wires, storage, land, maintenance and financing all exist and all have to be paid for. But once those assets are built, the cost of using an electron that would otherwise be curtailed is a very different number from the cost of building new generation to do the same job. Timing turns into an economic resource, a new thing to have.

THANK YOU, AND MOVE ON

Fossil fuels were a civilization-changing arrangement.

They concentrated several hundred million years of solar energy into portable deposits and made possible an industrial scale no earlier society could reach. Nearly everything in this book that got measurably better between 1850 and 1970 was, at some point in its causal chain, burning something. Pretending otherwise isn't moral seriousness. It's bad accounting with good intentions.

Whether combustion was best was never the question. What matters is whether routine combustion remains the best available arrangement once atmospheric consequences, water demand, rising global energy demand and the falling cost of the alternatives are all inside the same ledger.

Nuclear energy goes through the same test, and it should be run without the fear that it often evokes. So do carbon capture, geothermal, storage, transmission, efficiency, and less mature technologies. No technology gets a moral exemption and none is banned by rhetoric. The ledger is physical: what does it deliver, at what scale, with what energy, materials, water, land, waste, reliability and opportunity cost?

WHAT IS NOT IN THE BOX

There's a category of thing the box doesn't hold.

Phosphorus is the clearest case. Every living organism needs it, no organism can substitute for it, and unlike nitrogen it cannot be pulled out of the atmosphere, because phosphorus has no meaningful gas phase. There's no breakthrough waiting to be invented for it. The only route to commercial quantities is mining phosphate rock, and roughly 70 percent of the world's reserves sit in a single country, Morocco.[52]

Helium is similar in kind and worse in temper: produced by radioactive decay over geological time, essential for MRI machines and semiconductor manufacture, and light enough that any of it released simply leaves the planet.

Those are the real limits, and they aren't the ones people usually name. Rarely is the genuine constraint raw energy or raw material. It's the small set of substances for which no arrangement will do, together with the political fact that they are unevenly distributed and everybody needs them.

Which leaves a physical floor that is more generous than the pessimists claim and less generous than the enthusiasts assume, with a short list of exceptions that deserve rather more attention than they get.

Nature put the ingredients and the laws in the box, left out the instructions, and did not sort them by consequence. Intelligence discovers the arrangement. Memory makes the next arrangement better.

CHAPTER 18

The Public-Health Economics of Peace

How civilization can preserve security while converting reciprocal threat into productive capacity

"Right, as the world goes, is only in question between equals in power, while the strong do what they can and the weak suffer what they must."

Thucydides, History of the Peloponnesian War, the Melian Dialogue, trans. Richard Crawley

The worker is not the weapon, and security is not the arsenal.

The one handle on Eisenhower's machine was a narrower definition of defense: not the largest destructive capacity a nation can afford, but the minimum credible capacity that makes aggression fail. Add the rest of the sum: a rising floor gives people something to lose, abundant energy widens the range of productive alternatives, and better information makes consequences easier to see. So the question can be run again, this time as an economic test: if civilization can preserve security with less reciprocal destructive capacity, what happens to the people, factories, capital, engineering, and knowledge presently assigned to producing the excess?

First accounting correction, and it's simple. A machinist employed by a missile program isn't a benefit unique to missiles but productive capability. So are the engineer, the software team, the research laboratory, the supplier network, the factory, the machine tool, and the pool of investment capital. Counting those assets as though they disappear when the product changes assumes that productive capacity belongs to the weapon rather than to the people and institutions that learned how to make it.

That assumption creates a false choice, weapons jobs or unemployment. A defense engineer redirected toward grid resilience has not been economically erased. A precision-manufacturing plant converted to power electronics, water equipment, robotics, medical devices, transportation systems, communications, or resilient infrastructure has not vanished. The assignment changed. Much of the accumulated competence survived. Retooling costs money, and some skills don't transfer cleanly, but those are transition costs, not proof that destructive output is the only valuable use of the asset.

TWENTY THOUSAND WARHEADS, AND THE LIGHTS STAYED ON

None of that is hypothetical, though it sounds it.

In February 1993 the United States and the Russian Federation signed an agreement with an unlovely name and an extraordinary content. Russia would take five hundred metric tons of weapons-grade highly enriched uranium, material drawn from roughly twenty thousand nuclear warheads, blend it down into low-enriched uranium suitable for civilian reactors, and sell it to the United States. The first shipment arrived in June 1995. The program ran for two decades and finished in December 2013, having converted the full five hundred tons into more than fourteen thousand tons of reactor fuel.[53]

For most of those twenty years, something approaching a tenth of all electricity generated in the United States came from fuel made out of Soviet warheads.

Consider the arrangement rather than the sentiment, because the sentiment is the least interesting part. Nobody underwent a moral conversion. Russia needed hard currency badly and had a great deal of a commodity it no longer wanted to store. American utilities needed enriched fuel and didn't much mind where it came from. The transaction worked because both parties were pursuing their own

interest, and it survived twenty years of deteriorating relations between the two governments because commercial contracts are considerably more robust than goodwill.

What it establishes is exactly the accounting point. The asset was never the warhead. The asset was several decades of accumulated enrichment capacity sitting inside a metal shell, and it turned out to be worth more as fuel than as threat, to the seller, in cash, immediately.

It's no template. One-off, conducted with a collapsing superpower in urgent need of money, over material already withdrawn from deployment; it did not disarm anybody who did not want disarming. More than once, the price nearly killed it. And there were people who paid for it. The flood of cheap blended-down material depressed the uranium market for years and hollowed out parts of the domestic enrichment industry, a cost that belongs on the ledger rather than in a footnote.

But it happened. Five hundred tons, twenty thousand warheads, into the grid, during twenty years, for money. Any argument that industrial conversion at civilizational scale is a fantasy has to explain that away first.

KEEP WHAT IS YOURS

No theory of uniquely bad nations is needed to explain the modern weapons system. Twentieth-century American-Soviet rivalry built a security architecture in which each side could make an irrational case for capabilities, which the other side then had a rational reason to answer. Defense of its larger population led China to develop another major military-industrial capacity.

Alliances, neighbors, smaller states, and manufacturers respond to the capabilities around them. What results is a recursive security dilemma: because you possess it, I need it; because I acquire it, you need more.

Worth identifying, and it has no villain in it. Individual decisions can each remain perfectly understandable while their combined result becomes expensive and dangerous. Blame doesn't improve the calculation. Widening the field of view does. What matters is whether a country can stay as safe while building less of what only answers the other side's building.

What a lower-weapon world needs isn't complicated, even though building institutions that can enforce it is hard: keep what is yours; do not take what is theirs. It is the bound on the man walking up his neighbor's path, at the size of a country. Borders can still be disputed through law, negotiation, referendum where lawful, arbitration, and diplomacy.

For a president, premier, general, or citizen to accept substantial reductions in redundant weapons, the bargain has to be credible in both directions. Your population, territory, and ordinary life are protected if you leave other populations and territories alone. If a state attempts conquest, the wider system acts quickly enough that aggression does not pay. That response can begin with verification, financial isolation, trade consequences, denial of military advantage, support for the attacked population, and collective defensive force when force is necessary. Its purpose isn't revenge against what a leader loves. Its purpose is to make the aggressive act fail.

So the agency of 8 billion people is best understood as a backstop, not a clenched fist waiting to strike. Its strength comes from universality. The rule must apply to friends and rivals, large states and small ones, wealthy populations and poor ones. A collective security principle that excuses conquest by the powerful merely recreates the old mathematics under a new name.

Keep what is yours. Do not take what is theirs. Make conquest a reliably losing transaction.

Here civility becomes more than courtesy. A civil system can distinguish the opponent from the relationship, resist an act without requiring permanent hatred of the people on the other side, verify compliance, preserve records, reopen trade, revise an agreement, and allow yesterday's enemy to become tomorrow's partner. Those capacities reduce the amount of security that has to be purchased through permanent reciprocal threat.

That humanity can even discuss such an arrangement is evidence of the path already traveled. In the cave, there was no treaty inspection regime, no satellite verification, no international banking system, no global communications network, no shared scientific record, and no near-instant ability for billions of people to see an event. None of those tools guarantees wisdom. Together they change what is technically possible. Civilization can know quickly that a border has been crossed, coordinate more broadly in response, and remember more accurately what happened afterward.

CONVERSION AS PUBLIC HEALTH

A mature society routinely spends money now to avoid paying a much larger recurring cost later. Vaccination is the clearest example. The vaccine is an expense, but the disease is more expensive. Clean-water systems, sewage treatment, fire prevention, building standards, disease surveillance, and lead removal follow the same logic. They don't punish yesterday's victims or villains. They change the probability and cost of tomorrow's harm.

Industrial conversion can be treated similarly. If improved defensive technology, credible border protection, verification, collective refusal to reward conquest, and stronger international coordination make some offensive capacity unnecessary, then helping defense manufacturers convert isn't charity and isn't a fine on successful industry. It is a public-health expense for the global system. Government can preserve employment, finance retooling, retrain workers, help suppliers migrate, and give investors a credible path

into replacement markets. Society would be paying to reduce the friction of changing what an extraordinary productive system produces.

Energy shows how an arrangement can lose market share because a better alternative gets cheaper, not because everyone involved undergoes a moral conversion. In its 2025 cost review, the International Renewable Energy Agency reported that more than 90 percent of newly commissioned utility-scale renewable projects produced electricity below the cost of new fossil-fuel alternatives, while renewable generation avoided hundreds of billions of dollars in fossil-fuel costs.[54] Storage, transmission, materials, land, financing, reliability, and flexibility remain real constraints. Abundance is not infinity. What has changed is that the strategic resources increasingly become coordinated electrons rather than a continually purchased combustible scarcity.

As useful energy becomes cheaper, cleaner, and easier to coordinate, more human capability becomes affordable. Water can be purified or moved. Buildings can be heated and cooled. Machines can work. Transportation can electrify. Information can travel. A factory converted from destructive output doesn't need to enter a market of charity. It can enter enormous markets for energy, water, housing, transportation, robotics, medicine, communications, and resilience.

Here peace and prosperity become the same argument. A higher floor means more people able to take part. More capable participants mean more workers, customers, inventors, entrepreneurs, teachers, caregivers, scientists, artists, and problem-solvers. Whoever is at the bottom is an unrealized participant, not merely a cost center, and whoever is at the top isn't automatically opposed to raising that floor. Owners of productive capital can benefit enormously when the market becomes larger, more stable, and able to transact for longer.

So the question for a civilization isn't mainly how much value should be taken from the ceiling and moved to the floor. It's what arrangement raises the floor while expanding the ceiling that could be reached. Housing, health, energy, education, truthful information, transportation, peace, and lawful security can set free people whose abilities are going to waste. Those people become customers and competitors, labor and invention, demand and supply. It makes MORE possible.

Whether informed good can become more profitable than bad depends on complete enough accounting. Bad often looks profitable because the ledger stops early. A missile sale records revenue but may omit the resulting destroyed bridge, lost transactions, displacement, reconstruction, retaliation, fear, and human capability removed from the system. A fossil-fuel transaction records the fuel price while some health, environmental, and security costs land elsewhere. And a weapons job is counted as a benefit, while the worker's alternative productive employment is treated as though it never existed.

FEASIBLE IS NOT INEVITABLE

Some questions sound moral until you run them as engineering. Can children everywhere go to sleep without hunger? Humanity already produces more than enough calories to feed everyone. Hunger remains because food, income, access, transport, conflict, prices and politics don't line up for everyone at the same time.[55] So the physical question and the practical one are different. If feeding everyone is physically possible, the useful question is what keeps the possible from becoming ordinary.

Peace deserves the same discipline. It doesn't have to be proven inevitable to be shown feasible.

Expand the accounting boundary and the result can change. A destroyed bridge creates replacement demand. A functioning bridge creates commerce. Both appear as economic activity. Only the second preserves and expands capability. Which is why gross activity isn't

enough, and why the mathematics of civilization needs a wider ledger that asks what survived, what was created, what had to be replaced, what harm was avoided, who retained or gained agency, and how long the resulting value can compound.

EARTH VALUE

Useful output + productive capacity created + harm avoided + security preserved - resources consumed - external costs - value destroyed.

War should therefore be examined like a preventable disease, without pretending that every conflict is preventable or that deterrence has no value. Ask instead whether some portion of present destructive capacity can be converted without reducing the minimum security necessary to make aggression unsuccessful. Whether all military expenditure is bad was never the question. When the answer becomes yes, conversion belongs on the prevention side of the ledger.

A vaccine doesn't defeat evil inside the virus. It changes the conditions under which disease can spread. A better security architecture need not identify and morally reform every dangerous leader. It can change the payoff structure around aggression. Secure borders lower the return to conquest. Renewable abundance lowers strategic dependence on scarce fuels. Broader prosperity gives more people something valuable to preserve. Larger markets make peaceful participation more profitable. Truth infrastructure makes manipulation easier to detect. Persistent memory makes it harder for yesterday's failure to disappear before tomorrow's decision.

Here the recursive formula is being run at its largest available scale. Inherited: national arsenals, alliances, military industries, borders, fears, treaties, factories, and memories of actual invasion. New evidence changes what defense technology can do and what productive alternatives are available. Correction changes the assignment of resources. Agency expands when more people gain

energy, health, housing, knowledge, security, useful machines, and trustworthy information. Expanded agency then generates the next set of questions.

The recursion requires no permanent authority deciding what's good for everyone. It does require enough shared evidence to show which arrangements build up what people can do and which keep using it up. Was the weapon necessary for deterrence? Could the same security have been achieved with less destructive capacity? What happened to worker income after conversion? Did the factory enter a larger market? Did cheaper energy increase productive participation? Did raising the floor diminish the ceiling or enlarge it? What did the consequences teach us? What is the better next question?

> *The achievement is not that humanity built weapons. The achievement is that, thousands of years after the cave, humanity can ask whether it still needs so many of them.*

If a country can stay safe while turning weapons factories into ordinary factories, the change doesn't have to make anyone poorer. It can create a more efficient use of the same inheritance. Engineers remain engineers. Factories remain factories. Capital remains capital. Security remains a public function. What changes is how much of the world's extraordinary ability to make things goes into destroying them, and how much into the market where value can grow.

THE EQUATION OF CONVERSION

Preserved security + industrial conversion + abundant coordinated electrons + more capable participants = larger markets with less value destroyed.

CHAPTER 19

When Capability Outruns Obligation

Robots, private money, and the statutory gap

For thirty-five years, what time it was in the United States was settled by an industry nobody had given the power to settle it.

Consequential risks don't always come from a dangerous technology. They may come from two useful technologies discovering each other.

Humanoid robots and private digital money are each interesting on their own and neither is obviously alarming. A token can be created at almost no technical cost. A robot converts capital, energy, software, and materials into labor.

Put them inside a single vertically integrated market and something appears that was in neither of them: the owner of the productive machines may also control the unit of account required to buy what the machines make.

Not a technological problem. A category problem, and category problems are the ones law handles worst.

THE DAY THERE WERE TWO NOONS

On Nov. 18, 1883, at noon, the time changed for a large part of North America.

Until then, time was local and honest: if the sun stood overhead, that was twelve o'clock, and the town down the line was four minutes behind you. That worked for centuries and became unworkable the moment people were able to move faster than a horse. Railroads ran on dozens of incompatible times. A station in a busy junction city might display several clocks, each correct for a

different line, and a passenger changing trains had to do arithmetic to avoid dying in a collision that timetables were supposed to prevent.

Previously Congress had looked at national standard time in 1809, in 1870, and again in 1872, and had declined to enact it each time.

So the railroads did it themselves. They divided the continent into zones, agreed on the boundaries among themselves, and on an appointed Sunday simply began using the new time zones. Cities that resisted found that the trains didn't care. Within a few years the country had reorganized its clocks around a schedule set by private companies for their own operational convenience, and it worked, and almost everybody's daily lives quietly rearranged themselves to fit.

Congress ratified the arrangement in the Standard Time Act of 1918.[56]

Thirty-five years. For thirty-five years, time in the United States was a fact determined by an industry, exercising a power nobody had granted it, over a matter Congress, the public authority, had three times refused to touch. The railroads weren't villains. They had a real coordination problem and they solved it well; the zones they drew are close to the ones in use now. The railroads stepped into an empty space, did a public job well, and the law caught up a generation and a half later.

Time zones, though, are non-rival: the railroads gained nothing by making anyone else's clock worse, and once the standard existed everybody could use it for free.

BILLIONAIRE BUCKS

Call the thought experiment Billionaire Bucks. A generic name on purpose; it doesn't say any one person is running such a system.

An AI token becomes consequential only when its issuer also controls something people cannot easily do without.

So imagine a highly automated service that becomes a genuine bottleneck, in manufacturing, logistics, farming, construction, elder care, or any other part of ordinary life that stops being optional once everyone else has stopped doing it by hand. If one provider reaches dominant capacity and requires its own token for access, demand for that token is no longer speculative. It has been engineered into the productive bottleneck. Production feeds the money and the money feeds production, and the loop can finance its own expansion without ever going to a bank.

A corporation doesn't become a country by issuing a token. Legal sovereignty stays with states. But economic sovereignty has intermediate forms, and the intermediate forms are where people actually live.

What matters is exit cost. How expensive is it to leave and carry on an ordinary life somewhere else? Every other question, market share, valuation, whether the founder is charming on stage, is downstream of that one.

PRIVATE DIGITAL SEIGNIORAGE

There's an old word for what an issuer in that position collects, and it's worth recovering, because the word carries an obligation the modern version has quietly dropped.

Seigniorage is the profit a sovereign makes from making money. A mint took in bullion, stamped it, and handed back coin worth more than the metal and the labor that went into it; the difference belonged to the seigneur, which is where the word comes from. It was real revenue and often a large one. It was also never free, because the issuer who took it carried the other half of the arrangement: keeping the standard honest, standing behind the unit when somebody presented it, maintaining the courts in which debts denominated in it could be enforced, and answering to a population

that had no practical way to hold anything else. Debase the coin and the consequences arrived at the issuer's own door: in the price of everything, in the cost of the next loan, and occasionally in person.

Some forms of private digital money can create a version of the same capture while carrying less of the public obligation historically attached to sovereign money. The mechanisms differ. A speculative token, a regulated stablecoin, a loyalty point and a privately issued unit required to enter a productive network are not economically identical. The useful question is functional: what value is being created, who captures the return from issuing the instrument, what obligations come with that privilege, and who bears the cost if confidence or convertibility disappears?

That is the ledger test. If a money-like instrument draws real savings, captures seigniorage, or becomes necessary for access to productive activity, the transfer should become visible enough to measure. Better forensic tools will make some of those relationships harder to hide behind vocabulary. The point is not that every crypto asset is counterfeit. It is that a private claim functioning as money should not become invisible simply because the noun is new.

None of that needs cryptography. The clearest working example runs on airline miles.

An airline mile is a private currency issued by a company. It is created at a keystroke, sold in bulk to a bank, handed to a passenger as a reward, and redeemed later at a rate the issuer sets and can change without asking. In June 2020, with its aircraft parked and its revenue down something near ninety percent, United Airlines went to the credit markets with the most reliable asset it had left. It moved the MileagePlus brand and its member data into two bankruptcy-remote subsidiaries and pledged them, together with the accounts into which the program's revenues are paid, against a five-billion-dollar term loan facility.[57] The offering materials then disclosed what the program actually was: 5.3 billion dollars of cash from mile sales

in 2019, 1.8 billion of that as earnings, and seventy-one percent of the cash arriving not from flying anybody anywhere but from selling miles to a bank. United valued MileagePlus at $21.9 billion, twelve times those earnings, comfortably more than the stock market was then willing to pay for the whole airline. The financing closed the following month at $6.8 billion.[58]

Here's the sequence in the order it happened. A company printed a currency, sold most of it to a bank, and then borrowed against the nonexistent printing press for more than the market thought the underlying business was worth. Nobody called it seigniorage. Nobody needed to. Every step was disclosed, audited and appeared lawful, and the arrangement is a good deal more transparent than most of what has been built on a blockchain since.

Nobody is paid in airline miles. Nobody has to hold them in order to eat, and a passenger who concludes that the program has stopped being worth anything walks away at the cost of a worse holiday. A mile is a claim on one company's spare seats rather than a unit of account. Exit cost measures the whole distance between a loyalty scheme and a currency, and exit cost is precisely what closes as the issuer's product stops being optional.

Which leads to the question to put to any such arrangement. Not whether the instrument counts as a security, which is a question about paperwork, but whether the issuer has arranged to collect the seigniorage of a currency while carrying only the obligations of a coupon.

THIS HAS BEEN RUN BEFORE

Nothing hypothetical about the experiment. It ran for fourteen years, thirteen miles south of Chicago, and it has a visitor center. I grew up in that city, and the name was still a word people used.

In 1880 George Pullman built a town for the workers of his Palace Car Company. It was, by the standards of nineteenth-century industrial housing, genuinely good: brick houses, indoor plumbing, gas, a library, a theater, drainage, parks, no saloons. Visitors came to admire it. Pullman's theory was explicit and not cynical: workers living in clean and orderly surroundings would be better workers, and the town would return six percent on its capital while proving the point.

He owned all of it. The houses, the shops, the church building, the water, the gas. Rent was deducted from wages before the wages were paid. Employees weren't permitted to buy their homes.

Nothing in that arrangement is illegal, and for a decade nothing about it was even scandalous. Trouble arrived when the two systems moved in opposite directions. In the depression that followed 1893, Pullman cut the workforce by roughly a third and the remaining wages by about a quarter. He didn't cut the rents. There was no reason he should have. The housing company and the manufacturing company were separate accounts, and the housing company was doing fine.

For the workers, however, the situation was not fine. Their income fell and their rent, their largest fixed cost, didn't, and it was deducted first. Some men collected pay envelopes containing a few cents for a fortnight's work.

They walked out on May 11, 1894. The American Railway Union boycotted trains carrying Pullman cars, the boycott spread across the country, and it ended with federal troops.

What finally dismantled the arrangement wasn't the strike. In 1898, the year after Pullman died, the Illinois Supreme Court found that a manufacturing charter did not authorize the company to own a town, and ordered it to sell everything not strictly used for manufacturing.

The company didn't comply until 1907. The last of the buildings went in 1909.[59]

Count the intervals. Twenty-eight years from the town's construction to the beginning of compliance. Eighteen years from the arrangement becoming visibly dangerous to the arrangement being undone. Nine years between the highest court in the state saying *this is not permitted* and the company acting as though it were.

And the mechanism that eventually worked wasn't a law about company towns, because there was no law about company towns. It was a corporate charter provision about what a manufacturer is allowed to own, a rule written for entirely different reasons, retrofitted to a situation nobody had anticipated, a generation late.

Every difference from the modern case runs the wrong way. Pullman's residents could walk out of the town, and eventually did. They shared a physical place, which made them easy to organize. The company's assets were land and buildings, which sit inside one jurisdiction and can be ordered sold by one court. A productive bottleneck with a proprietary unit of account has none of those handles. Its customers are dispersed, its balance sheet is portable, and there is nothing for a sheriff to stand in front of.

THE PAYROLL PROBLEM

Robots expose an older dependency hiding inside public finance, and this one is arithmetic, not conspiracy.

Schools, police, courts, counties, pensions and social insurance are funded in substantial part by taxes attached to human work: wages, payroll, and the business activity that surrounds employing people. Nobody designed that architecture. It accumulated, in an era when productive value and the employment relationship were nearly the same thing, and it has the enormous practical advantage of being easy to see. A paycheck is a legible event.

If productive capacity moves from payroll to machines, the value doesn't disappear. It is still being created, still being sold, still generating returns. But a tax system that knows how to see a paycheck may simply fail to see the relationship that replaced it, and a public balance sheet can hollow out while every headline number looks fine.

Whether a robot should get a W-2 isn't the question, and that joke has been made enough times. If a machine does the work a taxed worker used to do, and the company earns as much or more, who pays for what that worker's taxes used to pay for?

Because the obligation hasn't stopped being owed. The automated enterprise still runs on courts, roads, currency, educated people, public safety, communications, and a customer base that has to be housed and healthy enough to buy things. Benefits without obligations aren't merely unfair. The structure is unsustainable.

RELATIONSHIPS INVENTED TOMORROW

Law is a memory system, and it's a slow one by design.

Statutes name relationships. Courts test the edges. Regulators accumulate experience. Precedent records what happened the last time an arrangement met reality, and the record is what makes the next case cheaper to decide. Deliberately conservative, because the cost of getting it wrong is high and the cost of being late is diffuse.

Convergence defeats it, and defeats it in a particular way. A single arrangement performs functions that used to belong to several separate categories, each with its own body of law. A token can behave like property, payment, security, loyalty scheme, and access key, in that order, on any given day. A robot can be equipment, productive capital, service provider, workplace participant and data-collection system simultaneously. Every regulator involved can look at it, correctly identify the part inside their remit, and correctly conclude that the rest is somebody else's problem.

So adequacy has to be measured by function rather than by novelty. What does this system actually do? What productive service does it control? Can people refuse it? Can they leave, and what does leaving cost? What public infrastructure is it running on? Which costs are landing on people who never agreed to bear them?

Similar functions should carry similar obligations even when the nouns are new.

These cases were resolved eventually, and each was resolved late, thirty-five years for the clocks, twenty-eight for the town, and in both cases the resolution came through a rule written for something else entirely. A governor that slow would have wrecked the engine. There's no reason to think we've gotten faster. The tools are faster, which isn't the same thing, and the gap between those two sentences is where the next several decades are going to be spent.

Power should acquire obligations as quickly as capability acquires scale. It never has yet.

CHAPTER 20

The Velocity of the Ordinary

The interval between miracle and infrastructure is shrinking

"The thing that hath been, it is that which shall be; and that which is done is that which shall be done: and there is no new thing under the sun."

Ecclesiastes 1:9 (King James Version)

The remarkable thing is not only how far civilization has traveled. It is how quickly the scenery now changes.

My retailing family sold some of the appliances that made modern domestic life ordinary, and I watched several of them stop being remarkable.

Within one life, computers moved from institutional rooms to desks, then to pockets, then to cars and kitchens, and now into conversation.

None of that is unusual.

THE HALF-LIFE OF ASTONISHMENT

Call it the half-life of astonishment: the interval between an achievement that seems miraculous and the same achievement becoming an ordinary expectation, complaints included.

Electric light, refrigeration, powered flight, antibiotics, television, computing, satellite navigation, video calls, machine translation: each made the trip from wonder to infrastructure, and the trips got shorter. Powered flight took about fifty-five years to go from a wooden aircraft at Kitty Hawk to jets carrying paying passengers across the Atlantic as a scheduled service. Satellite navigation went

from military system to a thing in a car in roughly a decade, and in rather less from a thing in a car to something people are annoyed about when it is briefly wrong.

The current situation is instructive mostly for how absurd it is. ChatGPT was released at the end of November 2022. By the following January, according to a UBS analysis widely reported at the time, it was estimated to have reached 100 million monthly users, a number that had taken TikTok about nine months and Instagram about two and a half years.[60] Two months from novelty to mass utility, and roughly another two before people began describing it as disappointing.

Complaint is the tell. Nobody complains about a miracle.

EXCEPT WHEN IT DOESN'T

Acceleration isn't a law, and treating it as one has embarrassed a great many forecasters.

On Oct. 24, 2003, the Concorde jet flew its last commercial passenger service from New York to London. It cruised at about twice the speed of sound and did the crossing in three hours and twenty minutes. It carried more than two and a half million people over its working life, and then it stopped, and nothing replaced it.

A passenger crossing the Atlantic this morning will do so at roughly 85 percent of the speed of sound, in something over seven hours, more than twice as long as the same journey took in 1976.[61] That has now gone on for over two decades and shows no sign of being temporary, and the reasons are entirely mundane: fuel cost, noise regulation over land, a small market, a crash, and the absence of anyone willing to fund the successor.

Fusion power has been arriving shortly for seventy years. Construction productivity in most rich countries has been roughly flat for decades, while manufacturing productivity has multiplied, which

is why a factory can produce a car for less than it used to and nobody can build a house for less than they used to. Nobody has walked on the Moon since 1972.

Some technologies accelerate dramatically, particularly if they are mostly information and can therefore be copied at nearly no cost. Others move at the speed of concrete, or of regulation, or of the slowest institution that has to say yes. Biology, infrastructure, law, culture and capital all run on different clocks, and a forecast that assumes one clock for all of them will be wrong in a specific and predictable way.

What has genuinely changed is that the platform from which the next question gets asked is now being rebuilt inside a single working lifetime, more than once. Faster is the least of it.

THE TRANSISTOR ECONOMY

An economy starts to resemble an integrated circuit when useful functions get placed close enough together to feed one another.

Local renewable electricity powers storage, transport, computation, pumps, factories, homes and microgrids. Machine intelligence raises the productivity of design, administration, maintenance, logistics, teaching and small enterprise. Flexible manufacturing moves some production nearer the customer while the knowledge required to do it travels globally and instantly. Each of those is useful alone. Together they compound, and the compounding is the point.

Integration is productive when it removes friction without removing exit. Put it in the negative, because the positive version is easy to fake. Does the arrangement make it easier for a person to compare, enter, leave, create and connect? Or has it merely become extremely efficient for whoever owns it?

The same convergence that can produce a closed loop of private money and private production can produce a low-friction participant economy in which a person in a small town has the tools that used to require a corporation. The technology doesn't decide which. Architecture decides, and architecture is chosen by people who usually aren't in the room when the rest of us find out.

THE PEDESTRIANIZATION OF EXPERTISE

There's another way to describe the velocity of the ordinary. A capability starts out specialist, expensive and astonishing. Then the tools get cheaper, the controls get easier, and something once reserved for an institution ends up in an ordinary person's hands. That is pedestrianization. It happened to photography, navigation, publishing, arithmetic and the encyclopedia, and it is starting to happen to parts of law, medicine, engineering, accounting and research. It doesn't abolish experts. It changes where their time is worth most. The calculator didn't abolish mathematics. It stopped wasting mathematicians on arithmetic everyone else could now do.

That's the direction I care about: give ordinary people more of the tools that used to require extraordinary access, and leave the judgment with them.

THERE IS NO FINAL WONDERFUL

Once a new wonderful exists, it becomes furniture, and that isn't a failure of gratitude. It's how the next question gets asked.

A cured disease reveals the next disease, which was always there and was simply not the one killing you first. Cheap energy reveals the next thing worth doing with energy. A better answer reveals a deeper question, and the deeper question is usually harder, more expensive, and more interesting.

So the future is an accelerating capacity, not a list of products: a capacity to rearrange energy, knowledge, matter and institutions, unevenly, on several different clocks, with long stalls in the parts that

need concrete poured or a law changed.

Direction can be reasoned about far more confidently than destination. Anybody selling you the destination is selling you a mood.

Distance tells us how far we have come. Velocity tells us how fast the meaning of ordinary can change, and which parts of ordinary refuse to.

THE EQUATION

Shorter discovery-to-access time + wider participation + retained exit = a shrinking half-life of astonishment without surrendering agency.

CHAPTER 21

The Arithmetic of MORE

Eight billion mornings, a rising floor, and the widening mathematics of NEXT

Every morning, more than eight billion human beings awake. They do not wake with the same politics, wealth, language, ambitions, fears, or ideas. They do not need to.

A child wants to understand something that was explained badly yesterday. In Uttar Pradesh a girl is studying by the light of her father's phone again. A parent wants a safer home. An engineer looks at an arrangement on a screen and sees, with irritation and then excitement, that it could be better. A shopkeeper wants one more customer. A scientist asks a question nobody has answered, and a good number of them are asking the same question this week without knowing it.

Somewhere a person wakes with an idea that did not exist the night before.

Most of those ideas are small. Many are wrong. A few are wrong in ways that will take twenty years and a great deal of money to discover. A very small number will change what is possible for everybody. Nobody has to know in advance which is which, and nobody ever has. All the arithmetic requires is that more people be able to take part.

IT WAS HERE THE WHOLE TIME

Civilization is complicated, and there's a living in saying so.

It also reduces to one morning.

Billions of people wake up carrying yesterday's legacy and wanting something about today to be better than it was. They don't need to want the same improvements. They don't need to agree about God, government, property, culture, status, or what a good life consists of. Their disagreement is the input, the thing that makes the field of vectors a field rather than a line, and no obstacle to the calculation.

The common denominator is that they are all still participating in NEXT, which is a different thing from consensus.

Which is why raising the floor and keeping the ceiling open belong in one equation rather than two.

A person consumed by survival has fewer practical choices about what to contribute. A person with safety, health, knowledge, energy, trustworthy information and lawful agency has more. Multiply that difference across billions of mornings and what the whole world can make grows in both directions at once: more people can enter, and the most capable are building on a larger and richer field.

Humans waking up wanting better are not one chapter of the mathematics of civilization. They are the recurring input to every chapter.

The human beside the tamed fire possessed the same extraordinary instrument we possess. Same brain, same capacity for language, same stubborn preference for a warmer night. Nothing about the hardware has meaningfully improved.

What accumulated was everything between our minds: language, writing, mathematics, agriculture, sanitation, medicine, electricity, law, institutions, computation, and now machines that can hold more of the record in view than any person ever could.

We don't wake each morning and rediscover fire. We inherit it.

THE STRUCTURAL LIMIT OF TAKING

There is only so much Better available from somebody else having less. That isn't a political statement. It's arithmetic. If one person has a hundred and another has nothing, moving some of the hundred can change the second life dramatically, and sometimes it should. But the total is still a hundred. Redistribution can change who holds what already exists. By itself it can't explain how a hundred becomes a thousand.

History's largest improvements did something else. They created more. Electricity didn't improve civilization by deciding who was entitled to the existing candles. Antibiotics didn't redistribute the existing supply of health. Literacy didn't take knowledge away from the people who could already read, and the semiconductor didn't divide the world's computing power more fairly. Something new became possible. That is a different mathematics from the old one, which could only take and divide.

Redistribution is bounded by what already exists.
Participation expands what can exist.

When a child survives, learns to read, and gets access to energy, information, transport, health care and eventually a useful income, civilization has gained more than a recipient. It has gained another participant: another customer, another worker, another possible inventor, another person able to ask a question none of the rest of us thought to ask. That's why the largest remaining opportunity may be at the floor. There is a limit to how much anyone can gain by making somebody at the top poorer, and a much larger gain in making billions of people at the bottom more capable, which can make the people at the top richer too. A billion more customers enlarge markets. A billion educated minds enlarge the field of invention. A billion people with something to lose enlarge the constituency for stability.

So the ceiling can rise because the floor rises, and the floor can rise again because the larger market makes the next improvement profitable. That isn't goodness. It's a loop: raise the floor and people can do more; when they can do more, more of them take part; more people taking part means bigger markets and more invention, which creates more value and makes the next rise in the floor pay for itself. Nothing guarantees it, and the loop can be broken. But it changes the central question from who must have less to how many more people can become able to make and buy more. That is a much bigger number.

REINVENTION IS PART OF THE MATH

More doesn't have to come from somebody having less. Often it comes from reinvention. The washing machine didn't require somebody else to wash more clothes by hand. The shipping container didn't require somebody else to move freight more slowly. The semiconductor made uses possible that the vacuum tube couldn't have supported at any price.

Reinvention isn't painless. Businesses fail, jobs disappear, and some communities carry transition costs they didn't choose. A civilized society should see those costs rather than hide them inside the word progress. But freezing every successful arrangement in place would freeze yesterday against tomorrow. A better next doesn't preserve every existing arrangement. It preserves enough agency that people can build the next one.

If raising the floor enlarges the market, and enlarging the market raises the ceiling, then the people already doing best have an economic interest in helping more people take part. If it is better for everybody, it may become even better for those it is already best for.

MORE IS NOT UNIFORMITY

More participants doesn't mean identical participants.

Success doesn't require that nobody loses a competition or makes a ruinous investment. It requires only that the system not depend on a permanent class of people whose failure is structurally necessary to everyone else's success.

A market expands when yesterday's excluded participant becomes tomorrow's customer, producer, inventor, teacher or owner. It contracts when the exclusion is the business model.

They can say yes to exchange, invention, education, medicine, housing, cleaner energy, and the next useful arrangement.

They can also say no. No to conquest. No to information systems that deliberately poison the shared record for a fee. No to private arrangements that demand public benefits while declining public obligations. No to any structure whose success requires the field around it to become less valuable.

No isn't the lesser half. A civilization that could only say yes would have been sold something unsurvivable a long time ago.

Machine intelligence now gives an individual access to more accumulated memory and cross-correlation than any individual has ever had. At collective scale it may do something else as well: show where the vectors are actually pointing, including the ones that no institution currently counts. It is the student pilot's wind, on the panel at last.

MORE

Tomorrow morning begins with most of what civilization managed to remember from yesterday, plus another day of eight billion people asking what might be better.

More participants. More agency. More knowledge. More memory. More questions. More experiments. More people above the floor. More room at the ceiling.

MORE.

That isn't a promise of paradise. It is an observation about distance already traveled by people who mostly did not know they were traveling, who were cold on a Tuesday, and who thought there might be a way to fasten the hide.

We inherit yesterday's better arrangements. Somebody asks a better question.

And NEXT begins again.

THE EQUATION

Inherited capacity x more capable participants x better questions x expanded memory x correction = MORE possibilities for NEXT.

EPILOGUE

Take a Deep Breath and Smile

The distance already traveled is evidence enough to keep asking

If you are here to ask the question, it got better.

That sentence doesn't deny pain. It recognizes the chain of survival required for a question to exist. Ancestors endured long enough to have children. Communities preserved language. Someone taught you to read or listen. Agriculture, water, medicine, energy, roads, markets, law, and countless acts of care created the moment in which you can consider what comes next.

* * *

Televisions, washing machines, refrigerators: the machines that quietly returned hours of life to ordinary families. The question worth putting to every one of them is not what it does but what it gives back. A washing machine converts a Monday at the washtub into a Monday spent somewhere else. What it gives back is an hour. And what a person does with that hour, read, sleep, learn something, talk to a child, or spend it on something the machine's inventor never imagined, is the whole story.

I've been thinking about those hours. On the death rates of 1990, well over a hundred million children were going to die before their fifth birthday in the years since, and didn't, because of a vaccination schedule, an oral rehydration packet, a bed net, or a clean water pipe that arrived in their village sometime in the last thirty years.[62] Each of those lives is a returned hour at the scale of a civilization: a capacity that exists now and didn't before, belonging to a person who'll use it in ways nobody planned. Machines my family sold returned hours to

the households that bought them. Public health returned something larger, not hours but decades, not to households but to a species, and the return has barely begun to compound.

What does a species do with returned time? The same thing an individual does: sometimes squander it, often surprise itself, occasionally build something that outlasts the generation that built it. Every tool named in these pages — the printing press, the shipping container, the vaccine, the electric grid, the AI — is an argument about what humanity chose to do with capacity that had just been freed. Not always a flattering argument. Freed capacity has been turned into weapons, propaganda, monopoly, and spectacular cruelty with the same reliability that it has been turned into science, art, and cooperation. Capacity is like that, and it's no paradox. It goes where the values of the system that holds it direct it.

Which is why the question worth putting to a washing machine is still the right question to put to everything that follows. Not what does it do. What does it give back. And to whom. A peanut farmer in Americus, Ga., answered it for me when I was a boy: when the power went out, he could still call his relatives.

* * *

We are far past painting on the cave wall and still doing the same sacred work: leaving a mark that another mind can use. The difference is scale. A question can now travel through a planetary network, encounter centuries of memory, cross languages, and return with relationships no unaided person could assemble.

That is not only a change in degree, but a change in kind. When a question could only travel as far as a person could walk or a letter could be carried, the memory available to answer it was local, bounded by the library, the elder, the community's accumulated experience.

What was known elsewhere stayed elsewhere, and the gap between what was known and what could be accessed was the primary limiting factor on most human decisions for most of human history. Literacy, printing, telegraph, telephone, internet, AI: each one is a reduction of that gap. None of them eliminates it. Each of them moves the frontier of what a person can access without having to be in a particular place, know a particular person, or already belong to a particular institution.

The lights are on, but we decide where they point. We can use AI to sharpen domination, automate suspicion, and manufacture obedience. Or we can use it to increase agency, expose hidden costs, preserve disagreement, and help people see the common direction inside their differences.

Such choices get made continuously, in the design of interfaces and in the policy of platforms and in the curriculum of schools and in the regulatory posture of governments and in the questions that individual people decide to ask or not ask. Every one of those is a decision that a real person makes in a real week, usually under pressure, usually without knowing that it is a decision that will compound. The washing machine's designers didn't know they were returning Mondays to millions of households. They were solving a mechanical problem with the tools they had. But the decision they made, and the decision the people who sold the machine made, and the decision the policymakers made about electrification and about water, those decisions stacked. That stack is what we are living inside.

The past supports cautious hope. Humanity has repeatedly fallen and gotten up. It has converted suffering into institutions, error into science, scarcity into engineering, exclusion into wider citizenship, and memory into tools. Not always. Not everywhere. Never completely. Enough to prove that the next better is not fantasy. The future supports the same explanation. More people can now participate in perception, memory, comparison, and creation. As

individual agency increases, the human vector field becomes richer and more honest. As collective AI systems become better at cross-correlating that field, shared interests and hidden conflicts can become more visible without requiring anybody to surrender individuality.

* * *

None of this asks us to sneer at the systems that carried us here. The religions that taught duty and mercy, the civil institutions that made disagreement survivable, the markets that coordinated strangers, the armies that protected societies when protection was necessary, and the sciences that made doubt productive are all part of the inheritance. Gratitude is permission to understand what they accomplished well enough to keep the benefit while improving the arrangement, and no command to freeze them.

Which may be the simplest form of civility there is: honor what brought us this far, don't mistake what we inherited for the last word, and leave the next generation more than we received.

Nature didn't give humanity a finished instruction manual. It gave us the capacity to notice consequences, remember, communicate, and try again. Civilization is that recursion made visible.

Recursion is the key word. Each returned hour, each freed capacity, each better tool isn't the end of the story. It's the condition for the next question. The cave painting is returned to by the person who saw it and went home and told someone, who changed it, who built on the change, who eventually wound up, through a chain that is humbling and astonishing at once, here, reading this sentence, holding the accumulated memory of a species in a device that fits in a pocket and costs less in real terms than a washing machine cost in 1965. Extraordinary, the distance traveled. Extraordinary too, the distance remaining. And the ratio of those two is what makes it

possible to feel something other than triumph or despair, something closer to this: there's work to do, and the work is worth doing, and it has been worth doing the whole way.

History is no longer written only by victors. It is being written by vectors, and every person has one.

That is the mechanism: a field of independent human directions, each shaped by what that person can perceive and remember and value, adding up in aggregate to a resultant that nobody chose and everybody contributed to. The resultant has been, across the long view, toward more people surviving, more people literate, more people with access to energy and medicine and law and the capacity to argue back. The direction has not been uniform. It has not been without catastrophe. It has been a direction. Our task is to understand the mechanism well enough to tend it, and to tend it well enough that the direction holds through what comes next.

Take a deep breath. Smile if you can. Be grateful for the distance already traveled and responsible for the distance still ahead.

CODA

NEXT

There is no final wonderful

The leaves return after winter. Not the same leaves, and not on the same branch. But continuation has a shape.

Human life is difficult partly because each of us experiences it from inside a single lifetime. We arrive in the middle of a story and leave well before it is over, and the part we are given is never the part that resolves. I was born before there was a television in the house. Arbad, who read every chapter of this book, will see more of what it is about than I will, and he won't see the end of it either.

That does strange things to our sense of scale. Politics feels final because an election consumes years of attention. War feels final because lives end inside it, and they do not come back when the treaty is signed. Illness feels final because a person is not an abstraction and never has been. Mortality makes the present enormous, and it is not wrong to.

Civilization is running on a different clock, and it isn't a kinder one, only longer. Generations suffer, build, forget, rediscover, repair, and hand forward arrangements they will not see finished. Most that any of us manages is to live as fully as we can and leave more usable capacity than we arrived with.

So the mathematics of civilization is a description of recurring pressures, and the pressures are these. More capable participants make larger markets. Truth lowers the cost of correction. Trust stores efficiency the way a battery stores charge. Health returns people to the field. Knowledge gains value as more minds get to use it. Peace preserves the assets that war spends. And a stronger floor gives more

people something to lose, which turns out to be the most reliable reason anybody has ever found to protect the conditions of continuation.

Backing away from the tools is no answer. What answers it is making sure that power acquires obligations as fast as capability acquires scale, which has never once happened on time, and is therefore worth starting early.

The floor remains the test. Not because the ceiling doesn't matter. The fortunate have already shown, at length and in detail, what the ceiling can do. What's open is the floor. Do ordinary people wake up safe enough, healthy enough, housed enough, taught enough, free enough, connected enough and hopeful enough to take part in the next improvement?

Everything else is commentary on that question.

We do not control NEXT. We increase what NEXT has available to work with.

Be happy. Be safe. Leave more than you came with. Then let NEXT do what NEXT has always done.

APPENDIX

The Core Equations

A compact map of the argument

Each equation below is a shorthand for a chapter's central argument. The variables are defined at first use in the chapter and in the Glossary. They do not have numerical values. They have direction.

* * *

The Better Test (Author's Note, recurs throughout)

A choice is better when its benefits remain durable as more people, more consequences, and more time are included.

* * *

The Human Constant (Chapter 1)

Wanting better x inherited memory x freedom to differ x capacity to correct = civilizational motion.

* * *

Civilizational Motion (Chapter 2)

Need + memory + experiment + selection = the persistent pursuit of a better next.

* * *

From Scarcity to Abundance (Chapter 3)

Falling marginal cost + wider access + more participants = a shift from guarding value to growing value.

* * *

The Eight Billion Market (Chapter 4)

Human capability x stability x voluntary exchange = compounding prosperity.

* * *

The Wider Ledger (Chapter 5)

Local gain − hidden costs = fragile profit. Shared capability + trust + lower external cost = durable value.

* * *

Agency (Chapter 6)

Grounded knowledge + meaningful choice + feedback + memory = individual agency, bounded by the same wider ledger — an act that reduces another's agency fails the test at the level of the individual.

* * *

Energy Abundance (Chapter 7)

Renewable generation + storage + transmission + flexible demand = coordinated electrons, carrying both energy and information through the same medium.

* * *

Truth Infrastructure (Chapter 8)

Provenance + corroboration + uncertainty + value transparency + question refinement = better decision quality.

* * *

Super-Sentience (Chapter 10)

Human values + expanded memory + cross-correlation + user control = super-sentience as greater individual agency.

* * *

History by Vectors (Chapter 11)

More authentic vectors + shared visibility = more opportunity for voluntary alignment, bounded by: the measurement instrument must not become an input to the thing it measures.

* * *

Value-Structured Intelligence (Chapter 12)

Relevant knowledge + explicit value structure + persistent question evolution = a wider field of judgment.

* * *

Alignment Without Uniformity (Chapter 13)

Diversity + agency + translation + shared evidence + fair revision = alignment without uniformity.

* * *

Two AI-Literate Generations (Chapter 14)

Two AI-literate generations x agency-preserving tools x inclusive transition = a plausible path from collective memory to collective better judgment.

* * *

The Book's Composite (Epilogue)

Eight billion independent people + expanding agency + collective memory + better questions = humanity's next opportunity to become more aligned, peaceful, prosperous, and fully itself.

THE PROVENANCE OF VISIBILITY

Memory + independent comparison + visible provenance + human judgment = a larger field of view.

PERSISTENT ALIGNMENT INTELLIGENCE

Reliability first + visible alignment and misalignment + provenance + persistent memory + user control = truth infrastructure that learns without becoming an oracle.

THE PHYSICAL FLOOR

Water + abundant low-emission electrons + storage + intelligent timing + recycling = more durable physical capacity.

THE FLOOR TEST

A stronger floor + lower fear of falling + wider participation = greater civilizational durability and a higher attainable ceiling.

SOMETHING TO LOSE

A life worth protecting x credible opportunity tomorrow + inherited capacity = a larger peaceful stake in continuation.

PEACE CONVERSION (CHAPTER 18)

Preserved security + industrial conversion + abundant coordinated electrons + more capable participants = larger markets with less value destroyed.

THE STATUTORY GAP (CHAPTER 19)

Converging capability - attached obligation = the statutory gap.

NEXT

Inheritance + memory + agency + adaptation + correction = more capacity for whatever comes next.

GLOSSARY

Terms Used in This Book

Agency. The practical capacity to perceive choices, understand their consequences, act, learn from what happens, and revise the next choice accordingly. Agency in this book is bounded by the same wider ledger that defines better: an act that raises one person's capacity by destroying another's fails the test at the level of the individual, not only at the level of policy. Raw ability to do as one wishes — including the ability of the man with the gun — is not agency in this book's sense. Coerced compliance is not agency. The bound is not a moral addition to the definition; it follows from the definition itself. See Chapter 6.

Alignment. Sufficient shared direction for cooperation without requiring identical values. Alignment is not unanimity. It emerges when enough people can perceive common interests clearly enough to choose cooperation over conflict. It is more likely where people retain independent means of checking what they are told, and less likely where everyone sees the same thing simultaneously without being able to verify it. See Chapter 13.

Better. A comparative outcome that increases durable human agency or capability after wider costs, longer time, and more affected people are included. Better is not perfect, not uniform, not painless, and not morally pure. Among the available choices, it means the arrangement that increases sustainable human agency and leaves more capacity for the next improvement. Every use of the word in this book can be checked against that definition. Where a claim fails the check, it is the claim that is wrong and not the definition.

Collective AI vision. The ability of many independently controlled human-AI systems to share, translate, and cross-correlate evidence and perspectives across languages, disciplines, and time. It is not one

machine mind and requires no single coordinating authority. Its value depends on each component system remaining controllable by and accountable to the individual using it, and on its staying good for the people around that individual beyond the moment.

Electron bank. A metaphor for the growing ability to store, schedule, exchange, and coordinate electricity as a practical asset across homes, vehicles, batteries, buildings, and grids. The same physical medium — the electron — carries both energy and information, making distributed energy infrastructure and distributed knowledge infrastructure two expressions of the same transition. See Chapter 7.

Human vector. A direction and magnitude expressed through a need, choice, value, question, measurement, or action. It is not a complete representation of a person and makes no claims about inner life. The mathematical image is used because it has an actual operation — vectors add — and the aggregate of human vectors is what this book calls history. Authentic vectors produce more useful resultants than manufactured ones.

Mathematics of better. The recurring pattern by which systems that expand sustainable survival, agency, truth, participation, and error-correction tend to outperform systems that depend on hidden cost, exclusion, or permanent coercion. It is not a law of physics and does not guarantee that progress continues. It is a description of a selection pressure that has been operating long enough to be treated as a pattern.

Nature becoming sentient. A metaphor for humanity — nature's most self-reflective product — increasing its collective capacity to perceive, remember, compare, model, and deliberately respond. It is not a claim that Earth is a single conscious organism or that any machine has woken up. What it describes is the stack of instruments, libraries, institutions, and now AI systems through which a biological species has extended its capacity to notice itself and act on what it

finds. See Chapter 10. Persistent Question Evolution (PQE). A recursive process in which each answer exposes assumptions, contradictions, missing evidence, and affected values, generating a higher-utility next question. The value of an AI system is partly the answers it produces and partly the questions it helps a person ask that they would not otherwise have known to ask.

RAG. Retrieval-Augmented Generation. An AI architecture that retrieves external documents to ground or improve a model response. RAG improves factual grounding but does not, by itself, surface the value structure through which retrieved material is read. See VSI below.

Super-sentience. Human cognition augmented by large external memory, rapid cross-correlation, translation across languages and disciplines, and AI-assisted simulation. Super-sentience does not require machine consciousness and does not claim that the augmented system is wiser than the person using it. It is a description of reach, not of judgment. The judgment remains the human's responsibility and the human's risk.

Truth by vectors. Evaluation of a claim by comparing independent lines of evidence — source quality, directness of observation, recency, stated incentives, uncertainty, and internal contradiction. A claim supported by many independent sources that had no access to one another before they agreed is more reliable than a claim supported by many sources that share a common origin. Unanimity is a warning, not a result.

Value-Structured Intelligence (VSI). A discipline applied at the moment of retrieval, not a new retrieval mechanism. VSI places the documents that record a person's or institution's declared values in the index alongside factual material, and shows which values shaped the retrieved answer. The functional test is whether the values are legible to the reader and whether the conclusion changes when a different value set is applied. VSI is a research program rather than a

settled field: related work exists in constitutional AI, value retrieval, moral graphs, and explicit memory systems. Claims of novelty should be tested against competing approaches. The chapter's argument does not depend on the architecture being novel; it depends on the projection being labeled. See Chapter 12.

Civility. The capacity to keep disagreement inside a shared structure long enough for evidence, revision, restraint, and future cooperation to remain possible. Civility in this book is not politeness; it is error-correction without unnecessary destruction.

Earth accounting. A wider ledger that considers useful output, capability created, harm avoided, security preserved, resources consumed, external costs, and value destroyed across more people and more time.

Public-health economics of peace. The treatment of preventable conflict and excess destructive capacity as recurring system costs that may justify finite investments in prevention, verification, industrial conversion, and collective security.

FURTHER TERMS

Billionaire Bucks. A generic thought experiment describing privately issued money whose usefulness is amplified by an issuer's control of productive assets, platforms, infrastructure, or essential services. It is not the name of an actual currency.

Floor. The minimum practical platform from which a person can participate in tomorrow: safety, health, shelter, knowledge, lawful agency, and enough stability to plan.

Functional sovereignty. The capacity of a private system to perform enough economic functions traditionally associated with public systems that participants become materially dependent on it, even while legal sovereignty remains with the state.

Half-life of astonishment. A metaphor for the time required for a once-miraculous capability to become ordinary infrastructure or expectation.

NEXT. Continuation itself. NEXT is not necessarily better. It is the next arrangement produced by inheritance, choice, constraint, technology, nature, and time.

Persistent alignment intelligence. A truth-infrastructure design in which repeated reliability checks preserve provenance, alignment, disagreement, and uncertainty patterns as hypotheses for future checking.

Private digital seigniorage. Economic value captured by a private issuer through the creation, circulation, reserve structure, fees, appreciation, or network effects of a digital money-like instrument.

Provenance of visibility. The inspectable history of where an observation, claim, memory, or comparison came from; the chain that makes later cross-correlation more trustworthy.

Something to lose. A practical stake in continuation: a life, relationship, asset, opportunity, identity, or future that a person or community has reason to protect.

Statutory gap. The distance between a system's real economic capability and the obligations existing law reliably knows how to attach to it.

Technological convergence. The creation of a new capability when technologies developed for separate purposes are combined into a single operating system or market structure.

NOTES

Notes and Sources

Evidence for the principal factual claims

The argument of this book is interpretive; the sources below support the numerical and historical claims it rests on. Population, poverty, health, energy, connectivity and military-spending estimates are revised periodically, and the editions cited reflect what was available while the manuscript was completed in August 2026. Where a figure is disputed, the note says so rather than picking a side quietly.

* * *

PREFACE

1. Marc Levinson, *The Box: How the Shipping Container Made the World Smaller and the World Economy Bigger* (Princeton University Press, 2006). The *Ideal X* sailed from Port Newark on 26 April 1956 carrying fifty-eight containers and reached Houston five days later; the voyage is described in Chapter 2.

2. United Nations Department of Economic and Social Affairs, Population Division, *World Population Prospects 2024*, which runs from 1950: world population of about 2.5 billion in 1950 and 8.2 billion in 2024; global life expectancy at birth in the mid-40s around 1950 and 73.3 years in 2024. On the reversal: World Health Organization, “COVID-19 eliminated a decade of progress in global level of life expectancy,” 24 May 2024 — global life expectancy fell 1.8 years between 2019 and 2021.

3. UNICEF, *Levels and Trends in Child Mortality*: 37.4 under-five deaths per 1,000 live births in 2024, down from 93.5 in 1990. The figure for 1950, roughly a quarter of all children, comes from the long-run series in the UN *World Population Prospects* and its reconstructions (Gapminder; Our World in Data), which are estimates rather than registrations for most countries. The rate of improvement has slowed: the annual reduction ran at 3.9 percent in 2000–2015 and 1.5 percent in 2015–2024.

4. UNESCO, *World Illiteracy at Mid-Century* (1957), which put adult illiteracy at about 44 percent around 1950. Current figure: UNESCO Institute for Statistics, about 88 percent adult literacy in 2024.

5. Long-run estimates of extreme poverty: François Bourguignon and Christian Morrisson, "Inequality Among World Citizens: 1820–1992," *American Economic Review* 92 (2002); Michail Moatsos, "Global extreme poverty," in *How Was Life? Volume II* (OECD, 2021). Estimates for 1950 vary with method and poverty line; around 60 percent is within their range. Current share: World Bank Poverty and Inequality Platform, at the $3.00 (2021 PPP) line.

CHAPTER 1: THE HUMAN CONSTANT

6. The five gains in this sentence, in order. Child survival and literacy as notes 3 and 4. Connectivity: International Telecommunication Union, *Facts and Figures 2025* — approximately 6 billion people, 74 percent of the world's population, using the internet in 2025. Renewable cost: IRENA, *Renewable Power Generation Costs in 2024*. Extreme poverty: World Bank Poverty and Inequality Platform, at the $3.00 (2021 PPP) line. The poverty headcount is the least stable figure of the five — the Bank's 2026 nowcast gives about 826 million, its 2025 published figure about 831 million, and the UN's 2025 SDG report 808 million. They differ by vintage and method rather than by disagreement, and the text uses the trend rather than any point estimate.

CHAPTER 2: THE PERSISTENT PURSUIT OF A BETTER NEXT

7. The eyed bone needle and tailored clothing as a precondition for high-latitude occupation are discussed across the Upper Palaeolithic literature. The dating is genuinely contested, which the text says; the Siberian examples from Denisova Cave run to approximately 50,000 years before present.

8. Timur Kuran, *Private Truths, Public Lies: The Social Consequences of Preference Falsification* (Harvard University Press, 1995). Kuran's account of the 1989 collapses in Eastern Europe is the model for the argument here and returns in Chapter 11.

9. Ignaz Semmelweis, *Die Ätiologie, der Begriff und die Prophylaxis des Kindbettfiebers* (1861), is the primary account. See also Sherwin Nuland, *The Doctors' Plague* (2003). Semmelweis ordered chlorinated-lime handwashing in the First Obstetrical Clinic in May 1847; monthly mortality fell from 18.3

percent in April to 2.2 percent in June. The text deliberately gives no percentages, because the published figures move considerably by month and by clinic and the argument does not turn on which pair is quoted.

CHAPTER 3: SCARCITY WAS A STAGE, NOT A DESTINY

10. International Renewable Energy Agency, *Renewable Power Generation Costs in 2024*. The global weighted-average cost of utility-scale solar electricity was approximately USD 0.043/kWh in 2024, roughly 90 percent below 2010, and on average 41 percent cheaper than the lowest-cost fossil-fuel alternative. One nuance the text does not lean on: the 2024 figure was 0.6 percent *higher* than 2023, the first such rise in the series.

CHAPTER 4: EIGHT BILLION MORNINGS

11. UN DESA, *World Population Prospects 2024*, as note 2 — 8.2 billion people in 2024, with the peak projected around 10.3 billion in the mid-2080s.

12. Chennai Metropolitan Water Supply and Sewerage Board. The Nemmeli seawater reverse-osmosis plant supplies over a million residents of the southern suburbs; together with the Minjur plant it accounts for roughly 200 MLD of a city supply that has been held near 650 MLD during shortage. Rationing across parts of the city continues.

CHAPTER 5: GOOD IS MORE PROFITABLE THAN BAD

13. Estimates of the aggregate market value of enslaved people in 1860 range from about $2.7 to $3.7 billion, exceeding the combined value of the nation's railroads, factories and banks. See the Measuring Worth project's "Measuring the Value of a Slave" and the survey in *Econlib*, "U.S. Slavery and Economic Thought." On the war's cost: Claudia Goldin and Frank Lewis, "The Economic Cost of the American Civil War" (1975), estimating direct costs near $6.7 billion. On the death toll: J. David Hacker, "A Census-Based Count of the Civil War Dead," *Civil War History* (2011), giving about 750,000 with a range of 650,000–850,000.

14. E. R. A. Merewether and C. W. Price, *Report on Effects of Asbestos Dust on the Lungs and Dust Suppression in the Asbestos Industry* (1930), the first epidemiological survey of the trade. Britain's Asbestos Industry Regulations 1931 followed within a year and came into force in March 1932. The American liability wave dates from *Borel v. Fibreboard Paper Products Corp.*, 493 F.2d 1076 (5th Cir. 1973), the first successful suit by an insulation worker.

15. Jane Gleeson-White, *Double Entry: How the Merchants of Venice Created Modern Finance* (2011). On its epistemological limits with respect to the discount rate and unpriced costs: Herman Daly and John Cobb, *For the Common Good* (1989), Chapter 4.

16. Johns-Manville filed for Chapter 11 in August 1982 while profitable, facing roughly 16,000 pending asbestos claims. The company's internal knowledge from the 1930s — including the Simpson–Brown correspondence — entered the public record through litigation discovery.

17. RAND Institute for Civil Justice, *Asbestos Litigation Costs and Compensation: An Interim Report* (2002) and *Asbestos Litigation* (MG-162, 2005). Claimants received about 42 cents of each dollar spent; defense costs took 31 cents and plaintiffs' lawyers and related costs 27. Total spending through 2002 is put near $70 billion in the 2005 study.

18. Enron filed for Chapter 11 on 2 December 2001. The $74 billion figure is a loss *claimed* in the shareholder litigation rather than an established total, and the text says so. Arthur Andersen employed roughly 85,000 people worldwide, some 25,000 of them in the United States, at the time of its indictment; it surrendered its license to practice on 31 August 2002.

19. US Environmental Protection Agency, *The Benefits and Costs of the Clean Air Act, 1970 to 1990* (1997). Direct costs of $523 billion in 1990 dollars against a central benefits estimate of $22 trillion — within a range of $6 trillion to $50 trillion. The range matters and is given in the text.

20. The northern cod (*Gadus morhua*) collapse. Canada declared the moratorium on 2 July 1992; the 1968 peak reported catch was about 810,000 metric tons; DFO reopened a limited commercial fishery in 2024 at a total allowable catch of 18,000 metric tons. On the measurement failure specifically: Daniel Pauly et al., "Anecdotes and the shifting baseline syndrome of fisheries," *Trends in Ecology and Evolution* 10:10 (1995).

21. The Union Carbide India plant at Bhopal released methyl isocyanate on the night of 2–3 December 1984; quantity estimates in the literature range from about 27 to 45 metric tons. Officially recorded deaths number 3,787; every serious estimate of the eventual toll is a multiple of that. The settlement of 14 February 1989 was $470 million.

CHAPTER 6: AGENCY IS THE ENGINE

22. James Watt patented the centrifugal governor in 1788, adapting a device already in use to regulate the gap between millstones in windmills; Christiaan Huygens had built a version for clocks earlier still. James Clerk Maxwell, "On Governors," *Proceedings of the Royal Society* 16 (1868), gave the first stability analysis and explained hunting mathematically. It is a founding paper of control theory.

23. UK Department for Transport, *The Inclusive Transport Strategy: Achieving Equal Access for Disabled People*, 25 July 2018: "We recommend that Local Authorities pause the development of shared space schemes which incorporate a level surface while we review and update guidance." Local Transport Note 1/11 was withdrawn at the same time. The campaign was led by the National Federation of the Blind of the UK, Transport for All, and Lord Holmes of Richmond.

24. *Williams v. Walker-Thomas Furniture Co.*, 350 F.2d 445 (D.C. Cir. 1965), opinion by Judge J. Skelly Wright. The cross-collateral clause kept a balance due on every previously leased item until the whole account was cleared, so a default on the last purchase exposed all the earlier ones. The reverse of the contract recorded Ora Lee Williams's social worker and her $218 monthly stipend. The court remanded on unconscionability, defining it as an absence of meaningful choice together with terms unreasonably favorable to the other party.

CHAPTER 7: ELECTRONS ARRANGED AS ENERGY, KNOWLEDGE, AND MEMORY

25. Namoruputh Health Centre, Turkana County, Kenya. Modular solar with battery storage was installed in 2022; the facility serves approximately 40,000 patients annually, and the battery-backed refrigeration held the vaccine cold chain through outages lasting months. Documented by Emergency Relief and Development Overseas (ERDO) and *The Energy Pioneer*. For county-wide context see UNICEF Kenya and Amref Health Africa's Vaccibox deployment (2023–24) and Energy 4 Impact / OVO Foundation's Project Jua (2018–24).

CHAPTER 8: TRUTH IS INFRASTRUCTURE

26. The International Prototype of the Kilogram, a platinum–iridium cylinder held under three glass bells at Sèvres, was sanctioned in 1889. The General Conference on Weights and Measures redefined the kilogram in terms of a

fixed value of the Planck constant, effective 20 May 2019 — 130 years.

27. Clair Cameron Patterson, "Age of meteorites and the earth," *Geochimica et Cosmochimica Acta* (1956), reporting the result first obtained in 1953: 4.55 ± 0.07 billion years, from lead isotopes in the Canyon Diablo meteorite. Patterson's subsequent work on atmospheric lead contamination is the reason the lead-crime literature in this chapter exists at all.

28. The fifty-six percent figure is Jessica Wolpaw Reyes, "Environmental Policy as Social Policy? The Impact of Childhood Lead Exposure on Crime," *B.E. Journal of Economic Analysis and Policy* 7:1 (2007). See also Rick Nevin, "How Lead Exposure Relates to Temporal Changes in IQ, Violent Crime, and Unwed Pregnancy," *Environmental Research* 83:1 (2000). The correction described in the text is Anthony Higney, Nick Hanley and Mirko Moro, "The lead-crime hypothesis: A meta-analysis," *Regional Science and Urban Economics* 97 (2022), 103826, pooling 542 estimates from 24 studies, finding substantial publication bias, and reporting that correcting for it leaves lead explaining roughly 7 to 28 percent of the fall in homicide.

29. FBI Uniform Crime Reports. The US violent crime rate peaked in 1991 at 758.2 per 100,000 and had roughly halved by the mid-2010s.

CHAPTER 9: THE OLD MATHEMATICS OF POWER

30. Dwight D. Eisenhower, "The Chance for Peace," address to the American Society of Newspaper Editors, 16 April 1953, delivered six weeks after Stalin's death on 5 March; and the Farewell Address of 17 January 1961, which runs about a quarter of an hour. The frequently repeated claim that the farewell address lasted under ten minutes is wrong. On the U-2 photography from 1956 and the absence of a missile gap, and on the American National Exhibition at Sokolniki Park in July 1959, see the Eisenhower Presidential Library's document collections.

31. Stockholm International Peace Research Institute, "Trends in World Military Expenditure, 2025" (April 2026). World military expenditure reached $2,887 billion in 2025, up 2.9 percent in real terms; spending fell in the United States while rising 14 percent in Europe and 8.1 percent in Asia and Oceania.

CHAPTER 10: AI TURNS THE LIGHTS ON

32. International Telecommunication Union, *Facts and Figures 2025*: approximately 6 billion people, about 74 percent of the world's population, using the internet in 2025, up from 60 percent in 2020.

CHAPTER 11: HISTORY WRITTEN BY VECTORS

33. The Ohio Anti-Saloon League was founded in 1893 at Oberlin, Ohio, by the Congregational minister Howard Hyde Russell. On Wayne Wheeler and the League's method of single-issue pressure politics, see Daniel Okrent, *Last Call: The Rise and Fall of Prohibition* (2010). Wheeler died on 5 September 1927, six years before repeal.

34. The Eighteenth Amendment was submitted to the states by Congress on 18 December 1917 and ratified on 16 January 1919. Connecticut and Rhode Island were the only two states to reject it. The Twenty-first Amendment was ratified on 5 December 1933.

35. The Leipzig Monday demonstrations began on 4 September 1989 at the Nikolaikirche with a few hundred people. On 9 October more than 70,000 walked the ring road; on 16 October the figure was about 120,000; by late October it was in the hundreds of thousands. The Berlin Wall opened on 9 November. The mechanism is Kuran's, as note 9.

36. The London Millennium Bridge opened on 10 June 2000 and closed two days later, remaining shut for about two years while Arup investigated. The standard account is synchronous lateral excitation; that account is contested — see Igor Belykh et al., "Emergence of the London Millennium Bridge instability without synchronization," *Nature Communications* 12 (2021). Nothing in the argument turns on which mechanism is right. On marching resonance, the Basse-Chaîne bridge at Angers collapsed on 16 April 1850, killing more than two hundred.

CHAPTER 12: VALUE-STRUCTURED INTELLIGENCE

37. Gerardus Mercator, *Nova et aucta orbis terrae descriptio* (1569). James Gall described an equal-area cylindrical construction in 1855; Arno Peters promoted the same projection from 1973. Carl Friedrich Gauss proved that a sphere cannot be developed onto a plane without distortion — the *Theorema Egregium* — in 1827. Africa is roughly fourteen times the area of Greenland. On the politics of the quarrel: Mark Monmonier, *Rhumb Lines and Map Wars* (2004) and *How to Lie with Maps*, 3rd ed. (2018).

38. Four hydroelectric dams on the Klamath River — J.C. Boyle, Copco No. 1, Copco No. 2 and Iron Gate — were removed across 2023 and 2024, the last in the summer of 2024. It is the largest dam removal in United States history. The Yurok and Karuk tribes were the principal parties and salmon passage the central issue.

CHAPTER 13: ALIGNMENT WITHOUT UNIFORMITY

39. *Phytophthora infestans* reached Ireland in 1845 and found in the Lumper potato a population with no resistance. On the political economy: Ireland remained a net exporter of food through most of the famine years, with grain, cattle and butter leaving for British ports. Roughly a million people died and roughly a million emigrated over the following seven years. See Cormac Ó Gráda, *Black '47 and Beyond* (1999), and Cecil Woodham-Smith, *The Great Hunger* (1962).

40. The Community Reinvestment Act of 1977 and the Federal Housing Enterprises Financial Safety and Soundness Act of 1992. The 1992 Act required HUD to set annual affordable-housing goals for Fannie Mae and Freddie Mac; the low- and moderate-income goal began near 30 percent, reached 50 percent for 2001–2003, and 56 percent by 2008. On what the goals did and did not cause, see the note on the Federal Reserve and Commission findings below.

41. Robert B. Avery and Kenneth P. Brevoort, "The Subprime Crisis: Is Government Housing Policy to Blame?" Federal Reserve Board, Finance and Economics Discussion Series 2011-36 (30 August 2011), later in *Review of Economics and Statistics* 97:2 (2015): "We find little evidence that either the CRA or the GSE goals played a significant role in the subprime crisis," and areas disproportionately served by CRA-covered lenders "experienced lower delinquency rates and less risky lending." Financial Crisis Inquiry Commission, *The Financial Crisis Inquiry Report* (2011), Conclusions, pp. xxvi–xxvii: the affordable-housing goals "only contributed marginally to Fannie's and Freddie's participation in those mortgages," and the CRA "was not a significant factor in subprime lending or the crisis." The Commission's dissents, Wallison's in particular, argue the reverse. The text follows the majority and the Fed's economists, and the chapter's argument does not depend on the dispute being settled.

42. Pew Research Center, "Wealth Gaps Rise to Record Highs Between Whites, Blacks and Hispanics," 26 July 2011, analyzing Census Bureau Survey of Income and Program Participation data. Median wealth 2005–2009: Black

families −53 percent, Hispanic families −66 percent, white families −16 percent. The Federal Reserve's Survey of Consumer Finances covers the same period with different methodology; the racial breakdowns cited here are Pew's.

CHAPTER 14: TWO AI-LITERATE GENERATIONS

43. Finnish National Agency for Education, *National Core Curriculum for Basic Education 2016*. On the PISA trajectory: Finland's scores have declined since roughly 2006, a decade before the reform, and the 2016 curriculum was phased in gradually enough to have reached the 2018 cohort barely at all. The Agency's own account of the decline ("PISA results reflect broader changes in Finnish society") names socio-economic inequality between families, education funding, reading motivation, social media and adolescent mental health. The reform can be neither credited nor blamed on present evidence, which is what the text says.

44. Per Kongshøj Madsen, "The Danish Model of 'Flexicurity': A Paradise — with Some Snakes," in *Flexibility and Employment Security in Europe* (2006). Unemployment insurance replaces up to 90 percent of previous earnings *subject to a ceiling*, which makes the scheme markedly more generous to low earners. Benefit duration was cut from eight years to seven in 1993, five in 1996, four in 1999, and two in 2010. On the 2008 stress: OECD *Employment Outlook 2010*, Chapter 4.

CHAPTER 15: THE FLOOR IS THE MEASURE

45. Samuel Plimsoll, *Our Seamen: An Appeal* (1872). The Merchant Shipping Act 1876 made a load line compulsory — a twelve-inch circle crossed by a horizontal bar — but did not fix its position, which Parliament did only in 1894. See Royal Museums Greenwich, "Samuel Plimsoll and ship safety," and the Lloyd's Register Foundation Heritage & Education Centre.

46. Anandi Mani, Sendhil Mullainathan, Eldar Shafir and Jiaying Zhao, "Poverty Impedes Cognitive Function," *Science* 341:6149 (2013). The Tamil Nadu sugarcane farmers were tested before and after harvest; the authors put the gap at roughly 13 IQ points. See also Mullainathan and Shafir, *Scarcity: Why Having Too Little Means So Much* (2013). The headline effect size has been argued over since, which the text says; the finding that survives is that the deficit traveled with the condition rather than with the person.

CHAPTER 16: SOMETHING TO LOSE

47. The European Recovery Program ran from 1948 to 1952 and disbursed about $13.3 billion, roughly 90 percent of it as grants. On how much of the recovery it caused — considerably less than folk memory holds — see Barry Eichengreen and J. Bradford DeLong, "The Marshall Plan: History's Most Successful Structural Adjustment Program" (1993), and the survey at EH.net. On the 1919 settlement and the reparations mechanism financed by American lending until 1929: Britannica, "Treaty of Versailles: German reparations," and the documents at germanhistorydocs.org.

48. Kongō Gumi, founded in 578 to build Shitennō-ji at the invitation of Prince Shōtoku, operated under forty generations of the founder's descendants. Its assets were purchased by the Takamatsu Construction Group in January 2006 after debts incurred during Japan's property bubble; it continues as a subsidiary.

CHAPTER 17: NATURE PUT EVERYTHING IN THE BOX; MAN FOUND GOOD AND BAD

49. Vaclav Smil, *Enriching the Earth: Fritz Haber, Carl Bosch, and the Transformation of World Food Production* (MIT Press, 2001). Estimates of the share of nitrogen in human food fixed by the Haber-Bosch process run from about 30 to 50 percent; Smil's own estimate is that roughly two-fifths of people now alive would not be without it. Haber demonstrated ammonia synthesis in 1909; Bosch industrialized it by 1913.

50. On Haber's war work: he proposed chlorine as a weapon and was present at its first large-scale release outside Ypres in April 1915. His wife Clara Immerwahr, herself a chemist, died by suicide shortly afterward. He was awarded the Nobel Prize in Chemistry for 1918 over substantial objection, and later died in exile after being driven out of Germany for being Jewish. On the process's continuing cost: it consumes on the order of 1 to 2 percent of world energy supply and accounts for about 1 percent of global carbon dioxide emissions. See J. W. Erisman et al., "How a century of ammonia synthesis changed the world," *Nature Geoscience* 1 (2008).

51. US Energy Information Administration, from California Independent System Operator data. CAISO curtailed about 3.4 million MWh of utility-scale solar and wind in 2024, a 29 percent increase on 2023, of which roughly 93 percent was solar. April 2024 was the record month at 1.46 million MWh, about 18 percent of that month's grid-scale wind and solar generation.

52. Phosphorus has no significant gas phase, so there is no atmospheric route to fixing it and no substitute for it in biology; commercial supply comes from mined phosphate rock. Roughly 70 percent of world reserves lie in Morocco. See International Fertilizer Association / Argus, *Phosphate Rock Resources and Reserves* (2023), and the African Development Bank's critical-minerals phosphate factsheet.

CHAPTER 18: THE PUBLIC-HEALTH ECONOMICS OF PEACE

53. The "Megatons to Megawatts" agreement of 18 February 1993 between the United States and the Russian Federation. Five hundred metric tons of weapons-grade highly enriched uranium — material from roughly 20,000 warheads — was blended down into more than 14,000 metric tons of low-enriched reactor fuel. The first shipment arrived in June 1995; the program completed in December 2013. For two decades up to about 10 percent of US electricity was generated from fuel made with this material. See Centrus Energy's program history, the US Energy Information Administration, and the World Nuclear Association's "Military Warheads as a Source of Nuclear Fuel."

54. International Renewable Energy Agency, *Renewable Power Generation Costs in 2025* (2026).

55. FAO, IFAD, UNICEF, WFP and WHO, *The State of Food Security and Nutrition in the World* (annual), and FAO food balance sheets: world food supply has exceeded average dietary energy requirements for decades; hunger persists through access, income, conflict and price.

CHAPTER 19: WHEN CAPABILITY OUTRUNS OBLIGATION

56. North American railroads imposed standard time zones on 18 November 1883 — the "Day of Two Noons" — having agreed the boundaries among themselves after Congress declined to legislate in 1809, 1870 and 1872. Congress ratified the arrangement in the Standard Time Act of 1918. See the Library of Congress, "The Day of Two Noons," and the Union Pacific Railroad Museum's account.

57. United Airlines Holdings, Inc., Form 8-K filed 12 June 2020 (SEC Archives, accession 0001104659-20-073190). The filing describes the commitment letter under which Goldman Sachs Lending Partners, Barclays and Morgan Stanley agreed to provide Mileage Plus Holdings, LLC and Mileage Plus Intellectual Property Assets, Ltd. with a term loan facility of up to $5.0 billion, secured on a first-priority basis by the assets of those

subsidiaries and by the cash accounts into which MileagePlus revenues are paid by United and by the program's marketing partners. The MileagePlus intellectual property, including brands and member data, was contributed to MIPA before closing.

58. The 2019 figures and the valuation are United's own, from the investor materials that accompanied the June 2020 MileagePlus financing: $5.3 billion of cash from mile sales (about 12 percent of airline revenue), $1.8 billion of EBITDA, 71 percent of the cash from third-party sales rather than from miles earned by flying, and a program valuation of $21.9 billion at twelve times 2019 EBITDA. The financing closed in July 2020 at $6.8 billion — $3.8 billion of 6.5 percent senior secured notes due 2027 and $3.0 billion of term loans. On the comparison with the airline's market value: UAL's equity traded in the neighborhood of $10 billion through June and July 2020, so the program's disclosed valuation was roughly twice the market value of the whole company. Published comparisons of the same kind for Delta and American that summer run in the same direction. The market-capitalization figure moved considerably week to week in 2020 and the text says 'comfortably more' rather than naming a multiple for that reason.

59. George Pullman built the town of Pullman, Illinois, in 1880; the company owned the housing, shops and utilities, deducted rent from wages before payment, and did not permit employees to buy their homes. In the depression following 1893 the workforce was cut by roughly a third and wages by about a quarter without a corresponding reduction in rents. The strike began on 11 May 1894 and ended with federal troops. In 1898 the Illinois Supreme Court held that a manufacturing charter did not authorise ownership of a town and ordered the sale of all non-industrial property; the company did not begin complying until 1907, and the last buildings were sold in 1909.

CHAPTER 20: THE VELOCITY OF THE ORDINARY

60. A UBS analysis widely reported in early February 2023 estimated that ChatGPT reached approximately 100 million monthly users in January 2023, about two months after release — against roughly nine months for TikTok and two and a half years for Instagram. It is an analyst estimate rather than a company disclosure, which is how the text presents it.

61. Concorde flew its last commercial passenger service from New York to London on 24 October 2003, cruising at about Mach 2.04 and crossing in three hours twenty minutes; it carried more than 2.5 million passengers over

its working life. Subsonic long-haul aircraft cruise at roughly Mach 0.85, putting the same crossing at something over seven hours.

EPILOGUE: TAKE A DEEP BREATH AND SMILE

62. A derived figure, so here is the arithmetic. Applying the 1990 under-five mortality rate of 93.5 per 1,000 to roughly 130 million births a year gives about 12.2 million deaths annually under a no-improvement counterfactual. Actual under-five deaths fell from about 12.5 million in 1990 to 4.9 million in 2024, averaging near 8.7 million across the period. The difference, compounded over thirty-four years, is on the order of 115 to 120 million. "Several hundred million" — the phrase an earlier draft used — does not survive that calculation, because applying the endpoint rate backwards overstates a decline that was gradual. Sources as notes 3 and 4.

* * *

A NOTE ON VALUE-STRUCTURED INTELLIGENCE

VSI is presented in Chapter 12 as a research program rather than a settled field. Related work exists in constitutional AI, value retrieval, cultural alignment, moral graphs, explicit memory systems and graph-based retrieval. The specific combination proposed — an explicit, evolving worldview model consulted at inference alongside persistent question evolution — is a claim to be tested against competing systems, not a result to be cited. Anyone who finds it has been done better elsewhere should say so; the argument survives the correction, because the argument was never that the architecture is novel. It was that the projection should be labeled.

A NOTE ON THE EMERGING-SYSTEMS MATERIAL

Chapters 19 and 20 are strongest where they describe architecture and weakest where they would have to forecast, and they are written to stay on the first side of that line. The humanoid-robot and private-money sections are thought experiments about convergence, taxation, exit and public obligation. They require no prediction about how many machines will be deployed by a given year, and they assert nothing about any actual person or company beyond what that company has itself disclosed.

The water-and-energy discussion likewise avoids a technology prescription. Its claim is narrower: potable water, atmospheric consequences, energy, materials, timing, storage and waste belong in one physical ledger. Nuclear

power is not excluded. Carbon capture, renewables, storage, transmission, desalination, recycling and efficiency are to be judged by demonstrated system contribution rather than by category loyalty.

A NOTE ON AI ASSISTANCE AND AUTHORSHIP

Artificial-intelligence systems were used throughout development as research interlocutors and cross-correlation instruments — for retrieval, comparison, criticism, contradiction, alternative drafts and persistent question refinement. Richard Polk is the author and retains responsibility for every judgment and argument here. Arbad Tahir edited the manuscript foundation and materially shaped its structure and its rigor.

ACKNOWLEDGMENTS

The Vectors That Made This Possible

No book about collective intelligence belongs to one mind. These pages grew from a lifetime of family, retail, public service, technology, energy, civic debate, and persistent conversation. They were sharpened by younger collaborators building across continents and generations, by readers on LinkedIn willing to test an unfinished idea, and by the vast human archive now available to anyone willing to keep asking.

I am especially grateful for the tension between generations, cultures, political traditions, and technical disciplines. The disagreement is not an obstacle to the work. It is part of the intelligence.

Most of all, I am grateful to the people who continue to wake up wanting a better day. They are the evidence, the market, the vectors, and the authors of whatever comes next.

THE BETTER NEXT IS NOT UNIFORM. IT IS MORE PEOPLE WITH MORE AGENCY SEEING MORE OF THE SAME REALITY.

* * *

Ask better questions. Expand agency. Leave more than you came with.

Arbad Tahir deserves to be named directly. Our work across Boulder and Islamabad showed, in practice, that two people can pull the same way without being alike. His editing gave the manuscript a structure sturdy enough to carry the continuing argument.

I am also grateful to the artificial-intelligence systems used in the work - ChatGPT, Copilot, Claude, and others - not as authors or authorities, but as new instruments through which a single human

question can be tested against more memory, more contradiction, and more possible next questions than were previously available in one room.

These relationships are the provenance of visibility.

“Be happy. Be safe. Leave more than you came with.”

RICHARD POLK