Chapter 1

The Challenges

The 21st century has begun in a tangle. As a species, we are more technologically interconnected than at any point in history, and yet in many ways we are more disconnected than ever. Our approach to industry has demonstrated incredible scale and prodigious output, to the point where some parts of the planet are mired in waste and overconsumption while our climate system is destabilizing. Our financial systems and economy, which are our main means of global organization, are strained to the point of dysfunction, and we have become inured to the fact that major economic collapse occurs roughly once a decade. Work, once laid out as the most reliable path to dignity and stability, is becoming uncertain for millions.

What we are experiencing is not a single upheaval. It is the convergence of multiple, deeply interconnected ones:

- AI and robotic automation are no longer nascent technologies, and their accelerating adoption is driving a fundamental reevaluation of labor (what we work on), cognition (whose thinking shapes it), and agency (who gets to decide what is done) as they challenge core assumptions of our current systems.
- Climate change is not simply a problem of average temperatures increasing but a volatility crisis. Each spot on the planet is going through larger and more unfamiliar swings in temperature and hydrology that will drive a system-wide disordering of the environment we depend on.
- Our economic systems, which still run on mental models from the Industrial Age, are out of sync with the limits and needs of the biosphere. These dislocations[1] are being preyed upon for political gain[2] and are driving further political and social instability.[3]

Ultimately these many upheavals come from our choices. They are symptoms of a design failure—a civilization built on partial or clouded understanding. We have treated the health of nature as an externality, labor as interchangeable and disposable, and technology as neutral. Each of these misguided assumptions now threatens our collective future.

Yet against this backdrop, this is not a book of despair. Instead it aims to provide clarity, tools, and agency for the century ahead. The systems that got us here were all designed by people, and that means that any and all of them can be redesigned. I personally know many of the people who designed technological systems that make up our digital infrastructure and run the current moment. Although there are some virtues in what we as technologists have built, there is clearly the need for a lot of improvements.

This moment feels existential because we are at the end of running systems that cannot sustain themselves. But such moments are the beginning of a new conversation, of new design, of changes aimed directly at the world we want—and I hope to have as many people as possible gain the skills to be active shapers of the changes that are important to them.

Ultimately we will need to design a civilization that works with nature, not against it. To design an economy that provides broad opportunity[4] even as the form that work takes continues to evolve. And to design political structures that are a culmination of collective human wisdom and possibility.

No one has a complete blueprint for the future. But we do have the power to ask better questions, to listen more deeply, and to design more wisely. This book proposes ways of thinking—and, more important, ways of doing—to meet the profound potential of this moment. Whether we hoped for this or not, we are the generations that must step up to this responsibility, and the decisions we are already making will reverberate for millennia. It's time to become the generation that is ready for that responsibility.

AI, ROBOTICS, AND THE FUTURE OF ECONOMICS

As we sit in the first half of the 21st century, we are in the process of another transformation ushered in by new technology. The last 50 years has seen in breakneck succession the rise of the personal computer, internet, cell phones, mass digitization, software for every imaginable application, the upgrading of physical industry through sensing and robotics, and now the upgrading of cognitive industry via artificial intelligence (AI) and agentic tools.

My life almost perfectly overlaps those 50 years, and in my 30+ years in the technology industry, I've been fortunate to contribute to hardware, software, design, and architecture spanning all these major technical transitions. We are in the early days of the last two, but they have been driving a lot more anxiety than other recent waves of technological change.

In some ways there is no real difference between AI and autonomous robotics and the many other tools and technologies that have preceded it. Technologies magnify our ability to carry out intentions, and AI and robotics absolutely are able to do this. Technologies also tend to be designed for reliable utility for their chosen task. Whether you are pressing play on a Walkman or asking a voice assistant to pause a track, the technology should work reliably. Reliability has allowed a type of cognitive comfort with the role of such technologies. A toaster toasts. A scooter scoots. For most tools and technologies we use daily, we can see exactly their intended use and can quickly assess the general bounds of their capabilities.

Herein lies one of the first departures of this latest round of technologies. Autonomous robots like the self-driving car carry out sophisticated actions and reason through complex environments without having the flat cause/effect reliability of most technologies. Generative AI systems often synthesize unexpected combinations that will sometimes meet our expectations and other times confuse or defy them.

When a being does something unexpected, or reasons through situations without much input, we tend to ascribe consciousness and agency to the being. Of course we do, because throughout our biological history, it has always been a correct assumption whenever we've had that experience. Whether the behavior came from your pet dog or a remote coworker, it's the right assumption that the other

being is conscious and employing agency in the process of doing something capable or unexpected.

Furthermore, ascribing consciousness and agency is actually a bit comforting to us. We have a lot of practice interacting with other conscious beings, so making this assumption is also a way of trying to interact in a way we are well practiced in. A person who speaks one language well and another poorly might try to speak to a stranger first in their stronger language just to see if they can understand in the language the person is most fluent in. We are highly fluent in the interactions between conscious beings.

As someone who has built these tools and contributed to widely used versions, I can tell you this: The assumption of consciousness is not correct.

While the outputs of each type of system (whether it is rules-based, inference-based, reinforcement, reasoning, or generative) can sometimes be surprisingly interesting or successful to the task, the architecture of any approach also create large regions of what is technically impossible for a specific system: limitations that actual humans and other conscious beings do not have. For example, a large language model asked to explain the physics behind a phenomenon could grab and remix what others have written about the topic, but it doesn't reason from an understanding of the internal structure of physics; it just tries to surface enough relevant text for you to work it out. This lack of understanding doesn't prevent us from wanting to communicate and relate to these technologies as if we were talking to a physicist, and that difference between the desire and reality will always generate errors. We may even interpret that experience as one where we seemed to be interacting with a conscious being. Again, our bias is that if the technology can interact in a manner we are already fluent in, we tend to prefer that.

The effect is not widely different from that in a well-established field: ergonomics. In this field we study the range of human motion: What actions are comfortable? Which are stressful? How should the design of physical things be put together to work well with most humans that will need to physically interact? The leap here is that AI and robotic tools enable a mentally ergonomic interaction with our consciousness. A table or chair feels comfortable when it is well designed for the human body, and an AI system might feel conscious to us when well designed for comfortable interaction with the human mind.

When I was at Google X, I was part of the leadership team and ran the group that developed all our product experiences, including the experience of the self-driving car, which has since spun out of Google and is currently known as Waymo. Having led many software and hardware efforts in the past, I was well accustomed to designing, testing, and iterating on regular tools and technologies across a wide variety of domains, spanning from search engines, productivity software, to advanced hardware work in interactivity, robotics, and sensing. Because what's common to all these tools is that people use the tool to get a task done, we can conduct user research on how successful they were in doing the task. By comparison, with the product experience of a self-driving car, the passenger is succeeding at the task of going from point A to point B, but they have no "driving" task. As such, without a task to complete, assessing the usability and success level with typical study design was going to give us nothing.

So we designed an entirely different approach to testing.

We had a summer intern from Rosalind Picard's group at MIT Media Lab, which specializes in affective computing, the discipline of assessing and factoring in human emotion

(aka affect) into how digital and computation systems work. Our intern was one of the world experts[5] in using a combination of galvanic skin response, heart rate and variability, facial microexpression analysis, and more, to triangulate what people were feeling about a situation without them needing to verbalize directly.

This meant that even if people did not have to do the driving task themselves, we could learn a lot about how they were experiencing the technology. One early result was eye-opening. You see, when you have a car that is updating its understanding of the world 60 times a second and can make control decisions on how to steer 10 to 30 times a second, then it's not that hard to do something that we thought would be amazing: to drive a perfectly smooth trajectory in the center of the lane with absolute minimum deviance.

But when we tested such a system, we realized quickly that it created massive stress for riders. Why? Well, a steering wheel that is making little robotic driving adjustments 10 times a second, even if the ride is smooth, visually looks like a steering wheel that is shaking or jittering. When a human is driving, the only situations in which a steering wheel is shaking or jittering are ones where there is either a serious mechanical problem or an extremely anxious driver. Neither situation is safe or comfortable. On the flip side, user responses made us realize that if we were thoughtful about those visual and kinematic experiences, then we could inspire comfortable or useful experiences for riders.

We came up with the idea of robotic choreography. Choreography is the discipline of using motion to inspire emotion. We recognized that the motion of a car you are in absolutely can inspire emotion. Imagine the difference in feeling (even with your eyes closed) of riding in a car driven by an extra-careful grandma, versus a professional limo

driver, versus an NYC cab driver. The difference you feel in that motion inspires differences in emotional response.

We realized that in the future we might teach robots to move like us, or move in ways that are akin to conscious beings we are familiar with (e.g., cab drivers). By doing so, we would simplify the cognitive and experiential interface between the person and the many lines of code and robotic sensor and control systems that were directing the driving. The task of making our tools relatable and approachable is foundational to the task of making a technology useful.

Given this, I see value in anything that humanizes technology so it works better with what we are foundationally and who we hope to become. The risk, of course, is that if we mistakenly believe that these technologies already are conscious, we will absolutely make mistakes as we find out they are not. We will put technologies in settings where they are inappropriate to the task, and as a result there will be loss of reputation, of finances, and of life. As of mid-2025, when I am penning this sentence, it's clear that this string of errors has already started in earnest.

This is a decidedly different danger to AI from the one that has absorbed so many imaginations, which is that we will soon achieve true AGI, or artificial general intelligence. This is the state when our AI system possesses "human-level" (or greater) intelligence, meaning the system can perform any intellectual task that a human can: learn, reason, adapt, and solve problems across domains.

I'm a skeptic on this, but I do believe we will soon achieve systems that *feel* like AGI. That danger is very real; it's just a different danger from actually achieving AGI. If you had an iguana that looked a bit like a dragon, you could easily imagine that it would grow up to be a dragon and gobble

you up. But if it's not actually a dragon, you might make fatal mistakes assuming too much of its future capabilities. You might, for example, spend trillions of dollars to be the first nation to breed many iguanas or put iguanas in charge of your national defense, expecting them to be dragon strength soon. And these mistakes could be just as dangerous to your world as the fabricated fear of dragons would be.

That brings us to this moment in history. Even without being conscious, AI and autonomous robotics are at the point where they can displace large swathes of the current job market—specifically the jobs that feature repetitive physical or cognitive labor. This may come in the form of a complete one-to-one replacement of a role, or in the reduction of labor through productivity leverage, where one person can handle what 5 or 10 people previously handled. It also could involve the entire elimination of a job category, as the work is being handled differently upstream or downstream in such a way that it eliminates the role.

Whatever the form AI displacement takes, we're in a period of rapid change. We're going to need new tools and perspectives to make the most of this moment, both for the livelihood of our families and communities and also for the more profound questions of how we choose to collectively construct our world.

CLIMATE DESTABILIZATION

Through our actions, the Earth system is rapidly shifting outside of the narrow band of stability that nurtured human civilization. What's coming isn't gradual warming—it's high volatility, with spikes, unusual plateaus, and unexpected disruptions. Every year we're seeing unprecedented floods,

fires, and droughts, and they will intensify in their instability throughout the course of the century.

These aren't freak occurrences; they are the shocks and reverberations of a system losing its equilibrium. As much has been made of how virtual the modern world is, ultimately everything in the physical economy was either mined or grown, which means it came directly from the natural environment. Our economy is a subset of our ecology, though we don't treat it as such in our economic designs, and as our ecology is imperiled, so goes our economic stability. Between economic and environmental shock, we will continue to see degradation of human welfare and dislocations in the larger biosphere.

What is currently understood to be a crisis of temperature will soon be understood to be a crisis of water. Not because we will have less water, but because we will have *more*. The main source of the freshwater we experience comes in the form of rain that supports our agriculture and rivers and streams that support our industry as those rains make their way across the landscape back to the ocean. There is more evaporation on a warmer planet, and that warmer atmosphere can carry more water vapor. Specifically, for each degree Celsius of warming, the atmosphere is able to hold 7% more water.[6] At this moment in history as we are passing 1.5°C of warming, so the atmosphere has the ability to hold 10% more water. This incremental increase represents a volume that is 60% of the volume of all rivers and streams on the planet.

This means a lot of water will be falling on Earth in new patterns: too much in one spot, too little in another. As anyone who has ever dealt with flooding in or around their home knows all too well, water is an incredibly powerful force in terms of physical impact and destructiveness. Water can

reshape landscapes and tear down infrastructure in hours. It could come in the form of an atmospheric river flooding areas with water beyond their carrying capacity. It could come from hurricanes strengthened by the warmer, wetter air. Or it could come from patterns of drought or glacial-fed rivers (e.g., Andes, Himalayas) retreating and drying up so that humans can't sustain themselves on the many landscapes fed by glacial rivers.

Changes to water will also become changes to fire. We live on a planet with a 21% oxygen atmosphere, a chemical oxidizer that allows biomass to burn. Any place on the planet with substantial stocks of dry biomass can be sites of massive fires. The upshot is that changes to water that are reducing water availability in unprecedented ways set the scene for fires and fire damage of unprecedented scale. In the past few years, multiple towns (Paradise, CA; Lytton, BC; Lahaina, HI) have been wiped off the map by fire. In Los Angeles, an area twice the size of Manhattan burned to the ground, with economic losses exceeding $250 billion (in 2025 dollars).[7]

What we had once spoken of in abstract terms (graphs and charts accompanied by stern warnings from scientists) is now being felt in sharp, visceral ways. We also see all the organizations that care concretely about risk actively adjusting their approaches. Insurance companies are shrinking coverage areas,[8] the military is spending on climate risks to infrastructure that could affect readiness,[9] and farmers are adjusting how they grow to keep up with more losses and destabilization.[10]

Herein lies one of the major shifts that is happening around the climate equation. For some time, it has been argued that it would be too expensive to address the changes to atmosphere and climate, but we are passing into the era

where it is too expensive to *not* address climate destabilization. The $250 billion in losses incurred in a week in Los Angeles massively outstrips the <$100 billion in clean-technology investments that the United States did over the latter 2.5 years of the Biden administration. An either/or question arises: Do we want to spend less money to build the climate-ready resilient infrastructure we need, or do we want to spend astoundingly more money trying to clean up the best we can after the human and natural fallout of a series of tragic disasters?

This new financial reality has not sunk in for most yet, but it is a fact that the damage (financially, socially, politically) we will see is going to quickly outstrip the cost of addressing the problem directly by an order of magnitude or more, and we'll soon have to face deeper questions around how such obvious mistakes and excuses were allowed to persist.

The stakes are much greater than the negative economic and human costs. The changes we have been making to the planet are driving extinction at roughly 1,000 times[11] the base rate, meaning that with each generation, we are leaving the planet ever more bereft of biodiversity. The gift of global biodiversity is one that took billions of years to create, and we have been taking a chainsaw to the tree of life, largely in an unintentional way. We find a few organisms we'd like to farm and it drives us to decimate landscape after landscape as we scale its production, only to realize later what we've driven to extinction, whether from habitat loss, interrupted migration routes, pesticide runoff, or active hunting.

The living diversity of this planet represents a body of wisdom of how life can thrive in every nook and niche on Earth, and it is stored in the genetic legacy of these beings, in their behaviors, learned and instinctual, in their languages and

cultures, in all the ways that these organisms communicate and live in community.

Our actions are burning down this multi-billion-year body of wisdom—an incalculable loss that, if we cannot course-correct, will be the biggest and most devastating legacy of this era in history. Society will change, technology will change, but the collapse of this planetary wealth will leave our descendants in a long-depleted landscape, and the 21st century will be the one they cite for when this wealth was lost forever.

This is the legacy we are creating, and it's clear that the magnitude of what we have been unconsciously creating through our destruction is about to substantially outstrip the magnitude of benefit we've created in our intentions. We had intended to use some of the world's natural resources for the project of making a stable and prosperous global civilization. We have instead crippled the biosphere in a way that threatens that very promise of stability and prosperity.

GLOBAL GEOPOLITICAL HARMONY

Rising ecological and economic pressures have global implications. Dislocations and migrations are increasing as climate destabilization makes areas of the planet uninhabitable. Trust in institutions is eroding. If we don't manage this transition with intention, we could slide into intensified conflict—between nations, between economic classes, between people and planet.

But if we do it right, we could move toward something humanity has never seen before: a global civilization capable of extended periods of peace, which continues to improve life for all beings in a widening circle of compassion

while building an economy that allows and supports ecological balance.

In times of uncertainty, it is easy and currently politically effective to scapegoat minorities as the cause of ills, as doing so appeals in bald-faced ways to the tribal parts of our brain. This is the part of social organization that is driven by in-group and out-group dynamics, loyalties, rivalries, and dominance. We have not yet developed an effective immune system to this type of cognitive impairment, but we know that there are limits to the quality of societies that can be built on tribal allegiances.

It may seem naive to presume that any individual can make a difference in geopolitical harmony, but I am colleagues with a number of the folks who developed the major social media services and their dopamine-loop exploiting algorithms and interactive designs. I've seen that industry stoke geopolitical chaos, and, in a way, this was things working as intended. The algorithm's goal was engagement, not political stability, scientific accuracy, or public benefit. It doesn't mean that these systems always drive toward negative impact on these nongoals, but they will always prioritize goals over the nongoals, and if meeting the more engagement goal means more political division and instability, then the algorithm will push for engagement—the goals—at the expense of nongoals. This is the algorithm working as designed. Given the chaotic nature of our current geopolitics, it's clear that the tools that I and others in my field have worked on have an influence here. We have definitely demonstrated several pathways for negative influence, where misinformation and division can be easily magnified by algorithm and interactive design changes. On the positive side, online organizing tools were credited for helping to catalyze the 2011 Arab Spring,[12] but the additional tools to support democratic participation and encourage social

stability were not in place. Again, things pan out in arbitrary ways on nongoals like social stability.

If ignoring society-destabilizing factors seems like a short-sighted way to design a system, you are right. And I hope to make it clear that we need more folks to hold these decisions to account, to develop effective immune systems to such manipulation, and to build the skills and access to be in decision rooms in the future, whether it's at the scale of globe-spanning technologies or at a local neighborhood meeting. The roots of harmony come from folks recognizing what still needs to improve and making different decisions at these critical junctures. It comes from gaining fluency in the prosocial behaviors and ways of connecting and leading that so that we can be a step more skillful in navigating the tangle into which we've tumbled.

I was fortunate to be on the internet before most were—first on BBS (bulletin board) systems in the 1980s, then the web in 1991. Back then it was not just HTTP but a bundle of protocols including FTP, GOPHER, TELNET, IRC, and NNTP. This was well before high-flying internet companies like Yahoo, Google, Amazon, and Facebook were founded. The internet consisted of a smallish number of links between research universities, and everyone on it was exploring these protocols. Many of us felt we were on the cusp of a new world.

This proved to be true. In short order, millions of webpages started to appear on every topic, and people around the world were finding common interests. An open web with democratized access to reading or publishing is exactly what we hoped to create and protect. It makes me (and many others that were part of that early tech generation[13]) sad to have seen the trajectory we could have had, only to end up in a very different world. Everything is owned by just

a few companies, and algorithms are evaluating us to direct us toward advertising and platform benefiting behavior, as opposed to us looking toward what most fascinates us—the things that surprise us and expand us, and the experiences that allow us to go deep, learn, create, and change.

For those of us who lived through this missed opportunity, the experience only clarified how systems of governance and power built on outdated frameworks can steer us unconsciously unless we bring their mechanisms to light and consideration in the design process. Specifically, we saw unprecedented creative energy and connection emerging as huge proportions of the world were able to express themselves online to anyone else in the world for the first time, through home pages, blogs, and online profiles. We saw new forms of publishing and a vast expansion of the media landscape. Two decades later, the internet has become a machine for advertisers to purchase attention, and human expression and connection is emphasized only when it drives more ad conversions and platform revisits. Attention is being handled without our intention, and the mechanisms to commandeer our attention have become a consistent cognitive drag force in every moment of modern life. We missed it. What we didn't examine clearly enough and seriously enough is how the economic frame that seeped into the technological tools we developed would eventually make the goal of these systems economic, not serving the people. We missed it because for a long time these two goals seemed to overlap well enough, and when they started to diverge, we had no effective reflection, improvement, and governance mechanism. In our older team designs, we had a customer experience and product group that was generally trying to make things that improve life. Without it, there'd be no draw to the tools at all. In contrast, the monetization and ads teams were looking to move metrics that improved net profitability. In this team layout,

the product and design teams tended to keep the voice and interests of the customer in the mix. As we moved to more algorithmic content presentation and algorithmically suggested user journeys, the balancing force started to wane, because the algorithm splinters the customer experience into as many experiences as there are people. User experience assessments became more statistical on aggregations of customer behavior, meaning a slower response loop, as it takes longer, for example, to see that algorithms are affecting self-image of young folks[14] than it does to test that a particular screen flow is confusing. This movement to algorithmic user experience led to practical slips in governance design and strongly highlighted the algorithmic goals of maximizing engagement and ad conversions. The degradation of these systems was almost preordained by this combination of goal and governance design. Slips in goal and governance design lead ultimately to degradation in system design. These are lessons that we should not painfully repeat.

In the same way, the economic frame doesn't govern itself. Bad goals and bad governance lead to bad system outcomes. The closest thing we have to a governance mechanism in the economy right now is the moral integrity of people in executive leadership, in boardrooms, and if things go awry, the ability for government regulation to rein in the most societally detracting elements of a technology or business model.

Maybe some of you bristled at the idea of moral integrity for CEOs or in boardrooms, because it hasn't been going great on that front recently. I agree. That said, I have been an executive leader and served on many boards, and I can say from direct experience that questions of intent and moral integrity definitely are discussed, usually earnestly, and often the right decisions are made. Yet I am still certain that

the checks and balances in place are woefully inadequate. You don't need a lot of decisions to go wrong to create massive negative societal impact, just a handful of bad decisions in the critical spots. It's time to look at the redesign of all these systems. Let's talk about a simple new design concept that could help provide better governance.

THE MEMBRANE IS PERMEABLE

As mentioned, I've spent a career developing new experiences on new technologies, and subsequently I'm well steeped in the user research required to understand whether a new experience is serving people or not. Most people at technology companies do intend to improve human life, but when the business model starts to hurt people, it creates tension with staying focused on improving life. So imagine a simple mechanism where for every major technology company, there are quarterly studies where 1,000 of their customers are picked at random and they fill out a 5-minute survey that allows us to empirically see if the tools and technology are beneficial or detracting from their lives. The company would not be in charge of picking the 1,000 customers; that would be done at random, and the people would be told that their input would be used toward improving that tool for the future.

This methodology could provide an independent assessment of how a technology is adding or subtracting to society, and if human benefit scores fall below a critical threshold, then it moves into the category of regulatory review. This additional governance mechanism might go a long way toward helping folks in the boardroom stand firm against changes that provide short-term financial gains, but ultimately hurt the folks who use the service. Of course, this is just one possible governance experiment; if it was tried

widely enough, we'd likely find some aspects that work and other areas to improve toward the goal of ensuring an independent gauge of societal utility. I mention it not as the only and best recommendation of how we improve governance but as an example of how simple the changes that could make a difference are. How utterly available to us they are if we are interested in the process of redesigning our world toward collective health and thriving. A lot of regulatory approaches have a high cost to enforce, but something like this survey is incredibly inexpensive (especially in contrast to the potential scale of societal damage of tools and platforms run amok) and can be done at any time with an independent assessment body using transparent research protocols.

In that same vein, let's take this survey concept and imagine a modification of it as an experiment in political governance. One of the current malaises of our political system is that it is highly driven by short media impressions, villainization of political personalities, and misrepresentations of intent. At the end of day, this process is meant to result in policy that aims to improve life for citizens and society as a whole. What if we took the 1,000 random respondents and pointed that methodology toward whether people understand the intent and effects of the laws being proposed. Currently laws are often written to obscure their intent—using "weasel words,"[15] as Carl Sagan would say. They gain popular support given a misunderstanding of the law's intent. Imagine if the content of such laws are presented to 1,000 folks at random, and they are then asked questions on their comprehension of the law's intent and its mechanisms of action. If more than 30% of people cannot understand the law above an 80% comprehension level, then it cannot be passed. Something like this would be a governance check that would require politicians and policymakers to be more direct and clear about the intent and effect of their

proposed laws. This comprehension check is based on the straightforward idea that democracy works better when people know what they are voting on. Simple enough. Others may have different ideas here, and that's perfectly fine. The point is that we know there are some acute dysfunctions in our current systems, and the process of experimenting toward concrete improvements that could address or obviate those dysfunctions is not particularly hard.

Nothing I've mentioned is expensive to do, and if it were able to head off even one wasteful bill from passing or one technology from driving unintended negative consequences, the cost of the research would be paid off 100 times over. The reason we are not in the practice of improving in this way has nothing to do with the cost or the difficulty; it is likely because those who currently benefit prefer an inefficient governance system as the inefficiency insulates against change. All the while, the vast majority who do not benefit like the current winners in the system have their energy of inquiry and change shunted over to political conversations. While some efficacy can come from that type of engagement, I know from direct experience that it takes far fewer people and less effort to solve problems directly through experimentation and research of the form just described. The membrane to change our reality is permeable. Cynicism and polarization benefit folks in power in the existing system because they both lead to inaction. Even if you believe you have great reasons to be aligned with cynicism and polarized perspectives, understand that this is the least effective place to act from, because it does not encourage action, experimentation, and learning. This book aims to give you more of the tools that can help solve these problems directly on behalf of your own communities, society, and the biosphere. So let's dive in another step.