Simulate the World

A Primer · Early Working Draft

Anselm Hook

Version 0.4 · June 2026

Contents

Chapter 1

Preface: The World

Early working draft (v0.4, June 2026). This text is unfinished and changing fast: expect rough prose, unverified citations, and sections that will move or disappear. Please don't circulate it widely yet.

> We face challenges too large for any individual mind to hold. This is the book's thesis: that simulation is not a specialist's toy but a form of civic infrastructure, and that making it accessible matters as much as making it rigorous.

Our intent in this short preface is to put one idea in front of you before anything else: what stands between serious models and the people they describe is not computation but publication. Then we get out of the way and let the book make its case.

Limits to social cognition

In January 2025, wildfires swept through Pacific Palisades and Altadena. According Cal Fire Damage Inspection over sixteen thousand structures burned. This was a tragedy and two years later the physical and emotional scars remain profoundly raw for the community. Could this have been mitigated? Could we have done better?

Regional governments and agencies do have a reasonable understanding of the science behind these fires today: fuel, slope, the potential for the Santa Ana winds to funnel through particular canyons, the specific shape of those canyons and historical wind patterns. But there isn't a tangible 'civic' model. Citizens, people impacted by these risks, have to trust agencies, hear them clearly, and truly imagine or understand the risks. These models are in government databases, agency systems, research groups — public in principle - but invisible for non-experts.

After the event there is civic clarity; concepts like fuel management, defensible space, evacuation routes. Can such clarity exist before this kind of event? We argue there could be one more civic tool; a model of how fire moves through this terrain under these winds. In effect a model, or simulation, in 3d, that lets people playtest the future, get a visceral feel for the risks. Today we have the computing power to build tools like this; to use brute force machine computation, and human intuition, to create a hybrid of human and machine reasoning that can help us 'see' better; make better decisions, possibly reach consensus better. That's what this primer is about.

This scenario, the lack of tangible civic models, that engage the stakeholders at their level, repeats over and over. In city halls, in watershed boards, utility commissions, school district meetings. Governments and corporations do have models, citizens do not. Fundamentally we live in a complex world, with

Figure 1.1. Pallisades Fire.

many events at different scales, velocities, interacting. The future is hard to predict. Often, at best, the outcomes of land-use law and policy decisions defy intended consequences. Bringing more minds to bear, bioregional understanding, local wisdom - along with satellite data, simulation data, can help improve our decision making. It is especially true that the stakeholders, some of whom are not always entirely obvious nor enfranchised, can contribute perspectives. Right now the primary tool citizens have is rhetoric; even a small improvement over that seems warranted.

The obstacle here is not only political; it is also structural, technical, fundamental. How do we think together as a group? How is information held, shared, visualized and acted on? The goal of this primer is to explore these questions - to explore how we can see and act on the future better.

The thesis

The thesis has two parts:

Part one: simulation is argument. A model of a system is more than a prediction machine: it is a structured claim about how the world works (Forrester 1971; Meadows et al. 1972). When a watershed manager says "if we increase pumping by thirty percent, groundwater will drop to critical levels by 2031," that is an argument with premises (the aquifer geometry, the recharge rate, the consumption estimate) and a conclusion. The argument is stronger if the premises are inspectable, if the uncertainty bounds are visible, and if you can push on the assumptions and watch what happens. A simulation that lets you do that does not merely forecast; it reasons in public.

Part two: accessibility is the unsolved problem. The computational tools to build serious models (system dynamics engines, agent-based frameworks, integrated assessment models) are mature and widely used. The gap is not in computation; it is in publication. Excellent data journalism does exist: arXiv preprints, NOAA's public dashboards, the data desks at the New York Times and the Guardian, and at its best it makes hard science legible to a wide public. But these are things you read, not things you run. What is still missing is a format that lets the model itself travel as a living document, running in the reader's browser, its assumptions exposed so that people who did not build it can inspect them and push back (Victor 2011). The consequence is that simulations exist in a narrow circuit of specialists, and the people most affected by the decisions they inform never get to touch one.

The book is about closing that gap. We state the thesis here, once. The rest of the book does not restate it: it proves it in the case studies, builds it in the chapter on the publication layer, and gives it a political home in the chapter on participation.

What this book is and is not

This is not a textbook on numerical methods or a software manual. We will cover enough of the underlying theory to understand what models can and cannot do: what it means for a model to be calibrated, how feedback loops produce behavior that surprises their designers (Meadows 2008), why uncertainty is not a failure but a result. We do this because you cannot use a tool wisely without some theory of how it works.

The primary thread is practical and civic. We are interested in cross-domain models that can be published, inspected, and argued with by the people who have a stake in them. We are interested in the participatory governance traditions that give that kind of modeling a political home. And we are interested in the specific, real scenarios (a drought in California, a warming Middle East, a bike lane that removes twelve parking spaces) where a well-framed model adds something rhetoric alone cannot.

> A note on rigor. We resist the temptation to simplify away the mathematics. Where the underlying dynamics matter (feedback loops, tipping points, stock-and-flow accounting), we include the derivation. Where precision is not the point and structure is, we say so explicitly. The book layers information: quick conceptual summaries up front, technical boxes and appendices for readers who want the full story.

A word on the title

"Simulate the World" is deliberately immodest. We are not proposing to build a complete digital twin of civilization; that project is ongoing at institutions with budgets many orders of magnitude larger than ours (Bratton, n.d.). We are proposing to build accessible, argued, inspectable models of pieces of the world that matter to specific communities in specific moments. The "world" in the title refers to the aspiration: that modeling our ecosystems should be a civic practice, not an elite one. A primer, as the subtitle has it, is where you start: by "primer" we mean a starting point and an invitation, not a schoolbook.

The book then follows a deliberate arc: it begins with what we value, our emotions around what we love, what we care about, our anxieties, then turns to the challenges that threaten it, looks at our emerging super-powers, shows where simulation has already changed outcomes, walks through how such models are built and published, and ends with how we participate. The arc runs from values to rigor to agency: what we care about, then how to reason about it honestly, then how to act on it together.

Further reading

- Sterling, B. (2000). “Viridian: The PlaNetwork Speech.” Viridian Note 00155. A tonal ancestor of this preface: Sterling opens on the Los Alamos wildfire as the emblem of a technical institution unable to escape the crisis it of all places should have foreseen.
- Vaillant, J. (2023). Fire Weather: A True Story from a Hotter World. Knopf. A narrative account of the 2016 Fort McMurray fire.
- Epstein, J. M., & Axtell, R. (1996). Growing Artificial Societies. MIT Press.
- Raworth, K. (2017). Doughnut Economics. Chelsea Green.
- Tainter, J. A. (1988). The Collapse of Complex Societies. Cambridge University Press.

Open questions

- What is the minimum viable “publication format” for a civic simulation?
- At what level of simplification does an illustrative model become misleading?

- How do we distinguish a model that helps a community reason from one that is used to launder a predetermined conclusion?

TODO / White-space

- FIGURE: Participation loop diagram — from data to model to argument to decision and back.
- PLACEHOLDER: Interview with a civic-tech practitioner or watershed manager (permissions required).
- PLACEHOLDER: Indigenous community voice on traditional ecological knowledge and model assumptions (permissions required).

Chapter 2

Our World

> Before we can simulate the world, we have to care about it. This chapter begins with the thing itself — the planet as a living system, the specific and personal ways we encounter it, and the values that motivate the rest of the book. It is not a chapter you can skip on the way to the technical content; it is the reason for all of it.

Our intent here is to start with the thing itself — the planet as a living system, and the specific, personal ways we each encounter it — because that is where the motivation for everything that follows comes from.

The first photographs

In 1967 the first full photographs of Earth were transmitted back from space. They were not the first images taken from orbit, but they were the first that showed the whole disk, in color, against the dark. The effect was disorienting in the way that genuine revelations are: we could suddenly see something that had always been there but never been visible from outside.

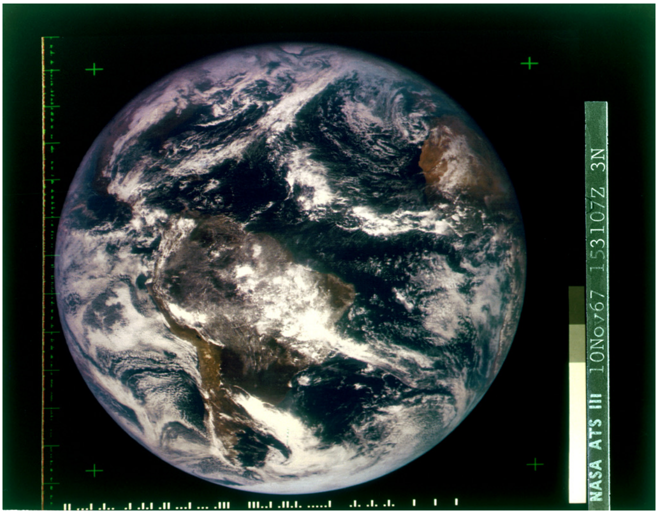

Figure 2.1. Earth from space, ATS-III satellite, November 1967: the first full-disk color photograph of the planet. NASA (public domain).

James Lovelock, working at JPL at the time, looked at those photographs and at the atmospheric chemistry data coming in from planetary probes, and coined the phrase that has been controversial and productive in equal measure: Gaia Theory — the hypothesis that life on Earth does not merely adapt to its environment but actively regulates it, maintaining the conditions for its own persistence through an intricate web of feedback and exchange (Lovelock 1979). The atmosphere's composition, the ocean's salinity, the surface temperature of the planet: all of these, Lovelock argued, are held within bounds that permit life to continue, and life is the mechanism doing the holding.

The controversy around Gaia has never quite died. The teleological flavor of the original formulation bothers people who prefer their science free of anything that sounds like purpose — and Dawkins and Gould both objected that evolution by natural selection produces no mechanism for planetary-scale altruism. The subsequent refinement, developed largely with Lynn Margulis, shifted the framing: Gaia as a regulatory property that emerges from the aggregate of self-interested organisms interacting with their physical environment, not a conscious goal (Lovelock and Margulis 1974). Geophysiologist Tim Lenton has developed this into the more rigorous concept of "Earth system regulation" — the observation that planetary variables like surface temperature and atmospheric composition have stayed within habitable ranges for billions of years despite perturbations that should have driven them out, and that biology is an active part of the mechanism maintaining them (2016).

As far as we know, this is the only place in the universe where that is true. The thin surface of this planet — a biosphere perhaps twenty kilometers deep from the deepest ocean trench to the upper atmosphere — is the only substrate we have found where the 3.8-billion-year experiment in organized complexity called life is underway.

That is worth starting with.

Life as a trajectory against entropy

The second law of thermodynamics describes a universe drifting toward disorder. Energy gradients dissipate; structures break down; information is lost. Life is not an exception to this law but a peculiar way of accelerating it locally while building order at the cost of exporting entropy elsewhere. A plant takes in sunlight, CO_2, and water and builds glucose: locally more ordered than the materials it started from, with the waste heat and oxygen paying the entropic price.

The trajectory from the first self-replicating molecules in the early ocean to the extraordinary diversity of the present biosphere — every beetle, every fungus, every phytoplankton bloom — is four billion years of compounding this trick. Each new form of life creates new gradients, new niches, new possibilities for still more complexity. The Cambrian explosion, some 541 million years ago, saw the sudden proliferation of animal body plans in the fossil record; it may represent a threshold above which multicellular complexity became self-amplifying. The appearance of eukaryotic cells with internal membranes and organelles; the colonization of land; the coevolution of flowers and their pollinators: each of these was a phase transition in the complexity budget of the biosphere.

Humans are a recent and peculiar addition to this story: a species that has externalized its nervous system through language, writing, and now computation, and that can build models of the systems it lives inside and act on what those models show.

Whether that capacity is enough to navigate what comes next is the question the rest of this book circles.

Individual experience as data

We are warmed by the same sun and breathe the same air. Simple acts, eating food or drinking water, connect us to an extraordinary chain of phenomena. Our reasoning and our values are shaped by this world.

Yet each of our experiences is unique. We each touch a different part of the world, and together those experiences might contribute to a larger picture. Some of us grew up near mountains, others in cities, some on prairies or in small towns or along coasts. As a result we each value different things, feel the loss of different things, and bring different evidence to the table when we reason together about what to protect.

> For me, it was growing up in Alberta, near the Rockies. The mountains provided the backdrop for my most vivid memories. That shapes what motivates this work. What are your most powerful memories? How many of your most vivid experiences are of times spent outside — exploring, hiking, running — as opposed to times at the office? Which would you like to protect most?

That question is not rhetorical. It is the appropriate starting point for any conversation about what a model is supposed to be optimizing for. The participation traditions we turn to later

Figure 2.2. The author's own backdrop: Lake Louise, Alberta, 2006.

(Chapter 9) are partly about creating structures where different pieces of local knowledge — the rancher's intuition about stream behavior, the indigenous community's multigenerational record of fish runs, the urban hydrologist's rainfall data — can be combined rather than overridden.

Going inside

There is another way to encounter the Earth system, one that is nearly the inverse of the satellite photograph. Instead of receding until the whole becomes visible, you descend until a small piece of it becomes overwhelming in its density and connection.

I came to scuba diving late, and it reconfigured something I thought I already understood. A coral reef is a city. It has districts and thoroughfares and specialists, cleaners and predators and prey, and a constant cycling of nutrients through roles that took millions of years to negotiate. A single reef head the size of a dining table sustains a community of hundreds of species with specific spatial arrangements and behavioral agreements that are invisible unless you slow down and look.

The first time I dived the waters off the Pacific side of Mexico, I surfaced from a cleaning station, a spot on the reef where small cleaner wrasse (finger-sized reef fish) pick parasites off larger fish, which queue patiently and open their gills to be inspected. I came up with a visceral understanding of what mutualism means: an arrangement between two completely different organisms, neither of which ever negotiated it, that reduces harm for both and has been running longer than our civilization. The reef keeps the ledger without writing anything down.

Later, at Socorro Island in the Pacific, I encountered oceanic manta rays, animals that spend most of their lives in open water but return regularly to particular reef cleaning stations. Watching a manta with a three-meter wingspan hold still while wrasse worked across its skin, I saw plainly that the open ocean and the reef are not separate places. They are coupled nodes in one system. The manta needs both; damage either and you damage the circuit.

Marine biologist Dr. Andrea Marshall, who has spent decades tagging and tracking pelagic mantas, describes their movements as tracing the invisible structure of the ocean: warm currents, productive upwelling zones, the reef stations they return to across hundreds of kilometers. Each manta is a moving data point in a planetary ecosystem model. What the satellite photograph shows from outside, the manta's track shows from inside: the Earth as a network of flows and exchanges, held together by relationships that have no natural human-scale representation.

Reefs are the second most biodiverse ecosystems on the planet (after tropical rainforests), cover less than 1% of the ocean floor, and support an estimated 25% of all marine species. They are also, as we will see in Chapter 3, among the most precisely documented tipping points in the Earth system. A tipping point is a threshold past which a system flips into a different state and does not easily flip back: warm a reef gradually and for a while little changes, until sustained heat stress crosses a limit and the coral expels the symbiotic algae it lives on, bleaches, and often dies. That limit is measured in degree-heating weeks, and it is now routinely crossed across large regions of the Indo-Pacific. Understanding the reef is, among other things, practice for understanding how tipping points work, and why we should care about them.

Do you feel our world is worth protecting?

It is worth spending time with this question before proceeding. The public conversation about what is at stake, and what to do, is loud and contradictory. Navigating it requires having defined, even roughly, what you actually care about.

All emotional responses to the crisis are reasonable, and none of them is, by itself, adequate to it. What we need is not just feeling but the ability to act on it, and for that we need tools that let us reason clearly about systems with many moving, interacting parts.

> Quick intuition — Complex adaptive systems and the value of diversity Ecology has a rough empirical rule: diverse systems are more resilient than simple ones. A monoculture responds uniformly to a pathogen; a polyculture has some individuals with resistance. The same applies to human communities and economic systems: variety is insurance. When we think about what to protect in the Earth system, biodiversity is not just an aesthetic preference — it is a functional requirement for stability under perturbation.

Planetary boundaries: a framework for what we are protecting

In 2009, Johan Rockström and colleagues at the Stockholm Resilience Centre proposed the concept of planetary boundaries — a framework identifying nine Earth system processes that regulate the stability of the Holocene-like conditions in which civilization developed, and estimating "safe operating spaces" for each (Rockström et al. 2009).

The nine boundaries: 1. Climate change (atmospheric CO_2 and radiative forcing) 2. Novel entities (synthetic chemicals, plastics, radioactive materials) 3. Stratospheric ozone depletion 4. Atmospheric aerosol loading 5. Ocean acidification 6. Freshwater change (blue and green water flows) 7. Land-system change (forest cover) 8. Biosphere integrity (functional and genetic diversity) 9. Biogeochemical flows (nitrogen and phosphorus cycles)

The 2023 update found that six of the nine boundaries had been crossed (Richardson et al. 2023). The concept is useful as a civic tool as much as a scientific one: it translates diffuse planetary processes into a legible dashboard that communities and governments can orient around.

Kate Raworth's doughnut economics framework, discussed in Chapter 9, builds the planetary boundaries ceiling into a two-sided model: social foundation below, planetary boundaries above, with a "safe and just space for humanity" in between. The doughnut is less a model than a frame for asking the right questions.

FIGURE: Planetary boundaries diagram — Stockholm Resilience Centre, 2023 update (placeholder)

The Gaia debate and Earth system science today

Current Earth system science treats the planet as a coupled system of physical, chemical, and biological processes with feedbacks operating at multiple timescales. The key insight retained from

Lovelock and Margulis is that biology is not merely embedded in the Earth's geochemistry but actively shapes it.

This matters for simulation because it means that models which treat the biosphere as passive — assuming that the physical processes drive everything and the biology just responds — are systematically wrong. The coupled feedbacks are the interesting dynamics:

- The forest-moisture feedback in the Amazon: tropical forest transpiration generates rainfall; deforestation reduces that rainfall, stressing the remaining forest, potentially triggering a tipping point at around 20–40% deforestation.
- The ice-albedo feedback at the poles: melting ice exposes darker ocean or land, absorbing more solar radiation, accelerating melting.
- The marine biological carbon pump: phytoplankton at the ocean surface take up CO_2 and, when they die, sink it toward the seafloor. Changes in ocean temperature and chemistry alter phytoplankton abundance and therefore the strength of this pump.

Each of these is a place where simulation earns its keep: the feedback dynamics are invisible to intuition, the timescales are long enough that consequences are deferred past any election cycle, and the nonlinearities can produce abrupt transitions that simple linear extrapolation would never predict.

Further reading

- Margulis, L. (1998). Symbiotic Planet: A New Look at Evolution. Basic Books.
- Raworth, K. (2017). Doughnut Economics. Chelsea Green.
- Wilson, E. O. (2016). Half-Earth: Our Planet's Fight for Life. Liveright.

Open questions

- Is the Holocene the right baseline for "safe" Earth system conditions, or is it too conservative given where we already are?
- How should non-human stakeholders — ecosystems, species — be represented in civic models?
- What is the relationship between place-based individual experience and planetary-scale modeling?

TODO / White-space

- FIGURE: Full-disk Earth photograph (ATS-III, 1967 — public domain).
- FIGURE: Planetary boundaries diagram (Stockholm Resilience Centre, 2023).
- FIGURE: Underwater reef photograph — coral city / cleaning station (placeholder; see assets/images/ for `Amed_Bali_Indonesia-14.jpg` as interim).

- PLACEHOLDER: Interview with ecologist or indigenous land steward on what "caring for the world" means in practice.
- TODO: Expand the Cambrian explosion / biodiversity section with more specific examples.
- TODO: Consider adding a short passage on the emotional register of encountering a living reef versus a bleached one — grief as data.

Chapter 3

The Emotional Arc: Mixed Messages, Real Feelings

> The way we feel about the planetary crisis shapes what we choose to reason about. This chapter examines the emotional landscape produced by the environmental media — the grief, anger, nihilism, and latent hope — and argues that naming these feelings is not a detour from the practical work but a precondition for it. The chapter also introduces the advocacy tone that runs through the rest of the book: not angry, but clear; not passive, but grounded.

Our intent here is to take the emotional register seriously — to name the grief, anger, and nihilism honestly — because reason is a tool we bring to bear on things we already care about, and what we let ourselves feel determines what we are willing to think about.

What we hear

Turn on the news on any given day and you will find the following: a beautiful and fragile planet; a litany of the things we are doing to it; a set of experts warning that time is running out; and a set of opposing voices insisting the crisis is exaggerated, politically motivated, or economically self-defeating.

Ocean acidification, pollution, sea-level rise, habitat loss, extreme weather, overpopulation, CO_2 buildup, drought, the Great Pacific Garbage Patch. We hear that corporate interests are cutting down the rainforests in the Amazon and Borneo, that monoculture farming is dangerous, that we are relying on antibiotics too much, that low-lying coastal regions are increasingly vulnerable to floods, that large-scale refugee migrations may destabilize countries. We hear about tipping points. We hear that methane hydrates trapped under the oceans may be released as the planet heats, suffocating life under a methane sky. There is even the popular idea of "the sixth extinction" — that the Earth has had many large extinctions before, some caused by similar kinds of self-induced criticalities, and that this may just be another one.

Even our experts signal concern in ways that feel more like counsel for escape than counsel for engagement. Stephen Hawking, in various talks, warned that "it will be difficult enough to avoid disaster on Earth in the next hundred years" and that the long-term survival of the species requires colonizing other planets. Elon Musk and Peter Thiel have funded space ventures partly framed in terms of backup copies of civilization.

Meanwhile, other people are hearing an entirely different message: that environmental regulations are blocking important economic projects, that climate science is politically captured, that the solutions being proposed will destroy more livelihoods than the problem. Either way, there is a sense of an opposing "other" — and the remarkable thing is that this "other" is, in significant part, ourselves.

How does that make you feel?

The question is not rhetorical and it is not therapy. It is epistemological.

There is a well-documented set of psychological responses to large-scale environmental challenges (Clayton et al. 2017; Pihkala 2022):

Eco-anxiety — a chronic fear of environmental doom that can shade into paralysis. A 2021 global survey of ten thousand young people (Hickman et al., The Lancet Planetary Health) found that 59% were very or extremely worried about climate change, and 45% reported that climate anxiety affected their daily functioning (2021). These are not marginal cases; this is the modal emotional state of an entire cohort of people who will live longer with the consequences than any generation before them.

Ecological grief — mourning for specific places, species, and ways of life that have already been lost. This is not irrational; it is a reasonable response to real loss. The challenge is that grieving and acting require different emotional orientations, and you can get stuck in the first without reaching the second.

Nihilism and fatalism — a sense that the situation is so far gone, or that individual action is so futile, that engagement is pointless. This is probably the most common form of disengagement in educated, well-informed people. It is a psychologically protective response to overwhelming information (Toffler 1970), and it is almost always factually wrong — the future is not yet settled, and the range of outcomes between the best and worst scenarios is enormous.

Anger — at the institutions, industries, or political actors who have blocked or delayed action. This is also reasonable, and sometimes clarifying. The risk is that it can become its own form of paralysis, or can generate an us-versus-them dynamic that makes coalition-building harder.

Quiet hope — at the scale and creativity of the response emerging in communities, research groups, and policy circles around the world. Often underreported relative to the catastrophe narrative.

> What I want to focus on here isn't the crisis itself but our response to it. Our response — or lack thereof — is ultimately all that matters. For whatever reasons there appears to be a certain kind of emotional reaction being produced. In many there is a kind of nihilism or melancholia, a sense of futility; in others, different reactions. Few of us are unmoved.
> For better or worse, our emotions drive what we choose to reason about. Rationalism is rarely in opposition to emotion; we choose to get rational and pragmatic about things that we are

> passionate about, that we feel empowered to tackle. It is worth spending time with these emotions rather than trying to avoid them. By looking at how we feel — by creating valid emotional space around our own reactions — we can get clarity on how we choose to act. It is almost as if we need a planetary therapist.

The advocacy register

This book has a tone. It is worth being explicit about it.

We are not neutral. We believe the planetary situation is serious, that the scientific evidence is solid, and that the challenges we face require serious collective effort. We are not going to pretend otherwise in the name of appearing unbiased, because there is no unbiased position — refusing to take a position is itself a position, one that implicitly favors the status quo.

At the same time, we are not angry. Anger is a legitimate response, but this book is primarily written for people who already understand the problem and are looking for tools — models, methods, civic infrastructure — to help address it. We are not trying to convince climate skeptics; we are trying to help the people who already know something is wrong reason more clearly about it and act more effectively.

The tone we are aiming for is something like this: urgent without being panicked, rigorously honest about uncertainty, willing to name what is at stake, and genuinely hopeful about specific tractable paths forward. We have things that actually work

— modeling tools, civic participation methods, open data — and those things deserve clear-headed advocacy, not either catastrophism or denial.

The emotional arc of the book follows the arc that exists in the original simulate.world presentation:

1. What we value — the world as it is, what we love about it, why it matters.
2. What is at stake — the specific challenges, their scale, their interconnections.
3. What we can do — the tools, models, governance methods, and civic infrastructure.
4. How we participate — the democratic traditions, from anticipatory democracy to computational governance, that give this work political meaning.

Each stage has its own emotional register. The first is warmth and appreciation. The second is honest confrontation with difficult information. The third is the satisfaction of rigorous problem-framing. The fourth is something like collective agency — the feeling of being part of a generation that faced a large problem and chose to engage it rather than look away.

Why the media framing is part of the problem

The dominant media narrative around environmental crisis has a structural problem that is independent of whether the underlying science is right.

The narrative consistently presents large, diffuse challenges (climate change, biodiversity loss, ocean acidification) as things that are happening to us, not things that are being collectively produced by decisions that could in principle be made differently. The villain, in the narrative, is usually industry, or government, or "humanity" in the abstract — entities too large to feel agency about. The scale of the problem is presented as a source of paralysis rather than a reason to mobilize (Morton 2013).

There is also a systematic preference for drama over information density. A wildfire is a crisis; the absence of a wildfire that was prevented by better land management is invisible. A flood is news; the slow, invisible depletion of an aquifer over thirty years is not, until the aquifer is gone and ten thousand farmers have no water. This produces a systematic bias in what the public knows about environmental systems: we know about the acute events and not the chronic processes; we know about the disasters and not the avoided disasters; we know about the failures and not the working solutions.

This is exactly the problem that civic simulation addresses. A model makes the chronic process visible. It shows the future state of the aquifer before the collapse, not after. It makes the avoided disaster legible. It connects the small decision (an almond farmer's water right) to the large system (the Central Valley's water budget) in a way that newspaper reporting cannot.

Superpowers we are only beginning to use

The previous section makes this sound bleak. It is not, quite.

Just as we are starting to have a real impact on the world, we are also developing surprising capabilities. These are powers that sit in deep tension with our sense of powerlessness.

Civic applications. The growth of social networks, distributed computing, and open data platforms has given ordinary citizens tools for collective sense-making and coordination that did not exist before. Crisis mapping after the Haiti earthquake, coordinated by thousands of online volunteers, produced better maps faster than the professional response. Wikipedia. OpenStreetMap, now more detailed than many national mapping agencies in the regions where volunteers are concentrated. iNaturalist, a citizen biodiversity observation platform with over 200 million species observations, has become a significant source of scientific data — several species have been rediscovered or newly described from observations submitted by amateur naturalists. These are not perfect, but they demonstrate that distributed human computation at scale is possible and can produce genuinely useful outputs.

Earth observation. Satellite data of the planet's surface, once the exclusive domain of superpowers with military space programs, is now available to researchers and (increasingly) to the public in near-real-time. Global Forest Watch shows deforestation as it happens, at 30-meter resolution, anywhere on Earth. Google Earth Engine provides a planetary-scale computing platform for environmental analysis, accessible via a JavaScript API.

Figure 3.1. Crisis mapping after the 2010 Haiti earthquake: thousands of online volunteers produced better maps, faster, than the professional response.

NOAA's Coral Reef Watch publishes daily global sea surface temperature anomaly maps that reef managers in the Indo-Pacific use to prepare for bleaching events. The ECMWF's ERA5 reanalysis dataset gives hourly global climate variables — temperature, wind, humidity, precipitation — going back to 1940. We have never had anything like this before.

Computation. Moore's Law has run long enough that what used to require a supercomputer now runs on a laptop. The reduced-complexity climate model FaIR — which can produce global temperature projections that track the IPCC's full ESM ensemble within their uncertainty bounds — runs in a Python notebook in under a second. System dynamics engines and agent-based frameworks are available as open-source software that a biology student can install in an afternoon. The friction between "I want to understand this system" and "I have a working model" has dropped by orders of magnitude. What remains high is the friction between "I have a model" and "I have a model that someone else can use" — and that is exactly the publication gap this book addresses.

Figure 3.2. Gordon Moore. Decades of his namesake law put supercomputer-class modeling on a laptop.

AI authorship. A new capability, still being absorbed: AI systems that can write, debug, explain, and translate simulation code. The barrier between describing a model in plain language and having a running implementation has been reduced to something close to a conversation. This is not a solved problem — AI-generated models introduce errors that require expert review, and the tendency of language models to produce plausible-looking but wrong equations is a real hazard in scientific work. But the direction is unambiguous: the technical skill required to build a first working simulation of a system is declining rapidly, which changes who can participate in civic modeling and how quickly.

The combination of better modeling infrastructure, accessible web publication, and AI-assisted authorship creates an opportunity that did not exist even five years ago. The tools are available. The data is available. The platforms for civic participation are available. What is missing is the connection — the layer that makes a simulation legible to the people who have a stake in it.

That is what this book is about.

Further reading

- Sterling, B. (2000). "Viridian: The PlaNetwork Speech." Viridian Note 00155. A model of the advocacy register this chapter aims for: wry, urgent, unsentimental about the crisis and unembarrassed about caring.
- Solnit, R. (2016). Hope in the Dark. Haymarket Books.
- Klein, N. (2014). This Changes Everything: Capitalism vs. The Climate. Simon & Schuster.
- Bendell, J. (2018). "Deep Adaptation." IFLAS Occasional Paper 2.

Open questions

- What is the most effective emotional register for civic simulation tools — is it hopeful, alarming, neutral?
- How do we design simulations that convey urgency without triggering the psychological avoidance responses that urgency often produces?
- Is there an ethical obligation to frame uncertainty in ways that do not amplify nihilism?

TODO / White-space

- FIGURE: The emotional arc diagram — values → crisis → tools → participation (placeholder).
- PLACEHOLDER: Interview with climate psychologist or community organizer on the emotional dimensions of environmental work.
- TODO: Add a closing paragraph connecting the superpowers section back to the book's thesis — the tools exist; what remains is building the connection to civic process.

Chapter 4

Existential Threats: The Scale of What We Are Facing

> This chapter confronts the problems honestly. We examine the specific nature of the planetary challenges we face — their scale, their nonlinearity, the psychological difficulty of holding them clearly in mind — and introduce the concepts needed to reason about them: tipping points, hyperobjects, the Anthropocene, and the historical record of extinction events. We are not trying to produce despair but to produce accuracy.

Our intent here is to confront the problems accurately rather than to produce despair — to give you the concepts (tipping points, hyperobjects, tail risk) that make these challenges thinkable instead of merely overwhelming.

A new class of challenge

The challenges we face are, in a meaningful sense, new. Not merely larger versions of the challenges that previous generations faced — drought, plague, war, resource scarcity — but structurally different: a new class of problem that has some properties that make them very hard to deal with using the cognitive and political tools we have.

Three properties distinguish them:

Scale. The relevant dynamics operate at the level of the whole Earth system — the atmosphere, the ocean, the global economy, the biosphere. No individual nation, community, or person produced these problems alone, and no individual actor can solve them alone. The spatial scale of action required does not match the spatial scale at which human political and cultural institutions operate.

Temporal displacement. The costs are deferred in time in ways that break our political systems. The benefits of burning carbon accrue now; the costs land decades hence. The benefits of aquifer depletion accrue to today's farmers; the costs land on the communities of 2050 who find no water. Democratic institutions discount the future sharply, because voters live in the present and demand attention to present costs and benefits. The long-range term is structurally invisible in most decision processes.

Nonlinearity. Small changes can produce catastrophic and irreversible shifts. The climate system, the Amazon forest, the Atlantic meridional overturning circulation (AMOC), permafrost methane: all of these contain feedback dynamics that can flip from a stable state to a very different stable state faster than linear extrapolation would suggest. A model that treats the system as linear — treating the next unit of CO_2 as having the same effect as the last one — will be systematically wrong near tipping points.

Hyperobjects

The philosopher Timothy Morton introduced the term hyperobjects to describe entities that are “massively distributed in time and space relative to humans” (2013). Global warming is his canonical example. It is not something you can see or touch directly. You can see its effects — a wildfire, a flood, a failed harvest — but the thing itself is too large and distributed to perceive. It occupies multiple timescales simultaneously and manifests differently depending on where you stand.

Morton argues that hyperobjects force a revision of our philosophical and cognitive tools. They are real — more real, in some sense, than the individual events we can observe — but they violate our normal categories. They are “sticky” (they adhere to things they come into contact with), they are “nonlocal” (their effects appear in multiple places that seem unconnected), they “phase” (you encounter different aspects depending on where and when you look), and they “inter-objectify” (they redefine the relationships between other objects).

For our purposes the key practical implication is this: hyperobjects require models to think about. You cannot observe global warming directly; you can only observe its local manifestations. You cannot predict a wildfire from looking at a day’s weather; you need to see the accumulated moisture deficit across seasons,

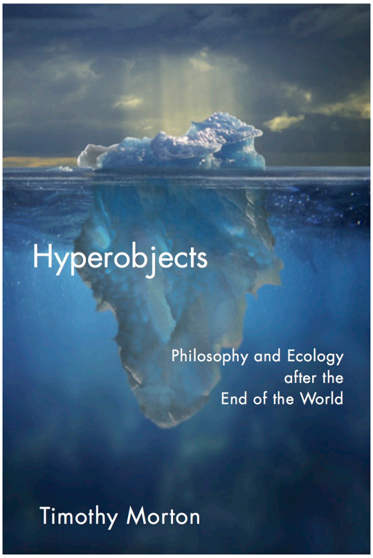

Figure 4.1. *Hyperobjects* (Morton 2013): some things are too large and distributed in time and space to perceive directly. Book cover.

the fire history, the fuel load, the wind patterns. Assembling that picture requires computation. It requires models. The model is how a community develops a shared perception of something too large to see.

The Anthropocene

We even have a label for our era: the Anthropocene.

The Anthropocene is an informal geologic chronological term for the proposed epoch that began when human activities had a significant global impact on Earth's ecosystems. Its defining signature: human activities have changed Earth's biodiversity and geology enough to be visible in the future geological record.

We have a planet-straddling civilization that uses two-fifths of the arable land for food crops, controls the flow of most major rivers, has cities so dense that we can see the outlines of the continents from space at night. We have chemically altered the balance of constituent elements in the atmosphere and are measurably altering the planet's surface temperature. The stratigraphy of the future will show this: elevated radioactive isotopes from nuclear testing, a global layer of microplastics, the mass die-off of large animals, a sharp increase in carbon in ice cores (Zalasiewicz et al. 2010).

The term has been proposed for formal ratification by the International Commission on Stratigraphy. The debate is not about whether the human signature is real but where precisely to place the start date: 1610 (the "Orbis spike," a carbon dip from the demographic collapse of the Americas post-colonization), 1945 (the first nuclear detonations), or the 1950s "Great Acceleration."

The practical implication: we are not separate from the Earth system; we are now one of its major drivers. The question is whether we will remain an uncontrolled perturbation or become something more like what the biosphere has always been — a complex adaptive system that regulates its own conditions for persistence.

The geological record of extinction

The history of life on Earth is punctuated by catastrophic die-offs. Jack Sepkoski and David Raup identified five mass extinctions by looking at the fossil record (1982):

Extinction	Period	Estimated species loss
Ordovician–Silurian	~443 Ma	60–70%
Late Devonian	~360 Ma	~70%
Permian–Triassic	~252 Ma	90–96%
Triassic–Jurassic	~201 Ma	70–75%
Cretaceous–Paleogene	~66 Ma	~75%

The largest, the Permian–Triassic event — the "Great Dying" — killed 57% of all families, 83% of all genera, and up to 96%

of all species. Life on land took thirty million years to recover in terms of diversity. The cause was probably a combination of Siberian Traps volcanism, ocean anoxia, and methane release — a cascade that began with one trigger and amplified through feedbacks.

The Cenozoic (our era) began with the recovery from the Cretaceous–Paleogene event that killed the non-avian dinosaurs. It is characterized by cooler temperatures and the emergence of mammals. We may be entering what some researchers call the sixth mass extinction: a global die-off with current rates of species loss estimated at 100–1,000 times background, driven by habitat destruction, overexploitation, pollution, invasive species, and climate change. The WWF Living Planet Index — which tracks the average abundance of vertebrate populations, not species extinctions — found a 73% average decline between 1970 and 2020 (WWF 2024). That is not a count of extinct species; it is a measure of how much of the living world that was present fifty years ago is still present. Seventy-three percent gone in half a century.

> A note on framing. The framing of a "sixth extinction" can cut two ways. Presented as historical inevitability — "mass extinctions are normal" — it implies fatalism. Presented as a measure of where we are in a process that might still change course, it is clarifying. We prefer the second reading. The range of outcomes between "business as usual" and "serious collective action" is not small; the difference matters, to millions of species and billions of people.

Tipping points and nonlinearity

The concept of a tipping point — a threshold beyond which a system shifts abruptly to a qualitatively different state — is central to understanding why linear extrapolation fails for complex Earth systems.

A simple example: an ice sheet is in equilibrium when the rate of snow accumulation equals the rate of melting. As temperatures rise, the melting rate increases. At some temperature, the ice sheet begins to contribute its own warming through the ice-albedo feedback (less ice → less reflection → more warming → less ice). The system has a tipping point: below it, warming produces gradual change; above it, warming is self-amplifying and the ice sheet can collapse much faster than the rate of forcing would suggest.

Several major Earth system tipping points are identified in the literature (Lenton et al. 2008; Armstrong McKay et al. 2022):

- Amazon dieback. Tropical deforestation reduces the forest's ability to recycle moisture, reducing rainfall, stressing the remaining forest. Estimates of the threshold at which the Amazon transitions from a net carbon sink to a net carbon source range from 17–40% deforestation (current: ~20%).
- AMOC collapse. The Atlantic Meridional Overturning Circulation — the "conveyor belt" of ocean circulation —

weakens as the North Atlantic freshens from Greenland ice melt. Collapse would dramatically change temperature and rainfall patterns across Europe and the North Atlantic. A 2023 study (Boers, Nature Communications) found statistical indicators of a loss of resilience in AMOC that are consistent with an approach to a tipping point (Boers 2021); a separate 2024 study (van Westen et al., Science Advances) using a higher-resolution model found a collapse possible within decades under continued warming (Westen et al. 2024). The science remains actively contested, but the direction of evidence is not reassuring.

- Permafrost thaw. Permafrost in Arctic and sub-Arctic regions contains large stores of organic carbon. As it thaws, microbial decomposition releases CO_2 and methane — a warming feedback. The methane is particularly concerning because of its high short-term warming potential.
- West Antarctic ice sheet. Marine-based ice that has been identified as potentially unstable under warming; its loss would add several meters to sea level.
- Coral reef systems. Tropical reefs host an estimated 25% of all marine species and are among the most precisely quantified tipping points in the Earth system. Their threshold is measured and exceeded in real time.

> Math derivation — Coral bleaching and degree heating weeks
> Coral bleaching occurs when sea surface temperature (SST) exceeds the local maximum monthly mean (MMM) — the long-

term average of the warmest month — by at least 1°C for a sustained period. NOAA quantifies thermal stress using Degree Heating Weeks (DHW):

$$\mathrm{DHW}(t) = \sum_{\tau=t-11\ \mathrm{wk}}^{t} \max(0,\ \mathrm{SST}(\tau) - \mathrm{MMM} - 1°C) \cdot \frac{1\ \mathrm{wk}}{1}$$

The thresholds that have been empirically validated across reefs worldwide:

DHW (°C-weeks)	Expected response
0–3	Minor stress; possible minor bleaching
4–7	Significant bleaching likely
8+	Severe bleaching and mortality likely

> The mechanism: corals are animals in a mutualistic relationship with photosynthetic algae called zooxanthellae (family Symbiodiniaceae), which live in coral tissue and provide up to 90% of the coral's energy. When water temperature rises, the zooxanthellae produce toxic reactive oxygen species; the coral expels them as a stress response, turning white ("bleaching"). Without their algal partners, corals slowly starve. At DHW ≥ 8, mass mortality typically follows.
>
> The tipping point character comes from the ecosystem recovery dynamics. A bleached reef takes ten to fifteen years to recover under favorable conditions. Back-to-back bleaching events — which are becoming more frequent as baseline SST rises — can kill a reef before it recovers, driving a phase shift from a coral-dominated to an algae-dominated state that can persist for decades. The Great Barrier Reef experienced mass bleaching events in 1998, 2002, 2016, 2017, 2020, 2022, and 2024; each reset the recovery clock (Hughes et al. 2017).
>
> This is a fold bifurcation of the same type described in the minimal tipping point model above: the "current state" (coral-dominated) and the "alternative state" (algae-dominated) are both stable, with a threshold between them. Once the threshold is crossed and the coral population crashes, recovery requires not merely reducing the thermal stress but returning to conditions substantially cooler than the tipping point — because the alternative stable state is self-reinforcing.

The relationship between these tipping points is also concerning. There is growing evidence that they are coupled — that triggering one can accelerate others — creating the possibility of a cascade that is self-sustaining beyond human control.

Math derivation — A minimal tipping point model
A tipping point can be illustrated with a simple fold bifurcation. Let x represent the state of a system (e.g., ice extent) and r a slowly changing parameter (e.g., temperature). The dynamics are:

$$\frac{dx}{dt} = r - x^2 + \text{noise}$$

At $r > 0$ there is no stable fixed point on the positive side; at $r < 0$ there are two: one stable (the current state) and one unstable (the tipping threshold). As r increases toward zero, the stable and unstable points approach each other until they collide and annihilate — and the system falls to a new state. The critical slowing down before the tipping point (the system takes longer to return to equilibrium after small perturbations) is one of the early warning signals researchers look for in observational data.

Under a Green Sky: methane and Black Swans

Paleontologist Peter Ward's Under a Green Sky (2007) describes a scenario that is worth knowing: evidence suggests that the Permian–Triassic extinction may have been partly caused by a methane inversion in the atmosphere — a rapid release of methane hydrates from ocean sediments, which converted the sky from blue to green as hydrogen sulfide-producing bacteria bloomed in a warming, anoxic ocean.

Ward's point is not that this will happen again soon. It is that there is a broader category of Black Swan events — low-

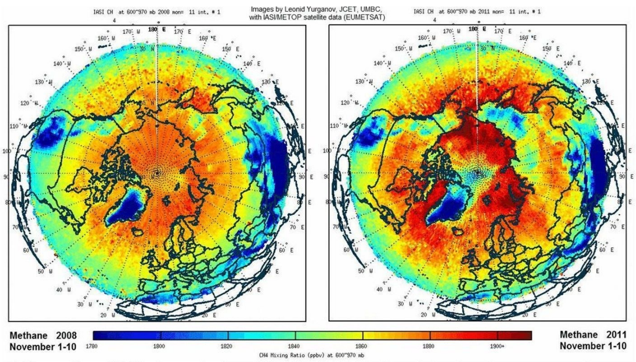

Figure 2. AIRS atmospheric methane concentration data showing an anomalous buildup of methane over the Arctic region from November 2008 to November 2011 Diagram from Yurganov 2012 in Carana 2012.

Figure 4.2. *Under a Green Sky*: Peter Ward's account of past mass extinctions and the methane-driven anoxic ocean.

probability, high-consequence, tail-risk scenarios — that our models and political processes systematically underestimate. We tend to plan around the expected case; Black Swan analysis asks what happens if the distribution is fat-tailed and the tail lands on us.

This matters for modeling because the most important things to model are often the scenarios that are not central cases — the ones where the distribution has a long tail. A civic simulation that shows only the median future is less useful than one that shows the distribution of outcomes, including the catastrophic ones.

The broader point: there are events possible in our future that are beyond our ability to anticipate. The goal is not to simulate everything but to build enough shared understanding that communities can respond more intelligently — and faster —

to the things that are actually happening.

Why we simplify the world

There is a reason these nonlinear systems feel so hard, and it is not a moral failing. Human cognition is metabolically expensive and individually bounded. We do not have the calories, the working memory, or the time to hold a coupled, nonlinear system in mind, and we have never found a way to network our individual minds into a group intelligence that could. A threshold the math crosses smoothly is, to us, a surprise we did not have the compute to see coming. So we do the only thing available: we cut the world down to something we can reason about.

The most consequential example is how we feed ourselves. A mature ecosystem is a dense web of higher-order interactions, plants and fungi and microbes and animals trading nutrients through relationships negotiated over millions of years. We cannot compute that, so we replaced it with something we can: the monoculture field, one crop, measured inputs, a measured yield, legible precisely because it has been stripped of the complexity we could not manage. The legibility is the whole appeal. The cost is resilience, soil, and most of the living web that used to be there.

It does not have to be this way, and the proof is that it has sometimes been otherwise. Knowledge accumulated over thousands of years can encode complexity that no individual could derive from scratch. The Amazonian terra preta soils (dark, charcoal-rich, microbially alive, still fertile centuries after they were made) are the residue of communities that built food forests rather than fields, biasing a rainforest toward the species they wanted without flattening it into a monoculture. Julia Watson's Lo-TEK catalogues dozens of such indigenous technologies: systems that work with complexity instead of erasing it, and that read as superstition only until you have a model rich enough to see what they are doing. (The Balinese water temples in Chapter 4 are exactly this: a coordination scheme that looked like religion and turned out to be optimization.)

This is the deeper argument under everything that follows. We simplify because, unaided, we have had to. But we are entering a period in which intelligence is increasingly available on tap, in which brute-force machine computation can hold more of a system's complexity than any person can, and in which the pairing of human intuition with that machine capacity might let us climb out of the local minimum that radical simplification represents. Nicholas Negroponte once named the shift of his era as a move from atoms to bits; the shift this book cares about is from systems kept simple enough to manage by hand to systems we can finally afford to understand closer to their real complexity. None of this guarantees the tools get used well. It is the more basic claim that, for the first time, the option is on the table.

Further reading

- Kolbert, E. (2014). The Sixth Extinction: An Unnatural History. Henry Holt.
- Tainter, J. A. (1988). The Collapse of Complex Societies. Cambridge University Press.
- Taleb, N. N. (2012). Antifragile: Things That Gain from Disorder. Random House.
- Armstrong McKay, D. I., et al. (2022). "Exceeding 1.5°C global warming could trigger multiple climate tipping points." Science, 377(6611).
- Watson, J. (2019). Lo-TEK: Design by Radical Indigenism. Taschen. (Indigenous technologies that work with complexity rather than erasing it.)
- Mann, C. C. (2005). 1491: New Revelations of the Americas Before Columbus. Knopf. (On terra preta and the engineered Amazon.)
- Negroponte, N. (1995). Being Digital. Knopf. (The "atoms to bits" shift.)

Open questions

- Is there a "safe" rate of planetary change, or do threshold effects mean that the question is binary?
- How should civic models communicate tipping point risk without triggering either panic or fatalism?

- What is the relationship between planetary-scale tipping points and local/regional tipping points that communities can actually act on?

TODO / White-space

- FIGURE: Mass extinction timeline diagram (Phanerozoic eon, major extinction events).
- FIGURE: Climate tipping points map (from Lenton et al. or similar).
- FIGURE: Bifurcation diagram for tipping point model.
- TODO: Add NOAA Coral Reef Watch DHW map as a figure — a live global heat-stress map that reads like a patient's vital signs.
- TODO: Add bifurcation diagram figure (programmatic — see TODO.md).
- PLACEHOLDER: Interview with paleontologist or Earth system scientist on the geological record as a model for future risks.

Chapter 5

Case Studies: When Simulation Changes the Outcome

> Abstract arguments for civic simulation are less persuasive than concrete examples of it working. This chapter assembles a set of case studies — from the Balinese water temples to California's water crisis to Yellowstone's wolves — that show what civic modeling looks like when it is built, argued, and acted on. These are the proof that simulation changes outcomes; they are also the pivot from the crisis of the previous chapters to the agency of the chapters ahead.

Our intent here is to make the argument concrete: to show real cases where a model changed a civic outcome, or where the absence of one let an avoidable harm happen — and to draw out what made each case work.

What makes a good scenario

Not every interesting question makes a good simulation. The ones worth building tend to share several properties:

Specific. A real place, a real issue, real stakes, competing stakeholders — not a generic case study. “Water management in arid regions” is not a scenario; “the effect of almond acreage expansion on San Joaquin Valley aquifer depletion, 2010–2040” is.

Intersectional. The interesting behavior arises from coupled dynamics that resist simple rhetoric. If a spreadsheet and a press release are adequate, a simulation adds nothing. The scenario earns its keep when the relevant dynamics cross domains — when a land-use decision affects hydrology affects food prices affects migration.

Diffuse benefits or deferred costs. These are exactly the cases where colloquial political reasoning fails systematically and a model can show something a room full of people cannot easily see.

Counterfactual leverage. The model’s job is often to make visible the thing that did not happen: the flood that was averted, the housing that was not built, the aquifer that will not be there in thirty years. These are structurally invisible to rhetoric-based politics.

Accessible audience. A model of Tuvalu’s freshwater lens has a smaller natural audience than a model of a bike lane on a street people live on. Both are worth building, but priorities matter.

Case study: the Balinese water temples

Figure 5.1. Balinese water temple. The temple network coordinated irrigation across a watershed for centuries. Source: arounddeglobe.com.

Probably the founding example of civic computation for environmental management is Steve Lansing's work in Bali (1991).

Lansing arrived in Bali in the 1980s during a period of disruption: the Green Revolution had brought intensive farming techniques to communities around the world, and the Balinese government had adopted them. But the new practices — higher-yield rice varieties, continuous cropping, synthetic fertilizers — were not working in Bali. Yields were declining; pest outbreaks were rising.

The traditional Balinese system had been organized around a network of water temples — hierarchical social institutions that coordinated the timing of water flows and fallows across thousands of farmers in a watershed. Priests at different levels of the temple hierarchy coordinated planting schedules that created synchronized fallow periods across the watershed, which, it turned out, starved pest populations and recharged soil nutrients.

The Green Revolution's "rational" approach — each farmer maximizing their own planting schedule — destroyed the coordination and amplified the pest dynamics.

Lansing built a computer simulation of the traditional system and was able to demonstrate that the temple-coordinated schedule was near-optimal for pest suppression across the watershed — a result that would have taken field experiments decades to produce but could be shown computationally in weeks (1993). The simulation became the argument; the argument changed government policy; the temples were reintegrated into water management.

The lesson is not that technology should defer to tradition. It is that emergent optimization in a social-ecological system can look like superstition from outside unless you have a model that makes the feedbacks visible. The farmers and priests had the coordination right; they lacked the language to explain it to outside observers; a simulation provided the translation.

Case study: California's water crisis

California's water situation is a case study in intersectional complexity — and in the failure of political systems to process information that arrives too slowly and is distributed across too many stakeholders.

A few numbers (circa 2015, but the structure has not changed) (USDA Economic Research Service 2009):

- Total water reserves (snowpack, rivers, reservoirs, groundwater, soil) are approximately 34 million acre-feet below normal.
- Reserves are declining at roughly 12 million acre-feet per year.
- Two-thirds of losses are attributed to groundwater pumping in the Central Valley.
- Agriculture accounts for 80% of water consumption.
- California produces over 90% of all US almonds, walnuts, pistachios, broccoli, strawberries, grapes, and tomatoes.
- Almonds alone consume ~600 gallons of water per pound — more than 25 times the water required per pound of tomatoes.
- When aquifers are emptied by pumping, the land sinks — up to a foot per year in some regions; damage from land subsidence in one area cost $1.3 billion.
- The water being pumped may be between 10,000 and 30,000 years old. When aquifers collapse, they often cannot be refilled.

At the same time, San Diego was spending a billion dollars to build a desalination plant while farmers were using aquifer water to grow almonds at prices that could not justify the cost of desalinated water. The resource allocation was, by any systems analysis, incoherent — but the political system had no mechanism for making the incoherence visible, because each actor was responding to their own local incentives.

A civic simulation of California water would need to model: - Snowpack and seasonal runoff from the Sierras - Reservoir storage and release schedules - Groundwater recharge and extraction by crop type and region - Economic return per acre-foot for different crops - Water rights allocation under the prior appropriation system - The feedback between depletion and land subsidence - Future climate scenarios affecting snowpack

This is not a trivial model, but most of the data exists. The reason it has not been built as a public civic tool is not that the computation is too hard; it is that nobody has prioritized building the publication layer that would make it legible to the actual stakeholders.

Case study: wolves in Yellowstone

In 1995, fourteen gray wolves were released into Yellowstone National Park, the first of several reintroductions after a seventy-year absence; the park's wolves had been hunted out by 1926. The immediate controversy was predictable: ranchers near the

Figure 5.2. Gray wolves in Yellowstone. Whether their return "changed the rivers" is now a live scientific debate. Source: savebiogems.org.

park feared livestock losses. The result that made the reintroduction famous took longer to appear, and it has turned out to be more contested than the famous version admits.

The famous version goes like this. With wolves back, elk stopped lingering in exposed valleys and riverbanks; willows, aspens, and cottonwoods regenerated; the recovered vegetation stabilized the banks; the rivers narrowed and deepened; beaver returned, and their ponds created habitat for fish, otters, and ducks. A single predator, reintroduced, reorganized the hydrology of a landscape. This is the story told by the video "How Wolves Change Rivers," watched tens of millions of times, and by the foundational trophic-cascade papers of William Ripple and Robert Beschta (2012; 2016). It is a textbook illustration of a trophic cascade: a change at the top of the food web propagating downward in ways direct-effect reasoning would never predict.

It may also be too clean. Over twenty years, a Colorado State team led by N. Thompson Hobbs and David Cooper ran a field experiment built to test it, fencing willow stands to exclude elk and building artificial beaver dams to restore water tables, then watching what actually limited recovery. The finding (K. N. Marshall, Hobbs, and Cooper 2013; Hobbs et al. 2024) was that removing the elk was not enough. Willows recovered only where the hydrology recovered too, where the dams raised the water table to what it had been when beaver still engineered these streams. The seventy-year absence of wolves had pushed the system into what the authors call an alternative stable state, and reintroducing the predator did not, on its own, reverse it. Ripple and Beschta dispute the experiment's design and reach (a formal Comment and Reply ran in Ecological Monographs in 2026 (Beschta et al. 2026; Hobbs et al. 2026)), so the argument remains live. Nobody claims the wolves did nothing; the elk herd fell by more than half. What is contested is whether a single causal chain (wolves to willows to rivers) captures what happened, or papers over a messier, hydrology-first, context-dependent reality.

That contest is the reason this case belongs in a book about simulation, and it is the opposite of an embarrassment. The clean cascade is exactly the kind of story that travels well and models badly. Real ecosystems are not chains of pairwise links; they are sets of interactions that fire only under particular combinations of species, water, and disturbance, what a recent review frames as higher-order interactions, better represented by hypergraphs than by simple networks (Vidiella, Duran-Nebreda, and Valverde 2026) (see Chapter 5). A model that hard-codes "wolves → willows → rivers" will reproduce the appealing story and inherit its blind spot. A model worth publishing does the harder thing: it holds the competing hypotheses (browsing-limited versus hydrology-limited recovery) against the same field data, and shows which one the evidence supports and where it remains genuinely uncertain.

So the civic value here is not a tidy lesson about predators. It is a method. The rancher's concern (livestock loss) is local and immediate; the ecologist's argument (ecosystem recovery) is diffuse, delayed, and, it turns out, contested even among ecologists. A shared, inspectable model does not settle that by decree. It makes the disagreement legible: here are the mechanisms, here is the data each side leans on, here is what changes if hydrology rather than browsing is the limiting factor. That is a better starting point for a decision than a video that has already made up its mind.

> Reintroducing wolves is contentious, and not only between ranchers and ecologists. A model's job in a case like this is not to certify the appealing story but to expose the assumptions underneath it, so that a community arguing about predators is at least arguing about the same evidence.

FIGURE: Yellowstone — the popular trophic cascade beside the hydrology-limited alternative (two causal diagrams of one system) (placeholder)

Case study: the Grand Coulee Dam

Figure 5.3. The Grand Coulee Dam. Built without a fish ladder, it ended salmon migration over 1,100 miles of river. Source: popularlogistics.com.

The Grand Coulee Dam on the Columbia River, completed in 1942, is a monument to the infrastructure values of its era: hydroelectric power, irrigation, flood control, economic development. It flooded 21,000 acres of prime bottomland where Native American communities had lived and fished for thousands of years. Because it was built without a fish ladder, it permanently blocked salmon migration, removing over 1,100 miles of natural spawning habitat.

The catch at Kettle Falls — a primary Indigenous fishing ground that was inundated — had been over 600,000 salmon per year. The loss was permanent and total. In one Army Corps of

Engineers study, annual losses were estimated at over one million fish. Chinook, Steelhead, Sockeye, and Coho salmon were eliminated from the upper Columbia. The Canadian government's official position was that "there is no commercial salmon fishery on the Columbia River, and hence Canadian interests in that respect will not be affected."

The parks in the interior British Columbia watersheds are named "Grizzly" — presumably for the bears that once populated the region. The grizzly, dependent on salmon, are gone. Nobody who built the dam considered the grizzly. Nobody modeled the trophic dependencies. Nobody quantified the long-term loss of protein to Indigenous communities against the short-term economic gain from hydroelectric power.

This is not a historical curiosity. The structural failure — large infrastructure project, diffuse and invisible costs, concentrated and visible benefits, absent stakeholders with no voice — is still the standard pattern in public works decision-making. Civic simulation does not fix the politics. But it can make the invisible stakeholders — the absent salmon, the future residents of the flooded valley, the communities downstream — visible enough to demand a response.

Case study: Rajasthan's wells and the Johad

A more hopeful case: in Rajasthan's Alwar district, wells had dried up, pushing communities into abject poverty. A local or-

Figure 5.4. Alwar district, Rajasthan, where reviving traditional johad earthen dams brought dry wells and a seasonal river back to life. Source: delhi-carrental.com.

ganization, the Tarun Bharat Sangh, began reviving traditional earthen dams called johads to capture rainwater and recharge the underground water supply (EcoTippingPoints, n.d.).

The results cascaded:

- Groundwater levels rose and wells came back to life.
- With water came agriculture and economic recovery.
- The seasonal river Arvari, dry for decades, began to flow again year-round.
- Traditional community practices of collective water management revived.
- Wildlife returned; the spiritual and cultural life of the villages was restored.

Nearly a thousand villages have now followed Gopalpura's example. The water recovery was also a social and cultural recovery — poverty and environmental degradation had been coupled, and healing one healed the other.

The lesson: the relationship between environment and community welfare is bidirectional and strongly coupled. Models that treat ecological restoration as separate from economic welfare will miss the feedbacks that make restoration pay.

Case study: coral reef restoration and emergent structure

Coral reefs are the second most biodiverse ecosystems on the planet, and they are collapsing in real time. An estimated 50% of the world's reef cover has been lost since 1950; thermal bleaching events that were once once-a-generation occurrences are now recurring every three to five years in large portions of the Indo-Pacific (Hughes et al. 2017). The question that reef restoration practitioners face is not only "how do we grow more coral" but "how do we build reef structure that is resilient to the next bleaching event" — and that is fundamentally a simulation problem.

The organization Emergent Reefs deploys custom-designed artificial reef modules in threatened reef systems. The core design challenge is that a reef is not simply a pile of rock; it is a three-dimensional structure whose geometry determines current flow, sediment dynamics, light penetration, and the diversity of niches available to colonizing organisms. Structure precedes biology: get the physical form right, and the community assembles itself. Get it wrong, and the artificial reef becomes an algae-covered surface with low biodiversity.

This is an engineering problem with deep simulation potential. Three questions worth modeling:

1. Structural complexity and fish habitat. Reef rugosity — the ratio of actual surface area to projected horizontal area — is one of the strongest predictors of reef fish biodiversity. A reef with high rugosity has many crevices, overhangs, and passages that create distinct microhabitats: shelter for juveniles, ambush positions for predators, cleaning stations, spawning aggregation sites. Computational fluid dynamics (CFD) can simulate water flow through proposed reef structures, predicting where sediment will deposit (bad: it smothers coral), where current will concentrate food, and what thermal stratification will look like. The

output is a ranked set of geometries to test in the field.

2. Bleaching resilience through genetic diversity. Not all coral colonies bleach at the same temperature. Populations in historically warmer waters often host zooxanthellae strains with higher thermal tolerance; some coral individuals have genetic profiles that confer resilience. A restoration program that plants a genetically diverse set of coral fragments — sourced from thermally distinct environments — creates a reef population with a broader response distribution: some colonies die in a bleaching event, but others survive and become the seed stock for recovery.

An agent-based model of this dynamic is tractable: each coral colony is an agent with a stochastic bleaching threshold drawn from a distribution; a thermal stress event is applied; survivors reproduce at a rate that depends on the substrate available (Hughes et al. 2019). The question the model answers: what initial genetic diversity is sufficient to ensure reef persistence under a given warming trajectory? How does that threshold change under different bleaching event frequencies?

3. Connectivity and keystone species. Reefs do not function in isolation. Pelagic species — sharks, manta rays, large groupers — visit reefs as part of wide-ranging circuits, and their presence maintains predator-prey dynamics that prevent algal overgrowth. Dr. Andrea Marshall's work tracking oceanic manta rays at Socorro Island, Cocos Island, and Revillagigedo demonstrates that mantas maintain site fidelity to specific cleaning stations over decades, and that their presence is a proxy indicator of reef health (A. D. Marshall et al. 2011; A. D. Marshall, Dudgeon, and Ben-

nett 2011). A simulation that models the connectivity between offshore productive zones and reef cleaning stations would capture something that a reef-only model misses: the dependence of reef health on the open-ocean ecosystem that feeds it.

This is the civic modeling opportunity at reef scale: not just a thermal stress threshold model, but a full socio-ecological model that couples water temperature (from NOAA Coral Reef Watch DHW data), reef community composition (from field surveys or remote sensing), fishing pressure, and restoration intervention rate (Darling et al. 2019). The question it answers for a reef manager: given the warming trajectory for this specific reef system over the next thirty years, what combination of thermal refugia selection, genetic diversity in plantings, and fishing pressure reduction maximizes the probability that a living coral system persists?

> Quick intuition — Why structure precedes community A natural reef takes centuries to build because each organism depends on the physical substrate created by the previous generation. Artificial reefs short-circuit this by providing substrate immediately — but the geometry of that substrate determines what community can establish. The simulation opportunity is at the design stage, before concrete is poured, when the geometry is still a variable.

From historical cases to a forward-looking catalogue

The case studies above are historical — they show simulation changing outcomes, or its absence allowing harm. The natural next question is what to build now. That forward-looking catalogue — scenarios worth building, organized by scale (local, regional, planetary), with buildability estimates and a suggested build order — lives in the reference section (see Reference — Scenario Library). We keep it there rather than here so this chapter stays focused on what the historical cases teach, and so the catalogue can grow and be re-prioritized without disturbing the argument.

Further reading

- "How Wolves Change Rivers" (2014). Sustainable Human (video, narration adapted from George Monbiot). The popular telling whose tidy causality the Yellowstone section above complicates.
- Vaillant, J. (2023). Fire Weather: A True Story from a Hotter World. Knopf. Narrative account of the 2016 Fort McMurray wildfire; companion reading to the fire scenarios in the reference library.
- Reisner, M. (1986). Cadillac Desert: The American West and Its Disappearing Water. Viking.

- Ostrom, E. (1990). Governing the Commons. Cambridge University Press.
- Epstein, J. M. (2007). Generative Social Science. Princeton University Press.

Open questions

- When does an illustrative model become misleading? What obligations come with publishing one?
- How do we design civic scenarios that invite genuine engagement rather than just validating the builder's assumptions?
- How can scenario outputs be structured to support the absent-stakeholder problem — giving non-human or future interests a voice?

TODO / White-space

- FIGURE: Balinese water temple irrigation network schematic (placeholder).
- FIGURE: Yellowstone — two causal diagrams side by side (the popular wolves→willows→rivers cascade vs. the hydrology-limited / alternative-stable-state account).
- FIGURE: California water budget Sankey diagram.
- TODO: Add BP Deepwater Horizon case study (privatized risk, socialized loss).
- TODO: Add Lake Merritt case study (tidal reconnection).

- TODO: Add specific Emergent Reefs project data once available — deployment locations, species colonization timelines, bleaching survival rates.
- FIGURE: Reef rugosity diagram — high vs. low rugosity cross-sections with habitat niches labeled (placeholder).
- FIGURE: DHW thermal stress map for Indo-Pacific showing bleaching event frequency 1998–2024 (NOAA Coral Reef Watch public data).
- PLACEHOLDER: Interview with Balinese water priest or researcher on the relationship between traditional knowledge and computational modeling.
- PLACEHOLDER: Interview with Central Valley farmer or water rights lawyer on the California case.
- PLACEHOLDER: Interview with Andrea Marshall / Marine Megafauna Foundation on pelagic-reef connectivity and what a civic-accessible manta tracking model would look like (see Reference — Voices).

Chapter 6

The Landscape of Simulation: Paradigms and Lineage

> This chapter maps the intellectual and technical landscape of socio-environmental simulation. Where did these tools come from? What are the major modeling paradigms and when is each appropriate? And where is the gap — the territory that existing tools do not cover — that civic simulation occupies?

Our intent here is to give you a working map of the field — where these tools came from, the handful of paradigms worth knowing, and where the full-blown Earth-system models sit — so that the rest of the book has somewhere to stand. (The catalogue of specific tools, institutions, and people lives in the reference section.)

The field and its names

There is a large and well-funded world of simulation software — but most of it models factories, warehouses, hospitals, and logistics networks. The thing it optimizes for is throughput. A different and smaller tradition has always been interested in a harder problem: what happens to the world over time when humans and natural systems push on each other?

This field goes by several names — systems modeling, integrated assessment, socio-environmental simulation, complexity economics — and the practitioners rarely all know each other. What they share is a concern with feedback, delay, and nonlinear collapse that industrial simulation mostly ignores. The most consequential behavior in complex systems — overshoot, oscillation, irreversible transition — arises precisely from these features.

The lineage: Forrester, Meadows, and *Limits to Growth*

The intellectual origin is Jay Forrester at MIT in the 1960s, who invented system dynamics: a way of representing the world as stocks (things that accumulate) and flows (rates of change) connected by feedback loops (Forrester 1971). Delays and nonlinearities in those loops are where all the interesting behavior lives — overshoot, collapse, oscillation.

Forrester's students Donella and Dennis Meadows applied this to the whole planet. Their 1972 report for the Club of Rome, Limits to Growth, used the World3 model to simulate interactions between population, food production, industrialization, pollution, and resource depletion across twelve scenarios to 2100 (1972). It sold 12 million copies in 37 languages and remains the founding document of the field. The model has been updated three times; its central predictions — that overshoot followed by decline is the default trajectory if resource constraints are not addressed — have held up better than its critics expected.

The Club of Rome — founded in 1968 as a gathering of scientists, economists, and public figures concerned with long-range planetary futures — remains an institutional anchor.

> A note on the reception of Limits to Growth. The book was heavily and often dishonestly criticized. Many of the critiques mischaracterized what the model actually predicted — claiming it said we would run out of oil by 1992, when in fact it modeled resource depletion under a range of scenarios and did not predict specific dates. The critique was politically convenient for interests that preferred no limits. This history is worth knowing because it will recur with any civic simulation that reaches a broad audience: contested models attract contested attacks, and the attacks will often misrepresent what the model said.

Modeling paradigms

System dynamics (SD)

The Forrester/Meadows lineage. The world is stocks and flows; feedback loops drive behavior (Meadows 2008). Best for understanding delays, overshoot, and policy resistance at a macro level. Not well suited to individual actors or spatial heterogeneity.

Tools: Vensim (industry standard), Stella Architect (accessible; widely used in education), PySD (runs Vensim/Stella models in Python).

Good for: Climate-economy feedbacks, demographic transitions, resource depletion, supply chain oscillations, pandemic dynamics.

Agent-based modeling (ABM)

Complements system dynamics when you want discrete actors — households, firms, governments, species — making decisions under uncertainty. Emergent behavior arises from local interactions rather than being specified top-down (Epstein and Axtell 1996).

Tools: NetLogo (lingua franca of ABM research; huge model library), Mesa (Python ecosystem), Agents.jl (Julia; performance-oriented), GAMA Platform (strongest when geography is central; deep GIS integration), Repast HPC (large-scale), AnyLogic (commercial; combines SD, ABM, and discrete-event in one), FLAME GPU 2 (GPU-accelerated; millions of agents).

Good for: Disease spread, social network dynamics, market microstructure, urban pedestrian behavior, fishery management, species habitat models.

Input-output (I-O) analysis

Goes back to Wassily Leontief's 1940s work: an economy as a matrix where every sector's output is another's input. Shocking a row (a port closes, a crop fails) propagates mechanically through the matrix. Simple, transparent, and still the workhorse for tracing upstream and downstream impacts.

Tools: EXIOBASE (multi-region I-O with environmental extensions; 200 countries × 200 sectors), GTAP (the standard in trade policy analysis), UN Comtrade + custom matrices.

Good for: Supply chain disruption analysis, ecological footprint of consumption, trade policy impact, carbon accounting through supply chains.

Computable general equilibrium (CGE)

Extends I-O with price responses and market-clearing. More realistic economics, but heavier data requirements and less transparent dynamics. GTAP is the dominant framework for trade policy analysis.

Good for: Trade policy, fiscal policy, energy transition impacts on sectoral employment.

Integrated assessment models (IAMs)

Couple climate physics to economics — the modeling technology behind IPCC scenario analysis. DICE (Nordhaus), MESSAGE, REMIND, and GCAM are the major examples. They smooth over heterogeneity and spatial detail in exchange for tractability; they are the tools that produce the headline scenario numbers in IPCC reports.

Relationship to this book: IAMs are not civic tools — they require expert operation and do not publish in interactive formats. But they produce the scenario vocabulary (SSP, RCP) and temperature pathway data that simplified civic models can draw on.

> Reduced-complexity / simple climate models (FaIR, MAGICC, Hector). These emulate the full IAM/ESM ensemble with a handful of ordinary differential equations. The IPCC uses them to produce the large probabilistic ensembles behind the headline temperature ranges, because you cannot run a GCM ten thousand times. This is the buildable tier for civic applications.

Earth System Models: the top of the computational pyramid

Earth System Models (ESMs) and General Circulation Models (GCMs) are the full climate simulations run by NOAA, ECMWF, NASA, and similar institutions.

A GCM divides the Earth into a three-dimensional grid of boxes, spanning tens of kilometers horizontally and layered vertically through the atmosphere and ocean. At every grid point, supercomputers solve the primitive equations of geophysical fluid dynamics — based on the Navier-Stokes equations and laws of thermodynamics — calculating wind velocity, pressure, temperature, and humidity at every timestep.

Processes too small to resolve at the grid scale (cloud formation, turbulence, individual thunderstorms) are approximated using parameterizations — empirically derived equations that represent the aggregate statistical behavior of sub-grid processes.

An ESM extends the GCM by adding active biogeochemical cycles: how carbon sinks (forests, phytoplankton) and chemical processes interact with the physical climate. It is the most complete representation of the Earth system available.

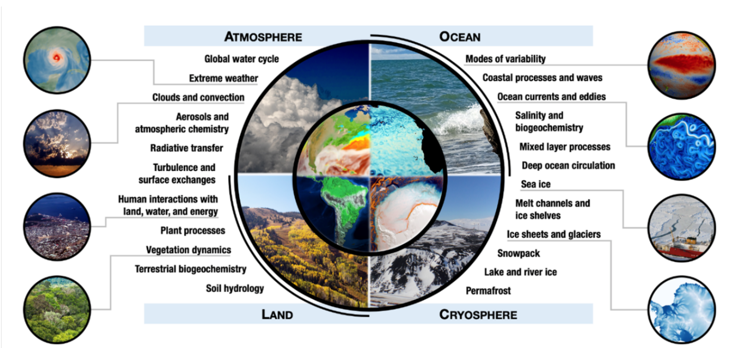

Figure 6.1. Schematic of an Earth System Model: coupled atmosphere, ocean, land, ice, and biogeochemistry. Source: US DOE.

ESMs are not browser-runnable, not even close. A single CMIP6 simulation run might take months on a top-tier supercomputing cluster. The CMIP6 ensemble — the set of simulations used in the IPCC Sixth Assessment Report — represents hundreds of thousands of CPU-years of compute. This is a tier of tool that informs civic modeling but does not replace it.

The right relationship between civic models and ESMs is: ESMs provide the physically-grounded scenario outputs (how much warming, how much sea-level rise, how much precipitation change, under what emissions trajectories) that simplified civic models then translate into local and sectoral impact assessments.

Beyond pairwise links: higher-order interactions

The paradigms above mostly represent a system as a set of two-body relationships: predator and prey, sector and sector, stock and flow. Many real interactions are not pairwise. An effect can appear only under a particular combination of species and conditions, and vanish when any one of them changes. The Yellowstone case in Chapter 4 is a live example: whether removing elk lets willows recover depends jointly on browsing pressure, the water table, and beaver-engineered hydrology, and no two-arrow diagram captures it.

A growing body of work represents these as higher-order interactions using hypergraphs, where one hyperedge links several species and environmental conditions at once instead of chaining pairwise edges. A 2026 review casts ecological structure this way: as feasible interacting sets that form, fragment, and reorganize under disturbance, a better fit for non-equilibrium ecosystems than the equilibrium-around-a-stable-state assumption that classical models inherit (Vidiella, Duran-Nebreda, and Valverde 2026). The practical takeaway for civic modeling: when a system's behavior depends on context, a model built only from pairwise links will look authoritative and mislead, and it is worth knowing a more faithful representation exists.

Emerging world models: AI and the digital twin

A newer strand brings machine learning to simulation itself. In deep learning, a world model is an internal representation of an environment that lets an agent predict future states and evaluate actions without testing them in the world (Ha and Schmidhuber 2018). As of mid-2026 most of this research concentrates on near-field physical systems (robotics, gaming, 3D scene generation), but the idea is being pushed outward, and two developments show where it is going.

The field is often described by two taxonomies:

Functional: - Renderers — generate realistic images or video; may not maintain consistent physics. - Simulators — emphasize structural and physical accuracy; useful for digital twins and physical AI. - Planners — use the internal simulation to evaluate sequences of actions and optimize toward a goal.

Architectural: - Latent dynamics models — compress the world into a latent state and predict evolution over time (Dreamer, RSSM). - Joint embedding predictive architectures

(JEPA) — predict representations of future states rather than pixels, reducing compute cost. - Generative interactive environments — train on large video datasets to create controllable worlds.

Spatial intelligence. Fei-Fei Li's World Labs, founded in 2024, argues that machines must "master space before language," and is building what it calls Large World Models (LWMs): systems that generate and reason about interactive 3D environments rather than flat images (2024). Its first product, Marble, builds explorable 3D worlds from a prompt. Today these are models of local physics, the gravity, mass, light, and fluid behavior of a room or a street. The label "world model" therefore collides with the sense this book uses, where a model of the world means coupled climate, ecology, and economy at planetary scale. That collision is temporary. Trained against a large enough corpus of real-world data, physics-grounded world models stop being demonstrations of local physics and start becoming faithful, queryable representations of how real places behave. A few years out they may fill the shoes the name promises, and converge with the civic-simulation aspiration from the machine-learning side.

Generalization across systems. Traditional forecasting builds a separate model per system: one for the Beijing metro, another for freeway sensors, another for bike-share demand. A 2026 framework, LLM-UTP, instead learns the shared structure of movement across eleven datasets and twenty-nine cities (rail, buses, taxis, bikes, freeways, high-speed rail), then fine-tunes to a specific system (Zhang et al. 2026). The language model is not writing text; it serves as a large pattern-learning engine for structured time series. The result points at a shift from bespoke models of single systems toward reusable models that transfer across cities, modes, and scales, on the wager that systems which look different on the surface often rhyme underneath.

Several large-scale digital-twin and planetary-data efforts are assembling adjacent pieces: Microsoft Planetary Computer, NASA's digital twins program, ESA/Copernicus, the WEF/DNV Global Digital Twin ambition.

A clarification, because it causes real confusion. The simulator behind this book is, today, mechanistic: brute-force numerical integration of many interacting agents over time. It is not a neural network, and it is not "LLM magic," though non-technical audiences often assume it must be. The choice is pragmatic, not ideological. Mechanistic integration is simple to reason about, its assumptions are inspectable, and inspectability is the whole point of a civic model. But these are techniques in service of an outcome, not articles of faith. As learned and physics-grounded world models mature, the natural path is hybrid: mechanistic where the structure is known and has to be argued about, learned where pattern beats first principles, with the boundary between the two made explicit rather than hidden. The question is not mechanistic versus learned; it is how to keep a model legible as more of it becomes learned.

This is still early. But the direction is clear, and these strands and the civic-publication work in this book are climbing the same mountain from different sides.

Where the existing tools stop

Every paradigm above produces outputs that are locked inside a desktop application, a Jupyter notebook, or a PDF figure. The closest things to public, interactive publication (NetLogo Web, some Stella online exports) are retrofits, not first-class targets: running a Vensim model requires Vensim, a GAMA model requires Java, a Mesa model requires Python. The audience that can do any of these is tiny relative to the audience with a stake in the questions being modeled.

This is the gap the preface named, and the reason it is worth stating where the field stands: the missing piece is not a better model engine — Mesa, Agents.jl, and NetLogo already exist — but a way to publish a model as an interactive argument. Chapter 8 takes that problem up directly. The practical posture it implies for the rest of this part: be willing to call a model "simple and illustrative" without apology, wrap existing tools where they are better, and treat calibration and data provenance as explicit variables rather than hiding them.

The specific tools attached to each paradigm above, the institutions advancing the field, and the people whose names recur throughout this book are catalogued in Reference — Tools, Institutions, and People.

Further reading

- Batty, M. (2013). The New Science of Cities. MIT Press.
- Colander, D., & Kupers, R. (2014). Complexity and the Art of Public Policy. Princeton University Press.
- Beinhocker, E. (2006). The Origin of Wealth. Harvard Business School Press.
- Li, F.-F. (2023). The Worlds I See. Flatiron Books. (On spatial intelligence and the trajectory toward world models.)

Open questions

- What is the right architectural model for a civic simulation layer that wraps existing runtimes?
- Is there a viable path from the current NetLogo/Mesa paradigm to browser-native publication without rebuilding everything?
- How do we calibrate civic models without the data infrastructure that national labs have?

TODO / White-space

- FIGURE: Modeling paradigm comparison table (SD vs ABM vs I-O vs IAM vs ESM — strengths, weaknesses, use cases).

- FIGURE: Computational hierarchy: ESM → IAM → reduced-complexity → browser-tier.
- FIGURE: World3 model structure diagram (from Limits to Growth).
- TODO: Add references to CMIP6 and the SSP/RCP scenario vocabulary.
- TODO: Expand on the Macrocosm / Farmer group work as the closest thing to the book's thesis at the research level.
- PLACEHOLDER: Interview with simulation practitioner (Oxford, SFI, or PIK) on what makes a model useful versus misleading.

Chapter 7

Data, Uncertainty, and the Question of Calibration

> The hardest part of any socio-environmental model is data, not the engine. This chapter addresses the deeper methodological question behind every model: how do you know whether it is right? We introduce calibration, ensemble methods, and the communication of uncertainty, and argue for a specific approach to honest modeling that distinguishes illustrative from predictive claims.

Our intent here is to be honest about the part that is actually hard — not the engine but the data and the calibration — and to make the case that legible uncertainty is more useful, and more truthful, than false precision.

Data is the constraint

A beautiful model runner means nothing without data that is accurate, current, and at the right spatial and temporal granularity. EXIOBASE, GTAP, UN Comtrade, IEA, FAOSTAT —

these took decades and tens of millions of dollars in funding to build (Franks et al. 2010). A new platform inherits none of that. The first version of any model built on a new platform will be demonstrably wrong, and the work to make it less wrong is data work, not platform work.

This is one of the strongest objections to building new civic simulation infrastructure, and it deserves a direct answer.

The answer is not that the data problem has been solved. It has not. The answer is that there is a spectrum of model uses that have different data requirements:

1. Predictive quantitative models (what will happen?) require calibrated parameters, uncertainty quantification, and validation against historical data. Data requirements are high.
2. Illustrative structural models (what is the shape of the dynamics?) require the structural relationships to be correct — the sign of feedbacks, the qualitative behavior near thresholds — but not the precise values. Data requirements are much lower.
3. Counterfactual argument models (what would happen if…?) require that the model can distinguish between scenarios — that the model is sensitive to the thing being varied — but need not be precisely correct about absolute magnitudes.

For civic purposes, categories 2 and 3 are often sufficient. A model that shows that groundwater depletion is self-accelerating (because falling water tables require longer pumps, which require more energy, which raises costs, which pressures farmers to pump more water to remain profitable) does not require precise aquifer geometry to be useful. A model that shows that eliminating a traffic lane increases pedestrian traffic that generates more retail revenue than the parking spaces it replaces does not need to know the exact demand elasticity to shift the burden of proof.

The honest move is to label what kind of model you are presenting — and to make the structural assumptions visible, so that people can challenge them rather than just disputing the output.

The data landscape, in brief

The raw material for civic simulation now exists, is public, and is largely free: trade and economic flows (UN Comtrade, OEC, EXIOBASE, GTAP), energy (IEA, EIA), food and agriculture (FAOSTAT), minerals (USGS), ocean and bathymetry (GEBCO, Global Fishing Watch), geography (OpenStreetMap, Natural Earth), climate (ERA5/Copernicus, NASA Earthdata), and population (WorldPop, SEDAC). A decade ago much of this did not exist or was not public.

The full catalogue — source, URL, access model, and the chapters that use each — lives in Appendix B: Data Sources and Code Links, which is the single place that list is maintained. What matters for the argument of this chapter is not the inventory but the consequence: the data exists, and yet building a correct model from it is still hard. That difficulty is the subject of the rest of the chapter.

Calibration: the genuinely hard problem

Calibration is the process of adjusting a model's parameters so that its outputs match observed historical data. It sounds simple. It is not.

A model of a coupled socio-environmental system may have dozens or hundreds of parameters. Many of them are not directly measurable — they represent aggregate behavioral tendencies, feedback strengths, or threshold values that can only be inferred from the behavior of the whole system. The parameter space is vast; many combinations of parameters will fit the historical data equally well (the problem of equifinality); and a model that fits history well may still project the future badly if it has overfit to specific historical features that will not repeat.

The Farmer group at Oxford has spent years on calibration methodology for agent-based macroeconomic models: - Surrogate/metamodel calibration: fit a fast, approximate model to the full ABM and use the surrogate for parameter search. - Black-box Bayesian inference: treat the model as a black box and use Bayesian methods to build a posterior distribution over parameters consistent with the data (Kennedy and O'Hagan 2001). - Differentiable ABMs (AgentTorch): make the agent model differentiable so that gradient-based optimization can be used for calibration.

For civic models, the honest synthesis is:

> Producing authoritative-looking outputs from uncalibrated models is arguably harmful (Saltelli et al. 2020). The appropriate response is not to avoid building models but to be explicit about what has been calibrated and what has not — and to frame outputs as scenarios or arguments rather than predictions when calibration is absent.

Uncertainty: ensembles and ranges

A single model run is not a forecast. Any projection is conditional on assumptions about initial conditions, parameter values, forcing inputs, and structural choices about how the model is built. The range of uncertainty across these assumptions is often larger than the central estimate.

The appropriate response is ensemble methods: running the model many times with systematically varied parameters (or varied forcing inputs, or varied structural assumptions) and presenting the distribution of outputs.

The IPCC uses this approach for its headline numbers (2021): the "likely range" and "very likely range" of temperature projections are not error bars on a single model; they represent the spread across the CMIP6 ensemble of different models from different institutions, each calibrated differently.

For civic models, the practical implication is: the question is not "what will happen?" but "what is the range of plausible outcomes, and how does that range depend on what we do?"

This framing is actually more useful than a point prediction for civic decision-making. A decision-maker who knows that the aquifer might be empty in 15–25 years depending on pumping rates (and here is what different pumping rates look like) has more useful information than one who is told “the aquifer will be empty in 19 years.”

> Quick intuition — communicating uncertainty
> A weather forecast does not say “it will rain on Thursday.” It says “70% chance of rain on Thursday.” That framing: - conveys that the future is genuinely uncertain - quantifies the confidence in the forecast - remains useful for decision-making (bring an umbrella if the probability is high enough to matter)
> Civic simulation should aspire to the same structure. The goal is not false precision but legible uncertainty.

Illustrative versus predictive: the honest taxonomy

Model type	Data requirement	Calibration required	Appropriate framing
Structural / illustrative	Low	No	“This is the shape of the dynamics”

Model type	Data requirement	Calibration required	Appropriate framing
Scenario / counterfactual	Medium	Partial	"Here is how the outcome changes if you change X"
Predictive / quantitative	High	Yes	"This is what we expect to happen, with these uncertainty bounds"
Policy optimization	Very high	Yes + validation	"This is the policy that best achieves these outcomes"

For civic publication, structural and scenario models are often the right tools. They are honest about their limits, useful for decision support, and accessible to non-specialists. The goal is not to replace the government agency's technical model — it is to give communities the ability to interrogate the structural logic of that model before accepting or rejecting its conclusions.

A final note on the word "available." This chapter treats data availability as a technical constraint, but availability is also a claim about rights. Some data about places belongs to peoples — Indigenous data sovereignty extends to data about lands, waters, and non-humans, and includes the right to withhold. That argument is made once, in Chapter 9 ("The right not to be modeled"); every data-sourcing decision in this chapter is subject to it.

Further reading

- Taleb, N. N. (2012). Antifragile. Random House. (On tail risks and robustness under uncertainty.)
- Saltelli, A. (2022). The Politics of Modelling. Oxford University Press.
- Oreskes, N., & Conway, E. M. (2010). Merchants of Doubt. Bloomsbury Press.

Open questions

- What is the minimum viable calibration for a civic simulation to be "honest enough"?
- How do we communicate probabilistic outputs to non-specialist audiences without collapsing to a single number?
- Is there a standardized "data provenance" format for civic models that would allow independent verification?

TODO / White-space

- FIGURE: Data source map — which data covers which domain (trade, climate, biology, population).
- FIGURE: Uncertainty fan chart — how ensemble spread grows over time.
- FIGURE: Calibration process diagram: observed data → model → parameter estimation → validation.
- TODO: Add specific examples of calibrated vs. uncalibrated models and the different conclusions they support.
- TODO: Add section on eDNA and emerging observational datasets.
- PLACEHOLDER: Interview with data scientist at national lab or climate service on data provenance challenges.

Chapter 8

From Physics to Practice

> This chapter bridges theory and practice. We walk from the top of the computational hierarchy (Earth System Models, IAMs) down to the tier that runs on commodity hardware, describing the architectural patterns, computational building blocks, and specific recipes that make the lower tiers tractable. We then describe the structural forces to model — from abiotic geography through climate to biosphere to the humanosphere — and the scoring problem: how do you evaluate a scenario's outcomes?

Our intent here is to show that you do not need a supercomputer to start — that a surprising amount can be computed from geometry and physics alone on ordinary hardware — while being clear about where that stops and the heavy models take over.

The computational hierarchy

Socio-environmental simulation occupies a wide computational range. At one end: Earth System Models requiring months on a supercomputer cluster to simulate a century of climate. At the other: a simple system dynamics model that runs on consumer hardware — a laptop, a phone, a browser tab — in milliseconds.

The hierarchy, from most to least computationally expensive:

1. Earth System Models (GCMs/ESMs) — 3D fluid dynamics on a rotating sphere; physically grounded climate physics; CMIP6 ensemble. Not browser-runnable; not even close.
2. Integrated Assessment Models (IAMs) — couple climate physics to economics; produce SSP/RCP scenario trajectories; operate at decadal timescales, global to regional resolution. Run on workstations in hours to days. Expert-operated; not civic tools in their current form.
3. Reduced-complexity climate models (FaIR, MAGICC, Hector) — emulate the IAM/ESM ensemble with a handful of ODEs (C. J. Smith et al. 2018). Run in seconds on a laptop. The IPCC uses these to produce probabilistic temperature ranges. This is the commodity-hardware tier of climate physics.
4. Domain-specific models (wildfire, hydrology, crop yield) — purpose-built for specific physical processes. Many are calibrated to specific regions and validated against historical events. Medium data requirements; often runnable in minutes on a laptop.
5. System dynamics models — stock-and-flow models of feedback dynamics; analytically tractable; run quickly. Best for qualitative dynamics and policy analysis rather than precise spatial forecasting.

6. Agent-based models — heterogeneous actors with behavioral rules; emergent dynamics. Computational cost scales with number of agents; can be made browser-runnable for small populations.
7. Input-output / matrix models — linear algebra; essentially instantaneous. Good for supply chain shocks, ecological footprints, economic sector impacts.

The civic simulation tier lives primarily at levels 3–7. The strategy: use the richest model your data and compute budget support, clearly label its tier and limits, and connect it to higher-tier outputs where available.

Physical building blocks: what the Earth gives you for free

Several quantities can be computed from geometry and physics alone, without storing large datasets. These are the foundation on which simple civic models can be built quickly.

Insolation and irradiance

Given latitude, longitude, elevation, and time, the direct and diffuse solar irradiance at a point can be computed analytically. The solar zenith angle θ_z as a function of latitude ϕ, solar declination δ, and hour angle H:

$$\cos(\theta_z) = \sin(\phi)\sin(\delta) + \cos(\phi)\cos(\delta)\cos(H)$$

The solar declination varies over the year:

$$\delta = -23.45^\circ \cdot \cos\left(\frac{360}{365}(d+10)\right)$$

where d is the day of year.

Without clouds, the direct normal irradiance at the surface is:

$$I_{DNI} = I_0 \cdot e^{-\tau/\cos(\theta_z)}$$

where $I_0 \approx 1361\ \mathrm{W/m^2}$ (solar constant) and τ is the clear-sky optical depth.

This gives you: solar resource potential at any location and time, shadow maps from terrain, the diurnal and seasonal temperature cycle as a first approximation. The expensive part is cloud shadowing, which requires a volumetric atmospheric layer — without clouds, the computation is fast.

DEM-derived hydrology

A Digital Elevation Model (DEM) is a raster of elevation values. From this, purely locally:

- Slope and aspect — first derivatives of the DEM surface. Slope drives surface runoff; aspect drives differential insolation and soil moisture.
- Flow direction — the direction of steepest descent from each cell (D8 algorithm: one of eight cardinal/diagonal directions). Defines drainage structure (Wang and Liu 2006).
- Flow accumulation — how many upstream cells drain

through each point. Where accumulation is high, streams form.

- Topographic wetness index (TWI): $TWI = \ln(a/\tan(\beta))$ where a is the upslope contributing area and β is the local slope. High TWI → water pools; low TWI → water runs off quickly.
- Viewshed — which cells are visible from a given point (relevant for fire spotting, wind exposure, solar shading).

USGS provides 1/3 arc-second (approximately 10m) DEMs for the US; SRTM provides global coverage at 30m; Copernicus DEM at 30m globally; GEBCO for ocean bathymetry.

First-order temperature and ice

A first-order global surface temperature can be estimated from latitude, elevation, land/ocean fraction, and season, without a climate model:

$$T \approx T_{\text{equator}} - \Delta T_{\text{lat}} \cdot \sin^2(\phi) - \Gamma \cdot z$$

where $\Gamma \approx 6.5$ °C/km is the environmental lapse rate and z is elevation.

This approximation is not accurate enough for policy work but is sufficient for planetary-scale illustrative scenarios: where snow persists year-round, rough ice extent, which regions approach critical temperature thresholds.

The forces to model: a taxonomy

Any serious civic model eventually needs to represent the following layers, at varying resolution:

Geography — abiotic

- Digital elevation models (DEMs)
- Soil type, soil carbon, porosity
- Rock type and aquifer structure

Climate

- Solar radiation: insolation, irradiance (mostly analytic)
- Orbital parameters: eccentricity, axial tilt, precession (for long-range scenarios)
- Radiative transfer: atmospheric gas absorption (CO_2, H_2O, O_3), albedo
- Atmospheric thermodynamics: convection, advection
- Fluid dynamics: Coriolis, Hadley/Ferrel/Polar cells
- Pressure systems and ocean-atmosphere coupling
- Cryosphere: sea ice, glaciers, permafrost
- Chemistry: CO_2, ozone, aerosols

Hydrology

- Ocean, lakes, soil moisture, plant transpiration
- Condensation and precipitation
- Runoff, rivers, streams (DEM-derived)

- Aquifers, groundwater, porosity
- Turbidity and water quality
- Tidal forcing

Biosphere

- Microbial populations and soil biology
- Plant communities (forests, grasslands, wetlands)
- Animal populations (megafauna, fish, invertebrates)
- Species interactions (predation, competition, mutualism)

Humanosphere

- Human populations and demographics
- Land use: roads, dams, fences, housing, zoning
- Industrial infrastructure: oil refineries, power plants, factories
- Energy consumption and waste products
- Food consumption and agriculture
- Farming practices: soil exhaustion, fertilizer runoff, irrigation
- Economic dynamics: water rights, commodity prices, labor
- Disease dynamics
- Governance and policy

Not all layers are needed for all scenarios. A wildfire model needs terrain, climate (wind, temperature, humidity), and fuel load (vegetation). It does not need economic dynamics. A water rights conflict model needs hydrology, agriculture, and governance. It does not need microbiology. The art is identifying which layers matter for the dynamics being modeled.

The coupling problem

Multi-domain models require connecting models from different layers. A climate model drives a crop model; a crop failure drives a migration model; migration changes labor markets which affects land use which feeds back to climate.

The coupling problem has several aspects:

Spatial resolution mismatch. A climate model might operate at 25km resolution; a land use model at 100m resolution; an economic model with no spatial component at all. Downscaling and upscaling across resolutions is non-trivial and can introduce artifacts.

Temporal resolution mismatch. Climate models output daily or monthly averages; economic models might operate at quarterly or annual timescales. Synchronizing these requires decisions about how to aggregate and interpolate.

Feedback directionality. Some couplings are one-way (climate drives crop yield, not vice versa at short timescales). Others are two-way (land use affects climate, climate affects land use). The directionality must be explicit in the model architecture.

State initialization. Each component model has a state that must be initialized consistently with the others. If the climate model starts in one state and the economic model in another, the system may exhibit spurious transients at the start of the simulation.

For civic models, the practical advice: start with one-way coupling and explicit inputs/outputs. Define what each model component takes in and puts out. Only add two-way coupling when the feedback is central to the question being asked.

A useful special case of one-way coupling is offline coupling: run one model to completion, save its output as a dataset, and feed that saved run into another model as a fixed input rather than computing both live. A food-pressure model can read a pre-baked run of world demographics and migration instead of simulating population dynamics itself. In reality everything affects everything, so a fully live, two-way coupling would be more faithful. But a cached, well-validated run of the domain you are not focused on is often good enough, and far cheaper, than re-deriving it badly inside every model that needs it. It also lets a library of pre-computed runs (a demographic trajectory, a climate pathway, a trade-shock cascade) become shared inputs that many downstream scenarios draw on. The cost is staleness: the cached run cannot respond to the model consuming it, so offline coupling is the right move precisely when that back-reaction is weak, or slow on the timescale of the question being asked.

The EMA Workbench and deep uncertainty

The EMA Workbench (Python) (2017) is worth knowing as an architectural pattern for civic models dealing with deep uncertainty — situations where the model structure itself is uncertain, not just the parameters.

The workbench supports: - Scenario discovery — find the combinations of uncertain parameters that produce different types of outcomes. - Robust decision-making (RDM) — evaluate policies not by expected value but by performance across a wide range of scenarios. - Many-objective evolutionary algorithms (MOEA) — find Pareto-optimal policies across multiple competing objectives.

For a civic water model, this might mean: rather than asking "what is the best policy?" (which requires knowing what "best" means and how certain you are about the model), you ask "which policies perform acceptably across all the scenarios we are uncertain about?"

Event injection: driving scenarios from the real record

A scenario does not have to be hand-authored. A second pattern is to drive a simulation from a stream of real events: feed a curated feed of dated incidents (a blockade, a coup, a port closure, a flood) into the model as parameter shocks at the moments they actually occurred. Public event databases make this tractable. ACLED catalogues political violence and protest with dates and locations; GDELT extracts events from global news at scale. Pulled in as shocks, they let a model run against the historical record and produce a what-actually-happened baseline, which is the thing a counterfactual ("what if the response had been different?") gets measured against.

Two cautions. The feeds are noisy and biased by what gets reported, so a layer of classification and filtering sits between the raw stream and the model, and that layer is itself a modeling choice that should be made visible rather than buried. And event-driven runs answer a different question from parameter-swept runs: one asks "what did this particular history do," the other asks "what is the range of things that could happen." A civic model often wants both, kept clearly distinct so a reader knows which one they are looking at.

The scoring problem

How do you evaluate a scenario's outcome? What makes a future "good" or "bad"?

This is a values question, not a technical one — but models must make it explicit rather than hiding it in the objective function. The scoring problem has several dimensions:

Biodiversity. Measures of species richness, functional diversity, population stability. The Living Planet Index (WWF) tracks vertebrate populations; the IUCN Red List tracks extinction risk.

Both provide empirical benchmarks for "how is the biosphere doing?"

Human welfare. Food security, water access, health outcomes, economic stability. Often in tension with ecological measures in the short run.

Resilience. Not just current state but the capacity to recover from perturbation. A simplified measure: recovery time after a shock, or the size of the stable attractor.

Equity. Who bears the costs and who receives the benefits of a given scenario? A scenario that improves average welfare while concentrating costs on the most vulnerable is different from one that distributes improvements broadly.

Temporal fairness. Does the scenario trade future welfare for present gain? Discount rates in economic models implicitly answer this question; in civic models it should be explicit.

A civic model should present these dimensions separately rather than collapsing them into a single score. The weighting is a political decision, not a technical one; the model's job is to make the tradeoffs visible, not to resolve them.

A minimal commodity-hardware stack

For a first civic simulation, the following stack is tractable:

Climate inputs: FaIR (reduced-complexity climate model) or ERA5 climatology (historical) for temperature and precipitation.

Terrain: SRTM or Copernicus DEM for elevation; derived slope, flow direction, and flow accumulation.

Domain model: System dynamics (implemented in Python or JavaScript, running in-browser via Pyodide or compiled to WASM) or a simple ABM (Mesa in Python, compiled or server-hosted).

Publication: HTML + JavaScript, connecting to a server-side simulation via WebSocket (socket.io pattern), or running entirely client-side for simple models.

Visualization: Map layer (Leaflet, MapLibre, or Deck.gl) for spatial outputs; time series charts (Vega-Lite, Observable Plot) for temporal outputs; parameter sliders for scenario exploration.

This stack can produce a working interactive scenario in days to weeks, depending on the complexity of the domain model and the availability of the underlying data.

Further reading

- Batty, M. (2013). The New Science of Cities. MIT Press.
- Saltelli, A., et al. (2008). Global Sensitivity Analysis: The Primer. Wiley.

Open questions

- What is the right trade-off between physical realism and computational tractability for civic models?

- How do we build coupling architectures that are both flexible enough to extend and simple enough to understand?
- Is there a minimal viable "scoring framework" that captures the values questions in a civic model without imposing them?

TODO / White-space

- FIGURE: Computational hierarchy diagram: GCM → IAM → RCM → browser-tier.
- FIGURE: DEM derivatives: slope, flow direction, flow accumulation, TWI (with example raster).
- FIGURE: Coupling architecture: one-way and two-way domain coupling diagram.
- CODE: Minimal FaIR example: run a carbon pulse to 2100 in Python, plot temperature response.
- CODE: DEM flow direction algorithm (D8) in pseudocode.
- TODO: Add more on WASM/Pyodide for browser-tier Python.
- PLACEHOLDER: Interview with earth scientist on what physical primitives are well-characterized and what is still poorly understood.

Chapter 9

The Publication Layer

> The missing piece in civic simulation is not computation — it is publication. This chapter takes the gap the preface named and asks the constructive question: what would it take to publish a simulation as an interactive argument that a non-specialist can run, inspect, and challenge? It has two halves — what the format is, and who it has to reach — because a format nobody can open is not a publication.

Our intent here is to build the book's central claim in full: to describe the artifact (an explorable, inspectable, runnable model) and the reach (the audiences it must actually serve), and to argue that AI-assisted authorship is what makes both newly possible.

Taking the gap as given

The preface named the gap: there is no equivalent of a scientific paper or an op-ed for a simulation. Every significant tool in this field — Vensim, Stella, NetLogo, Mesa, AnyLogic — produces outputs that are locked inside a desktop application, a Jupyter notebook, or a PDF figure. The model lives on one person's computer. The outputs are static: a chart exported to a slide deck, a table in an appendix. The person who wants to understand the model's assumptions cannot change them and watch what happens; they must accept or reject the conclusions as presented.

This is not a minor inconvenience. It is a structural barrier to civic use. For a model to function as a civic argument — for it to change the quality of a planning meeting, a city council hearing, a school board deliberation — the audience needs to be able to interact with it. They need to be able to say "but what if we increase the pumping rate?" and watch the aquifer level respond. They need to see the difference between the baseline and the proposed policy. They need to be able to challenge the assumptions by changing them.

Closing this gap is a harder problem than building a better model engine. It has two halves. The first is the artifact: what a published simulation is, technically and in its intellectual lineage. The second is reach: who has to be able to open it, and what that demands of the design. They are the same problem from two sides.

Part one — what the format is

The intellectual lineage

Bret Victor and explorable explanations

The clearest statement of what the publication layer should be comes from Bret Victor's 2011 essay "Explorable Explanations" (2011). Victor argues that documents that make claims about dynamic systems — policy analysis, scientific arguments, engineering proposals — should embed the model they are making claims about, so that the reader can vary the assumptions and verify the conclusions interactively.

He calls these explorable explanations, and contrasts them with passive text that describes the behavior of a system without giving the reader the ability to explore it directly. A sentence like "if the pumping rate increases by 10%, the aquifer will be depleted in 17 instead of 22 years" is much weaker than a slider that lets the reader set the pumping rate and watch the depletion timeline update.

Victor's Dynamicland project extends this to physical space: a computing environment where any printed document can be an interactive object, manipulated by hand, with results projected back onto the paper in real time. The goal is to make computation as immediate and physical as drawing (Victor 2013).

The publication layer we are describing here is in the same intellectual tradition: a format where the model is the argument, where the reader is the experimenter, and where the gap between claim and evidence is made interactive.

Jupyter and the notebook model

Jupyter's model is a persistent server-side Python kernel connected to a browser client. The notebook is a JSON document containing a sequence of cells — either markdown or code. Code cells execute in the kernel, and their output (text, charts, widgets) is rendered inline. State accumulates in the kernel across cell runs. Google Colab is Jupyter hosted by Google, with a managed kernel and Drive-backed storage.

The constraint: cells execute sequentially, which fits exploratory data analysis but is awkward for a continuously-running simulation. The output model is oriented around snapshots, not streams.

Observable

Observable (observablehq.com) is the closer analog. Notebooks are JavaScript-first. Cells are reactive — if cell B depends on cell A, it re-evaluates automatically when A changes. You can embed canvas, D3, Three.js, and other browser APIs directly. There is no persistent server kernel; compute is in the browser.

Observable's reactive model is well-suited to parameter sliders and interactive visualizations (Bostock 2018). The tradeoff: it uses a proprietary reactive JS dialect that is different from standard JavaScript, and the format has quirks that make AI authorship harder.

Glitch and the frictionless deploy

Glitch (glitch.com) is worth naming as a spiritual predecessor. It offered instant Node.js hosting with a browser-based editor — you forked a project, edited files live, and your app was running at a public URL within seconds. No build step, no deploy pipeline, no GitHub required.

It was popular for small experiments precisely because the friction between “idea” and “running thing someone else can see” was near zero. The lesson: deployment friction kills experimentation. The platforms that replaced it (Render, Vercel, Netlify) are more powerful but all require a git push and a build step — enough friction to break the flow of a quick idea. For civic simulation, the equivalent of Glitch’s frictionless deploy is running a local server and sharing it over a tunnel (ngrok, Cloudflare Tunnel) — a public URL in ten seconds.

The architecture of a civic simulation platform

Based on the intellectual lineage above, the right architecture for civic simulation is:

The scenario is an ordinary HTML document. No special notebook format, no proprietary cell system. The author (human or AI) writes HTML with standard JavaScript. The document can include a small client library that wraps the simulation API. This degrades gracefully: the page can be a static document if the simulation server is offline.

The server is a thin bridge. The simulation logic lives in domain-specific modules (agents, manifests, data loaders). The server starts simulation instances on demand and streams tick events to browsers over WebSocket (socket.io). The server does not know about insolation, hydrology, or any domain concept — it just starts processes and forwards events.

Multiple clients share one simulation instance. A multiplayer model (a city council viewing the same scenario) is the same as a solo model, just with multiple WebSocket subscribers. The simulation does not know or care how many clients are watching.

Plain HTML is the right format for AI authorship. An AI can write, edit, and debug it. It composes with any browser API. The tradeoff vs. Jupyter/Observable is that there is no built-in reactivity or cell runner — the author writes the wiring. For a tool where AI is doing much of the document authoring, this is acceptable. (We return to AI authorship below.)

Interface requirements for civic simulation

A civic simulation interface needs at minimum:

A map layer. Most scenarios have spatial extent. The map provides orientation and makes the scenario legible to non-specialists who know the geography. Leaflet, MapLibre, or Deck.gl are the right tools; they are open-source, well-documented, and composable with JavaScript.

A time control. Most scenarios evolve over time. Users should be able to play, pause, step, and scrub through the simulation timeline. The time control also exposes the scenario's temporal scope — is this days, years, decades?

Parameter sliders. The key interactive element: let the user vary the scenario inputs and watch the outputs respond. The set of parameters exposed should match the question being asked (what if the pumping rate doubles? what if rainfall decreases by 15%?). Too many parameters is as bad as too few — it shifts the burden to the user to figure out which ones matter.

A narrative layer. The simulation does not speak for itself. The document needs text that explains what is being modeled, what the assumptions are, what the outputs mean, and why it matters. The narrative layer is the argument; the simulation is the evidence.

Uncertainty visualization. As discussed in Chapter 6, a single model run is not a forecast. The interface should show ranges or distributions of outcomes, not just a single trajectory.

Provenance display. What data was used? Who built the model? What are the key structural assumptions? This should be visible without requiring the user to read the source code.

Phase-transition display. This is the requirement non-specialists feel most and interfaces serve worst. A simulation computes smoothly: a reservoir slides toward empty, a population curve bends down by degrees. But people do not experience a system as a smooth curve; they experience thresholds, the points where a situation changes in kind rather than in degree. Water at 40 percent and water at zero are not two points on one line; one is "managed" and the other is "crisis," and the passage between them is what people actually care about. A civic interface should mark those qualitative state changes on the model itself: where a curve crosses a threshold, where a feedback flips sign, where the system leaves one regime and enters another. Naming the inflection points, rather than leaving them implicit in a continuous trace, is how a model meets human perception halfway. It is also the visual form of the book's argument that systems behave nonlinearly (Chapter 3): the threshold is not a presentation flourish, it is the part of the dynamics that matters most. If Chapter 3 is right that we are constitutionally bad at feeling nonlinearity, this is one of the concrete places the tool earns its keep. It shows the inflection a person could not compute on their own, and gives a group a shared object to point at and argue over, which is the small, specific way these tools help us reason past the limits of unaided intuition.

Cascade display. When a model couples several domains, the most useful thing to show is often not the final number but the chain that produced it: a strait closes, oil prices rise, fertilizer costs follow, a harvest falls, a food index crosses a line, and somewhere a community comes under stress. Showing that causal chain explicitly, rather than only the aggregate output, lets a viewer see where in the chain an assumption or an intervention bites. It also makes the model arguable at each link instead of only at the conclusion, which is the whole point of publishing it.

The widget library problem

A curated set of high-quality visualization components — heatmap canvas, time scrubber, cell table, scenario parameter panel — that a scenario page can include with a single custom element tag (`<orbital-heatmap>`, `<orbital-timeseries>`) dramatically lowers the authoring cost. Rather than writing a canvas renderer from scratch, the author writes one line. The bar for an AI-authored page drops considerably. This is part of what made Observable and Deepnote productive: a small set of polished built-in visualizations meant that most scenarios could be presented without custom code.

The root gallery: discoverability and browsing

A public civic simulation platform needs a root experience — something like shadertoy.org or Glitch's project gallery — that lets visitors browse public scenarios, understand what each does, and fork one as a starting point.

The minimal version: an index page listing all public namespaces and projects, with brief descriptions and (eventually) live thumbnails. The full version: search, filtering by domain (water, energy, urban, food), tags, a "recently updated" feed.

This is where the social dimension of civic simulation lives: you are not just building a model for yourself, you are contributing to a shared library that others can inspect, challenge, extend, and run.

Part two — who it must reach

Accessibility as a systemic property

The artifact above is only half the problem. The other half is reach: a published simulation is accessible only if the whole chain — from raw data to model to published interactive argument — is traversable by the people who have a stake in the question, not just by the people who built the model.

By this standard, every simulation tool in common use today is inaccessible — not only because it is desktop-bound, but because the people most affected by the decisions that simulations inform typically cannot access, run, modify, or challenge the models that produce the technical inputs to those decisions. They receive conclusions, not models. They must trust or contest expert judgment with no tools for the latter. This is not primarily a technical problem; it is a design problem (Norman 1988). The tools exist to make simulation accessible. The missing piece is the will and the resources to build the translation layer.

The three audiences

Civic simulation has at least three distinct audiences, each with different needs:

The general public — residents, community members, engaged citizens. They do not read Python. They may not understand what a feedback loop is. They need: a map they recognize, a question framed in terms they care about, sliders with plain-language labels, and outputs that mean something to them without technical translation. They also need the model to respond quickly — a ten-second delay between parameter change and output is too long for engagement.

Engaged practitioners and advocates — planners, environmental NGOs, city council members, journalists, teachers. They can follow a technical argument; they may be able to read a chart or interpret a statistical result. They need: the assumptions to be visible, the uncertainty to be communicated, the ability to drill down into the structure if they want to, and a publication format they can share. They are the primary audience for the participation tools in Chapter 9.

Technical specialists — model builders, data scientists, domain researchers. They need: access to the underlying model code, the data sources and their provenance, the calibration methodology, and the ability to fork and modify. This audience is the bridge between the civic tier and the research tier.

These audiences are not served by the same interface. The right approach is layered disclosure: a document that presents a simple, legible surface to the general public, allows practitioners to drill into the assumptions and uncertainty, and exposes the full technical stack to specialists.

Design principles for accessible civic simulation

Simple at the surface. The first thing a community member sees should be their territory, a question they care about, and a clear invitation to explore. Not a diagram of the model structure, not a list of data sources, not a warning about the model's limitations. Limitations, data sources, and model structure are important — they belong in the document — but they should be discoverable, not mandatory. The Distill journal's "hover for more" interaction, Observable's expandable sections, or a simple "Show assumptions" button are the right pattern.

Legible uncertainty. As discussed in Chapter 6, uncertainty is a result, not a failure. An accessible civic model presents uncertainty in terms a non-specialist can act on: not "the 95% credible interval is 17–23 years" but "under current pumping rates, the aquifer is likely to reach critical levels somewhere between 17 and 23 years from now — here is what different pumping rates do to that range." Fan charts, cone-of-uncertainty visualizations, and scenario comparison views are the standard tools. The key principle: the uncertainty visualization should never collapse to a single number, because that implies false precision.

Fail-graceful. A civic simulation document should work partially even when the simulation server is offline. The narrative text, the background maps, and the historical data visualizations should all be static and load from CDN. The live simulation layer is an enhancement, not a dependency. This matters for communities with unreliable internet, and for archiving — a document that depends entirely on a running server is not a document, it is a service, and services go down.

Provenance visible. Every number in a civic simulation document should be traceable to a source. The provenance layer —

which data was used, who processed it, when, with what methodology — should be accessible without requiring the reader to read the source code. A simple "data sources" panel, visible on request, is the minimum. Provenance also carries permissions: where data concerns Indigenous lands or knowledge, community-authored terms (such as Local Contexts TK and Biocultural Labels) belong in this layer, and the publication system must support withholding as well as disclosing — the full argument is in Chapter 9, "The right not to be modeled."

Plain language. Technical terms (aquifer recharge rate, trophic cascade, quadratic voting) should be defined on first use. Glossaries, hover-over tooltips, or linked sidebars are the right mechanism. The goal is not to avoid complexity but to ensure that complexity is opt-in rather than required for basic comprehension.

The publication layer as civic infrastructure

The most important structural intervention for accessible civic simulation is treating the publication layer as public infrastructure, not as a feature of a commercial product.

The analogy is open-source maps. Before OpenStreetMap, good geographic data was expensive and proprietary — only organizations with resources could access it. OpenStreetMap democratized the data layer; the communities that built it (and continue to improve it) include millions of contributors worldwide. The result is that a small community organization in a rural area can produce a map of its territory as good as anything a developer with a GIS team can produce.

An equivalent infrastructure for civic simulation would provide: - Open data pipelines for the major environmental and economic datasets - Browser-runnable model templates for common scenario types (watershed, urban heat, supply chain) - A publication format that is freely hostable on any web server - A gallery for sharing published scenarios with community review - A review and provenance system that distinguishes models with different levels of validation

This does not require a centralized platform. It requires open standards and shared tools — the same model that made OpenStreetMap, Wikipedia, and open-source software succeed.

Civic models as educational tools

A specific and high-leverage application of accessible civic simulation is education. Students who interact with a model of their watershed — who can change the pumping rate and watch the aquifer deplete, who can add a wildfire and watch the debris flow propagate — develop intuitions about coupled systems that no amount of textbook description can produce.

The Balinese water temple example (Chapter 4) is precisely this: a simulation that made visible the emergent optimization in a social-ecological system, providing the conceptual tool that allowed outsiders to understand what the community had always known.

Educational civic simulations have a different set of design requirements from advocacy simulations: - Explorability over accuracy. For education, the goal is to develop intuition, not to forecast. The model can be simplified in ways that would be inappropriate for a policy argument. - Multiple entry points. Students at different levels of preparation need different starting points. The layered disclosure approach works well here. - Play before understanding. Let students manipulate the model before they read the explanation. The surprising result (turn off the wolves and the river gets shallower) motivates the explanation, not the other way around.

AI as the primary author

The deeper implication of everything above is that the primary author of civic simulation documents is increasingly not a human programmer but an AI system directed by a human with domain knowledge. One distinction matters before we go further: this is AI as the author of the model, writing and wiring the document. It is a separate role from AI as the engine inside the model, the learned and physics-grounded world models discussed in Chapter 5. The two can combine, but they are different jobs, and conflating them is the source of much of the confusion about what "AI simulation" even means.

A human city planner should be able to say "I want a model of the groundwater depletion in our district under three different pumping-rate scenarios" and receive a working interactive scenario document, without knowing JavaScript or Python or any particular simulation framework. This has already become partially true. AI systems can write, debug, and explain simulation code. The barrier between "I can describe what I want to model" and "I have a running model" is a conversation, not a programming course.

This is the single most significant change in accessibility — not a new interface pattern, but the removal of the programming barrier as a constraint on who can build civic simulations. A community organization that cannot afford a data scientist can use an AI to build a simplified model of its watershed. A teacher who wants to demonstrate a trophic cascade can generate a browser-runnable model without knowing JavaScript.

It is not a substitute for expert review. AI-generated models still require validation, and the provenance question is if anything more important when the author is an AI — the tendency of language models to produce plausible-looking but wrong equations is a real hazard in scientific work. But the direction is unambiguous, and it sets concrete design requirements for the infrastructure:

- File system layout should be flat and explicit. AI tools read and write files; implicit state and convention-based frameworks introduce errors. A directory of named JS files with explicit exports is better than a framework with conventions an AI must infer.

- Manifest formats should be declarative. The scenario configuration should be human-readable JSON or YAML, not a Python class hierarchy.
- The API surface should be minimal. Every additional method or configuration option is a place where AI-generated code can be wrong. Prefer a small, well-documented core.
- The document format should be plain HTML. An AI can write and edit HTML reliably. Proprietary notebook formats, custom cell languages, and compiled frontends add friction that does not benefit AI authorship.

> The human role is curation, judgment, and presentation — deciding which scenarios matter, reviewing what the AI produced, and communicating results to an audience. The editing interface (VS Code or equivalent, plus a chat window) is the workbench. The scenario HTML page is the output.

Stylistic notes for civic simulation interfaces

Science aesthetic, not consumer product. Civic simulations are analytical tools. They should feel like scientific instruments — Tailwind or equivalent utility CSS, system light/dark mode, restrained color palette (Bringhurst 1992). They should not look like apps trying to sell something.

Mobile-first. A city council member reviewing a scenario should be able to do so on a phone during a meeting. The layout should work at small screen sizes.

No login friction for viewing. Public scenarios should be viewable without authentication. Authentication is required only for publishing, which is appropriate.

Half-globe convention. The home page can feature a globe (Three.js or Cesium) placed so that only the left half is visible — a design choice that signals both planetary scope and editorial restraint.

Further reading

- Matuschak, A., & Nielsen, M. (2019). “How can we develop transformative tools for thought?” https://numinous.productions/ttft/
- Hohman, F., et al. (2020). “Communicating with Interactive Articles.” Distill, 5(9).
- Tufte, E. R. (2001). The Visual Display of Quantitative Information. Graphics Press.
- Cairo, A. (2016). The Truthful Art. New Riders.

Open questions

- What is the right balance between narrative structure and open exploration in a civic simulation document?

- How do we handle the "filter bubble" problem — scenarios that show only the evidence supporting one stakeholder's position?
- Is there a standardized interface pattern for "here are the assumptions; challenge one"?
- What is the minimum level of technical literacy required to be a meaningful contributor to a civic simulation?
- Is AI authorship a net positive or net negative for the trust and provenance of civic models?

TODO / White-space

- FIGURE: Architecture diagram — server, WebSocket, client, HTML document, domain model layers.
- FIGURE: Interface wireframe — map, time control, parameter sliders, narrative panel.
- FIGURE: Layered disclosure diagram — surface / practitioner / technical specialist views of the same document.
- FIGURE: Widget library examples — heatmap, time series, parameter panel (placeholder).
- TODO: Add comparison table: Jupyter vs. Observable vs. custom HTML approach.
- TODO: Add specific examples of civic simulation used in educational contexts.
- PLACEHOLDER: Interview with a Bret Victor-tradition researcher or explorable-explanation author.

- PLACEHOLDER: Interview with an educator or community organizer on what "accessible simulation" means in practice.

Chapter 10

Participation, Governance, and Collective Futures

> Simulation is not neutral. A model that gives a community the ability to reason about its water supply, its housing stock, or its fire risk is also a political intervention — it shifts information, redistributes epistemic authority, and changes who can participate in a decision. This chapter covers the traditions of participatory governance that give civic simulation its political home: from cybernetics and anticipatory democracy through computational governance and regenerative finance, to the specific structural failure modes of democratic deliberation that simulation addresses.

Our intent here is to give the publication layer a political home — to show that the modeling and governance traditions are solving adjacent halves of the same problem, and that a model only matters if there is a process ready to use it.

Why collective reasoning fails: seven structural failure modes

Before cataloguing tools and traditions, it is worth being precise about what they are trying to fix. Democratic deliberation is not merely slow or politically contentious — it has structural failure modes that arise from information asymmetries, not just from bad faith or tribalism.

1. Expert–lay asymmetry. One party holds a model (the engineer's pipe-capacity spreadsheet, the seismic hazard map, the grid load forecast) and the other holds lived experience. The technical account is often correct but illegible to non-specialists; the experiential account is legible but partial. The failure is translation, not disagreement. Most infrastructure decisions live in this gap.

2. Distributional engagement asymmetry. Cost and benefit land on different people, and participation tracks cost. A bike lane removes twelve parking spots — those twelve people show up furious — to benefit three thousand cyclists who don't appear because each gains only a little. The process systematically overweights the concentrated loser. The information problem: aggregate benefit is invisible to the individuals who compose it.

3. Temporal asymmetry. Cost now, benefit deferred — or the reverse. Deferred maintenance, pension liabilities, wildfire fuel reduction, transit capital investment. Human cognition and political cycles both discount the future steeply, so the long-range term is structurally invisible at decision time. Counterfactuals

(what didn't burn, what didn't fail) are almost impossible to argue from in a room. A model can make the future term visible and legible.

4. Absent-stakeholder problem. In housing and zoning, the most affected party — the person who would live in the units if they were built — is not in the room, cannot vote, and has no formal channel. The same applies to future generations on climate and infrastructure, to future patients on health facility siting, to non-human species on land use decisions. A model can give absent and non-human stakeholders a computable voice by making their interests explicit rather than ignored.

5. Coupling asymmetry. Citizens experience issues as separate single-topic problems: traffic is bad, rent is high, the school is overcrowded. These are often coupled outputs of one system — adding housing without transit loads roads; fixing roads without addressing zoning raises rents. No single stakeholder sees the links. Interactive scenario tools let participants trace consequences across the whole system before committing.

6. Strategic asymmetry. Some stakeholders have direct incentive to misrepresent. A developer underestimates traffic impact; an opponent inflates it. Both are gaming a rhetorical process where numbers are assertions, not shared objects. A shared computational model doesn't prevent bias in the inputs, but it makes the model inspectable — opponents can challenge assumptions rather than trading competing assertions — and separates the structural question (what does this model predict?) from the values question (what should we do?).

7. Emotionally radioactive issues. Some decisions are too charged to surface in public deliberation — conversations shut down before they start. Structured and sometimes anonymous engagement tools can make these issues discussable by lowering the social cost of honest expression.

These are not problems that more voting solves. They are problems of information structure.

The lineage: cybernetics and governance

Norbert Wiener and feedback

Norbert Wiener coined cybernetics in 1948 — the science of communication and control in systems, whether animal or machine (1948). The core insight: all purposeful behavior involves feedback. Governance, as much as thermostats, is a feedback system. This framing invited a different question about democracy: not just who votes, but how information flows through a political system and whether it has the feedback properties needed to correct itself.

Stafford Beer and the Viable System Model

Stafford Beer took this most seriously. His Viable System Model (VSM) described the minimum structural conditions for any organization to remain viable — able to adapt to a changing environment (1972). Five recursively nested systems cover operations, coordination, control, intelligence (sensing the environment), and policy.

Beer argued that most real organizations violate the requisite variety principle (Ashby's Law: a controller must have as much variety as the system it controls) and are therefore doomed to fail in complex environments.

The most radical application was Project Cybersyn (1971–73) in Allende's Chile — Beer's attempt to build a real-time cybernetic nervous system for a national economy. Factory production data flowed to a central operations room; economic deviations triggered alerts; workers could flag problems upward. It was destroyed in the 1973 coup but remains the most serious real-world attempt to apply cybernetic principles to democratic governance. Eden Medina's Cybernetic Revolutionaries (2011) is the definitive account.

Anticipatory democracy

Alvin Toffler coined the term in the 1970s, building on his Future Shock (1970) thesis that social systems were being overwhelmed by the pace of change. His argument: democratic institutions needed to develop anticipatory capacity — systematic futures thinking built into governance — or they would always be reacting to crises rather than shaping them.

The anticipatory democracy tradition drew on futures studies, citizen participation methodology, and systems thinking: - Long-range scenario planning as a civic practice, not just a corporate one - Structured public participation in futures work (citizens' assemblies, futures workshops) - Institutions that explicitly represent future generations

The contemporary futures literacy movement — championed through UNESCO — is the current institutional expression: teaching people to use imagination about the future as a cognitive and civic tool.

Contemporary tools: the civic-tech layer

Pol.is

Pol.is is an opinion mapping platform built around a "wiki-survey": participants submit short statements; others vote agree, disagree, or pass. The vote matrix is reduced with SVD/PCA and clustered to reveal opinion groups.

The key output is not a poll result but a map: a 2D projection where people who voted similarly cluster together, and where each statement's position indicates whether it unifies or divides groups. The system can surface consensus statements — things agreed on across otherwise opposing groups — that standard Q&A or polling would never find.

The design is notable for what it omits: no threaded replies, no engagement metrics, no upvotes. Reply threading and engagement metrics reward conflict and virality; removing them changes what behavior gets amplified. A provocative statement gets no more attention than a careful one; you cannot dominate a Pol.is conversation by posting frequently.

Used in Taiwan's vTaiwan process (Uber regulation, 2015), by the Audrey Tang-era digital ministry, and broadly in civic consultation.

Decidim

Decidim (open-source) — participatory democracy platform used by cities (Barcelona, Helsinki, NYC) for participatory budgeting, planning, and policy co-design. Built on the premise that digital civic infrastructure should be publicly owned.

Citizens' Assemblies

Randomly selected citizens deliberating intensively on a defined question (climate, abortion, electoral reform) and producing recommendations. Ireland's experience (marriage equality, abortion) is the benchmark. The KNOCA network specifically tracks climate assemblies.

The key design feature: stratified random selection ensures that the assembly reflects the broader population, not just the people motivated enough to show up. This directly addresses the distributional engagement asymmetry.

Computational democracy: structural responses to named dysfunctions

A broader family sometimes called computational democracy uses mathematical or algorithmic structure to address specific failure modes of standard deliberative and voting processes.

Quadratic voting (QV) — the cost of casting n votes on a single issue is n^2, paid from a fixed budget of "voice credits." A voter who cares intensely can accumulate votes, but at increasing cost. This redistributes influence toward breadth of preference rather than intensity alone. Proposed by Glen Weyl and Eric Posner (2018); piloted in the Colorado state legislature and in Gitcoin public-goods funding rounds.

Liquid democracy — voters can vote directly or delegate their vote to someone they trust, who can in turn redelegate. The result is a weighted directed graph of trust. The motivation: direct democracy requires everyone to be informed about everything; representative democracy fixes expertise at the cost of granularity; liquid democracy allows fine-grained delegation by topic.

Conviction voting — votes accumulate over time; the longer you maintain a position, the more weight it carries. Rewards sustained belief over reactive rushes; prevents last-minute vote surges.

Futarchy — Robin Hanson's proposal: "vote on values, bet on beliefs." Policy is decided by prediction markets rather than majority vote. Voters define goals; markets aggregate distributed knowledge about which policy best achieves them. The proposal's intellectual clarity is not yet matched by deployment at scale, but it foreshadows the direction that coupling simulation models to governance processes would point.

Each mechanism addresses different failure modes: | Mechanism | Failure mode addressed | |—|—| | Quadratic voting | Intensity asymmetry | | Liquid democracy | Expertise gap | | Conviction voting | Temporal manipulation | | Pol.is | Dominance dynamics of discourse | | Citizens' Assembly | Distributional engagement asymmetry | | Simulation | All of the above (makes the model inspectable) |

MRV: connecting simulation to incentives

The gap between modeling and governing is partly a measurement problem. A model can tell you what a watershed is doing; governance needs a mechanism for rewarding people who improve it.

Regen Network's contribution is a methodology framework for MRV (Measurement, Reporting, and Verification) of ecological state. The idea: if you can measure ecological health credibly and at scale, you can attach economic value to ecosystem restoration and create market incentives for regenerative land management.

The MRV approach is significant because it attempts to bridge simulation (modeling what a system is doing) and incentive design (rewarding people for improving it). If simulation tells you the state of a watershed and MRV provides a verified measurement protocol, the combination creates the feedback loop that governance needs.

Adjacent efforts: Pachama (satellite and AI-based verification for forest carbon projects), Global Forest Watch (near-real-time deforestation monitoring), iNaturalist (~200M biodiversity observations as civic science).

The doughnut economics framework

Kate Raworth's doughnut economics provides the civic policy framework that most directly operationalizes planetary boundaries for local use. Two concentric rings:

- Outer ring: ecological ceiling — the planetary boundaries discussed in Chapter 1.
- Inner ring: social foundation — the minimum standards for human dignity (food, water, health, education, gender equality, etc.).

The "safe and just space for humanity" is the doughnut between them: above the social foundation, below the ecological ceiling. The City Portrait methodology (from DEAL, the Doughnut Economics Action Lab) applies this framework at the city scale, creating a dashboard that maps a city's performance against both rings.

This is the right conceptual frame for the scoring problem described in Chapter 7: a civic simulation that shows whether a proposed policy moves the community toward or away from the doughnut.

The gun problem: tools and their uses

There is a tension worth naming explicitly. The argument throughout this book has been that simulation tools lower the barrier to civic reasoning. But tools are recruited by whatever interests can use them. A high-quality watershed model can be used by a community group to challenge a developer's traffic impact study — or it can be used by the developer to produce authoritative-looking outputs that justify predetermined conclusions.

This is the "gun problem": a gun is a tool that is arguably apolitical, but quickly becomes recruited by different parties with different agendas because of its utility. Do the oppressed become entrapped by using the tools of the oppressor? Simulation is similar. The existence of civic simulation tools does not guarantee that they will be used to represent public interests — they will be used by whoever has access and motivation.

The response is not to withhold the tools, which would only leave the field to the parties who already have resources. The response is to: - Make the models open-source and inspectable, so that strategic misuse can be identified and challenged. - Build provenance and versioning into published models, so that changes can be tracked. - Create community review mechanisms — peer review for civic simulations — that surface contested assumptions. - Prioritize building tools for parties who currently lack them: community groups, indigenous communities, environmental NGOs, not just government agencies and developers.

The right not to be modeled

This book argues that data about places should be public and that the models built on that data should be open. The gun problem already qualified the argument: tools are recruited by power. This section confronts a harder qualification. For the Indigenous nations of the Pacific Northwest, where this book is written, being modeled — counted, surveyed, sampled, mapped — has more often been an instrument of dispossession than a remedy for it. A serious case for civic simulation has to hold that history at full weight, and accept the right that answers it: the right not to be modeled.

Earned skepticism

The skepticism is not abstract. Between 1990 and 1994, researchers at Arizona State University collected blood samples from members of the Havasupai Tribe for what the community understood to be diabetes research. The samples were then used, without consent, in studies of schizophrenia, inbreeding, and population migration — the last contradicting the tribe's own account of its origins. A tribal member, Carletta Tilousi, discovered the reuse in 2003 by attending a dissertation defense. The 2010 settlement returned the samples and paid $700,000 to forty-one tribal members; it set no legal precedent (Harmon 2010). In 1979, researchers surveyed alcohol use among the Iñupiat of Barrow, Alaska, and released the findings to the press before the community had seen them; the resulting coverage branded the town a “society of alcoholics” (Foulks 1989). The Ancient One — known to archaeology as Kennewick Man — spent twenty years as a contested specimen after his discovery on the Columbia in 1996, studied under court order over the objection of five claimant tribes, before DNA analysis confirmed what the tribes had said all along and he was reburied in 2017.

Linda Tuhiwai Smith opens Decolonizing Methodologies (1999) with the sentence researchers in this field learn first: “The word itself, ‘research’, is probably one of the dirtiest words in the indigenous world’s vocabulary.” The pattern she names has been measured. A global review of 140 climate studies drawing on Indigenous knowledge found that 87 percent followed an extractive model — outside researchers using the knowledge with minimal community participation or decision-making authority (David-Chavez and Gavin 2018). Climate research, the field closest to this book, is not the exception to the history. It is a current example of it.

Data about land is data about people

A civic simulation project might hope to stand apart from that history by modeling only geophysical data: terrain, water, vegetation, heat. The hope does not survive contact with the definitions. The CARE literature defines Indigenous data in three categories, and the first is data about environment — lands, skies, waters, resources, and non-humans (Carroll et al. 2020). A bioregional model of the Willamette Valley is a model of Kalapuya homeland whether or not a single person appears in it.

Location itself can be the harm. Federal law has recognized this for decades: the Archaeological Resources Protection Act (1979) explicitly exempts the locations of archaeological sites from public-records disclosure, because publishing a site's coordinates is an invitation to loot it. The tension is live in Oregon. The state's LiDAR is free to download — a policy this book would ordinarily celebrate — and that same data can reveal Kalapuya earthwork mounds in the Calapooia watershed, some of them burial sites, which the Confederated Tribes of the Grand Ronde Community of Oregon hold to be highly sensitive. Open data and sacred ground occupy the same coordinates. There is also a quieter, structural harm, which the legal literature calls a confidentiality catch-22: to defend a sacred place in an environmental review, a tribe must often describe it into the public record, exposing the very thing it is defending; staying silent forfeits the decision. And there is harm by omission — what the National Congress of American Indians calls the Asterisk Nation: populations too small for the statistical tables, replaced by a footnote, erased by aggregation.

Sovereignty, not courtesy

The frameworks that answer this history are claims of jurisdiction, not requests for sensitivity. Indigenous data sovereignty, as the United States Indigenous Data Sovereignty Network defines it, is the right of Native nations to govern the collection, ownership, and application of their own data — a right that derives from inherent sovereignty over peoples, lands, and resources, not from research ethics (Kukutai and Taylor 2016). The UN Declaration on the Rights of Indigenous Peoples (2007) grounds it at treaty level: Article 31 affirms the right to "maintain, control, protect and develop" cultural heritage and traditional knowledge, including knowledge of "the properties of fauna and flora"; Article 19 sets the standard of free, prior, and informed consent. United States practice runs on a weaker rail — Executive Order 13175 (2000) requires consultation, and the US endorsement of the Declaration in 2010 explicitly read consent down to "a process of meaningful consultation." Consultation is a meeting; consent is a decision. The gap between them is where most of the conflict lives. Oregon at least instruments the process: Senate Bill 770 (2001) obliges every state agency to maintain tribal-relations policies, reporting, and an annual government-to-government summit.

The operational frameworks are specific. OCAP — Ownership, Control, Access, Possession — was formulated by First Nations health researchers in Canada in 1998, and its fourth term is the sharpest: possession, physical stewardship of the data, "the mechanism by which ownership can be asserted and protected" (First Nations Information Governance Centre). The CARE Principles — Collective Benefit, Authority to Control, Responsibility, Ethics — were drafted at an Indigenous-led workshop in Gaborone in 2018 and published by Stephanie Russo Carroll and colleagues in 2020, framed to complement the open-science FAIR principles: be FAIR and CARE. The friction between those two acronyms — openness and reuse by default versus authority to control — is not a wrinkle to be ironed out. It is the honest shape of the problem (Carroll et al. 2021). The machinery for living inside the friction already exists: the TK and Biocultural Labels of Local Contexts (founded by Jane Anderson and Kim Christen in 2010) attach community-authored permissions and provenance to data as machine-readable metadata — governance that travels with the dataset.

Refusal is governance

Audra Simpson, in work that began as an ethnography of her own Kahnawà:ke Mohawk community, named the practice of refusal: declining to make some knowledge available at all, not as a gap in the record but as an exercise of jurisdiction over it. Sovereignty, she writes, "involves a calculus ethnography of what you need to know and what I refuse to write in" (Simpson 2007, 2014). Eve Tuck and K. Wayne Yang extend the argument to research itself, in the form of an axiom: "there are some forms of knowledge that the academy doesn't deserve" (Tuck and Yang 2014). Maggie Walter shows that even well-intentioned official statistics about Indigenous peoples arrive pre-framed as deficit — the five Ds of what she calls 5D data: disparity, deprivation, disadvantage, dysfunction, difference — so that being counted on someone else's terms is its own injury (Walter 2016; Walter and Andersen 2013). And the collective dimension breaks the open-data world's basic unit of permission: individual consent cannot protect a group, because one member's opt-in can expose the whole community — for a tribe, a relative's DNA or a neighbor's survey answer is never only theirs (Tsosie, Yracheta, and Dickenson 2019).

For a book about simulation, the conclusion is uncomfortable and unavoidable: the absent-stakeholder argument earlier in this chapter says a model can give the voiceless a computable voice, and the refusal literature answers that a voice imposed is not a voice. Both are true. A community that has been modeled by others for two centuries may reasonably regard one more model — however open, however well-meant — as the same old instrument with a friendlier license. The right not to be modeled is the data-sovereignty form of a right this book otherwise champions in every chapter: the right of a community to decide its own future. (The phrase is this book's own synthesis; its anchors are Simpson's refusal of legibility, C. Matthew Snipp's argument that decolonized data begins with the power to determine who is counted (Snipp 2016), and the group-privacy literature's finding that data harms attach to collectives, not only persons (Taylor, Floridi, and Sloot 2017).)

Refusal also has a positive, constructive form. Te Hiku Media, a Māori broadcaster, built state-of-the-art speech recognition for te reo Māori from three hundred hours of speech contributed by its own community in ten days, governed by the Kaitiakitanga License: the data is not owned but cared for, used only for Māori benefit, and not openly released. "Data is the last frontier of colonization," says Keoni Mahelona, who built the platform — and when OpenAI released a model trained on scraped te reo speech, Te Hiku published its refusal in so many words. The San peoples of southern Africa issued their own research-ethics code in 2017, with consequences attached: researchers who violate it are refused future collaboration. The Native BioData Consortium keeps Indigenous genomic samples on sovereign land in South Dakota, governed by the tribes they came from. Working models, working institutions — built without open release, under the community's own law. The lesson for civic simulation is not that openness is wrong; it is that openness is a choice that belongs to the people the data is about.

The ground truth in Oregon

Nine federally recognized tribes hold sovereignty in Oregon: the Burns Paiute Tribe; the Confederated Tribes of the Coos, Lower Umpqua and Siuslaw Indians; the Confederated Tribes of the

Grand Ronde Community of Oregon; the Confederated Tribes of Siletz Indians; the Confederated Tribes of the Umatilla Indian Reservation; the Confederated Tribes of the Warm Springs Reservation of Oregon; the Coquille Indian Tribe; the Cow Creek Band of Umpqua Tribe of Indians; and the Klamath Tribes. Within living memory the federal government attempted to delete most of them administratively: the Western Oregon Termination Act of 1954 severed federal relations with more than sixty tribes and bands at a stroke — termination was, among other things, a data operation, a closing of rolls and a liquidation of trust lands. Restoration took three decades of litigation and legislation, from Siletz in 1977 to Coquille in 1989. Skepticism toward the state's record-keeping is not a cultural posture here. It is institutional memory.

And the tribes are not waiting to be modeled. The Confederated Tribes of the Umatilla Indian Reservation manage their natural resources through the First Foods framework — water, salmon, deer, cous, huckleberry, in the order they are served in the longhouse — articulated by Eric Quaempts and published in the peer-reviewed literature (Quaempts et al. 2018). Read it as a modeler and the recognition is immediate: an explicit objective function, a defined state space, monitoring touchstones — hydrology, geomorphology, connectivity, riparian vegetation, aquatic biota — running a real natural-resource department and a formally adopted climate adaptation plan (2022). It is not traditional knowledge waiting to be ingested into someone's model. It is a model, theirs. The Columbia River Inter-Tribal Fish Commission, founded by four treaty tribes in 1977, runs its own fish-genetics laboratory in Idaho and an ocean-prediction center in Astoria, and contributes data to the region's shared StreamNet system on terms its members set — "as practicable," the agreements say, which is partial openness as policy. The Klamath Tribes hold water rights with time-immemorial priority that survived their own termination (United States v. Adair, 1983), drove the largest dam removal in United States history, and in October 2024 watched the first salmon in 112 years swim past the old dam sites. The question was never whether these nations can do data, models, or restoration at scale. The question is who controls the model, and the answer these institutions give is: they do.

What this asks of the publication layer

The publication layer of Chapter 8 acquires obligations here, and they are buildable:

- Uninvited is the wrong default. Do not model Indigenous lands as a service performed for communities who did not ask. A model of a place can carry the knowledge of the people who have stewarded it longest — and only they can decide that. Where a tribal nation is interested, enter through its own institutions — its natural-resource department, its climate program, intertribal bodies like the Affiliated Tribes of Northwest Indians — under its protocols, with free, prior, and informed consent as the standard, not consultation (Climate and Traditional Knowledges Workgroup 2014).

- Publishing must include not-publishing. A publication layer that can only disclose is an instrument of the catch-22. Sensitive-location redaction, access tiers, and community-controlled layers are first-class features, not exceptions; ARPA has treated location confidentiality as good law since 1979.
- Governance travels with the data. Provenance metadata should carry Local Contexts TK and Biocultural Labels natively, so a community's terms are machine-readable wherever the data goes. OCAP's possession principle has an architectural meaning: community-held deployments, not just community-flagged records in someone else's database.
- Expect and respect refusal. A declined invitation is a governance outcome, not an obstacle. And notice that redaction has its own information cost — a conspicuous blank can advertise what it conceals — which is a design problem to be solved with the community, not for it.
- Reciprocity is measurable. The 87 percent extraction figure has a complement: in community-initiated studies, findings flowed back to the community almost always; in externally initiated ones, about 15 percent of the time (David-Chavez and Gavin 2018). Which kind of project this is shows up in the data flows, not the intentions.

The gun problem ended by arguing that the answer to misuse is not to withhold the tools. This section is the bounded exception that proves the rule: sometimes withholding is the tool, and the hand on it is sovereign. A civic model exists to serve the community it depicts. The community does not exist to serve the model.

Further reading

- Laloux, F. (2014). Reinventing Organizations. Nelson Parker.
- Raworth, K. (2017). Doughnut Economics. Chelsea Green.
- West, G. (2017). Scale. Penguin Press.
- Bateson, G. (1972). Steps to an Ecology of Mind. Chandler.
- CTUIR (2022). CTUIR Climate Adaptation Plan. Confederated Tribes of the Umatilla Indian Reservation. (The First Foods framework in operation.)
- Cifor, M., Garcia, P., et al. (2019). Feminist Data Manifest-No. manifestno.com.
- Local Contexts — TK and Biocultural Labels: localcontexts.org.
- Te Hiku Media — the Kaitiakitanga License: kaituhi.nz/kaitiakitanga; and Hao, K. (2022, April 22). A new vision of artificial intelligence for the people. MIT Technology Review.
- San Code of Research Ethics (2017). South African San Institute.
- Native BioData Consortium: nativebio.org.

- Zong, J., & Matias, J. N. (2024). Data Refusal from Below. ACM Journal on Responsible Computing, 1(1).
- Status of Tribes and Climate Change (STACC) Reports, Vols. 1–2. Institute for Tribal Environmental Professionals, 2021 / 2025.

Open questions

- Under what conditions does a civic model increase equity rather than reinforce power asymmetries?
- How do we build peer review mechanisms for civic simulations that are lightweight enough to actually be used?
- What is the relationship between the MRV framework in regenerative finance and the calibration problem in simulation?
- What does free, prior, and informed consent mean for a planetary-scale model whose grid cells cover every homeland on Earth? Is there a scale at which the consent question changes character — and who decides?
- When a community refuses, what does the model show? A conspicuous blank can advertise exactly what it conceals; how is refusal rendered without becoming disclosure?
- The absent-stakeholder argument and the refusal argument both claim the moral ground of voice. Where exactly is the line between giving a voice and imposing one?
- As of this writing there is no peer-reviewed literature on Indigenous data sovereignty and digital twins or civic simu-

lation specifically; the nearest work governs environmental research data (Williamson, Provost, and Price 2023). Civic simulation will arrive in Indigenous homelands before its ethics literature does. What obligations follow from going first?

TODO / White-space

- FIGURE: VSM diagram (Viable System Model — five nested systems).
- FIGURE: Pol.is output visualization — opinion map with cluster structure.
- FIGURE: Doughnut economics diagram (planetary boundary ceiling + social foundation floor).
- FIGURE: Computational democracy mechanisms table (failure mode vs. mechanism).
- PLACEHOLDER: Interview with Audrey Tang / vTaiwan participant on digital civic participation.
- PLACEHOLDER: Interview with Citizens' Assembly facilitator on the deliberative process.
- PLACEHOLDER: Indigenous data-sovereignty practitioner or tribal natural-resource staff voice for "The right not to be modeled" — approach through tribal office or an intertribal body (ATNI), per the protocol notes in 83-reference-voices; do not approach individuals directly.
- TODO: Expand on the commons governance literature (Ostrom) and its relationship to watershed models.

- TODO: Verify before print — Havasupai sample count (sources range 200–400; settlement figures are solid); SB 770 “first state in the nation” phrasing; Warm Springs adaptation plan date; Simpson “refuse to write in” page number (cited as p. 72) against the print Junctures issue. Walter’s 5D wording is verified verbatim against Walter 2016 (ANU Press, ch. 5). See primer-devlog/20260612-data-sovereignty-research-notes.md for sources.
- TODO: A map of the nine tribes’ lands seems like the obvious FIGURE for this section, and the section itself explains why it is not obvious. Decide with care — existing public maps (e.g., Oregon Blue Book) may be appropriate to reference rather than redraw.

Chapter 11

What Now

> The book has moved from what we value, through what threatens it, to the proof that simulation can change outcomes, the craft of building and publishing such models, and the participatory traditions that give them a political home. This closing chapter asks what to do with all of it — what the world looks like if we build this, what it will honestly take, and how to take part.

Our intent here is to close the loop the book opened: to return to the room from the preface, show what changes when the model is finally in it, and turn the argument into an invitation.

Back to the room

We began in a room like this one: residents of a fire-scarred neighborhood working through fuel management and evacuation routing, with the modeling that bears on those decisions nowhere they can reach it. Where it exists at all, it sits on agency servers and under consultants' contracts — public in principle, inaccessible in practice. One kind of knowledge lives in models, another in lived experience, and neither can be properly examined alongside the other.

Imagine the same room with the gap closed. The fire-spread model is open on a shared screen and on the phones of the people in the room. A resident who has watched the wind funnel through her canyon for thirty years changes the wind assumption and shows the modeler where the model is wrong; the modeler changes the fuel layer and shows the resident what defensible space actually buys her; a council member compares two fuel-reduction budgets and sees, with honest uncertainty bands, what each one moves. Nobody has been turned into an expert. The argument has simply become a shared object that all of them can push on. That is the whole proposition of this book, reduced to a single room.

Nothing in that scene is science fiction. The fire-spread algorithms are standard, and the terrain and fuel data are largely public. The publication layer that puts the model on everyone's phone is the one buildable missing piece — and, as the last chapters argued, it is now buildable.

The world if we build this

Scale that room up and the change is not that decisions become automatic — models do not make decisions, and should not. The change is in the quality of the disagreement. Today, contested public decisions are largely arguments between assertions: a developer's traffic estimate against an opponent's, a price tag against a fear, a confident projection against a vague dread. When the model is a shared, inspectable object, the argument moves to where it belongs — to the assumptions. You stop trading conclusions and start examining premises. The values question (what should we want?) separates cleanly from the structural question (what does this system actually do?), and both get clearer for being pulled apart.

In that world, the chronic and the deferred become visible. The aquifer's depletion shows up before the wells fail, not after. The avoided disaster — the fire that did not jump the ridge, the flood the wetland absorbed — becomes legible enough to defend in a budget. The absent stakeholders — future residents, downstream communities, species with no vote — get a computable voice, because someone made their interests an explicit part of the model rather than an externality outside it.

This is not a promise that the right side wins. Tools are recruited by whoever can use them, and a polished model can launder a predetermined conclusion as easily as it can illuminate an honest one. What the publication layer changes is not who holds power but what can be hidden: an open, inspectable, versioned model is one that opponents can challenge at the level of its assumptions instead of its authority. That is a smaller claim than "simulation fixes democracy," and a more defensible one.

It is worth being plain about the standard here, because it is lower and more honest than "accurate prediction." We will not produce perfect models, and we do not need to. The thing we are measuring against is not a perfect forecast; it is rhetoric, two people connecting a proposed change directly to a hoped-for outcome and leaping clean over the rigor in between. Against that bar almost anything helps. A rough, shared, inspectable model gets people used to thinking in systems (Meadows 2008), gives them a feel for how a thing responds when you push on it, and turns “trust me” into “try it and see.” The deepest value of this work may be less any single answer than the practice itself: getting communities used to arguing over a chessboard with the pieces in view, rather than across a table that has nothing on it but conviction.

What it will honestly take

We have tried throughout to be candid about the difficulties, and this is the place to gather them.

Calibration remains the hard problem. A model that is wrong but confident is worse than no model. The discipline the book asks for — labeling clearly whether a model is illustrative, scenario-based, or predictive; showing uncertainty as a range rather than a number; making data provenance a visible variable — is not optional polish. It is the difference between a civic tool and a more sophisticated way to mislead.

The data and the maintenance are real work. The open datasets exist, but turning them into a calibrated, current model of a specific place is labor, and keeping it current is more. A publication layer lowers the cost of authorship; it does not abolish the cost of being right.

Adoption is a social problem, not a technical one. A model in the room only helps if the process is ready to use it — which is why the participation traditions of Chapter 9 are not a coda but a load-bearing part of the argument. The tool and the process have to arrive together.

The incentives are not yet there. Nobody is paid to build the publication layer the way they are paid to build the model. This is the gap the book is, in part, an argument for filling — through open infrastructure, public funding, and the same volunteer energy that built OpenStreetMap and Wikipedia. Sustaining shared resources of this kind is a governance problem with a long literature of workable, locally crafted solutions (Ostrom 1990).

None of these is a reason not to begin. They are the actual shape of the work.

An invitation

This is a primer in the original sense: a starting point, and an invitation to take part. The work is genuinely distributed, and different readers hold different pieces of it.

- If you build things — the publication layer, the widget libraries, the open templates, the scenario gallery are mostly unbuilt. The hardest and most valuable work is not a better model engine but the connective tissue that lets a model travel.

- If you work in a domain — a watershed, a fishery, a city's heat map, a supply chain — the scenarios worth building are the specific ones you already understand. The reference scenario library is a place to start, not a limit.
- If you teach — a model students can break is worth more than a chapter they can only read. Civic simulations are, among other things, the best instruments we have for building intuition about coupled systems.
- If you organize or decide — ask for the model. Ask for it to be open, inspectable, and in the room. The most powerful thing a community can do is refuse to accept conclusions detached from the models that produced them.

You do not need permission, and increasingly you do not need a programming course. The barrier between describing a model and running one is now a conversation. What remains is the choice to treat modeling the world as a civic practice rather than an elite one.

The world is still worth it

We opened with what we value — a four-billion-year experiment in organized complexity, encountered each of us through some specific and personal door: a mountain range, a coral reef, a river that came back. The argument of this book has been technical in places, but its root is not. It is that this world is worth the trouble of understanding clearly, together, at the scale the challenges actually operate on.

The tools for that are, for the first time, within reach of ordinary people. Whether we pick them up is not a question simulation can answer. It is the one question that was always ours — and hope, as Solnit reminds us, is not a forecast of the outcome but a commitment to act inside the uncertainty (2016).

Further reading

- Raworth, K. (2017). Doughnut Economics. Chelsea Green.
- Bratton, B. H. (ongoing). A General Theory of Simulation. Antikythera.
- Russell, B. (1999). "Headmap Manifesto." http://dpya.org/en/index.p Headmap_Manifesto(1)_-_Ben_Russell — an early, prescient vision of location-aware computing: the physical world annotated with information, every place carrying data and stories. A spiritual ancestor of treating place itself as something legible and authorable.

Open questions

- What is the minimum viable institution — public, nonprofit, or commons — that could sustain a civic simulation publication layer?

- How do communities build the habit of asking for the model, and what changes when they do?
- What does success look like in ten years, and how would we know we had it?

TODO / White-space

- FIGURE: The reimagined room — the council meeting with the model in it (callback to the preface).
- PLACEHOLDER: Short closing voices — practitioners and community members on what they would build first.
- TODO: Decide whether this chapter ends the book or is followed by a brief afterword / acknowledgments.

Reference — Scenario Library

> A curated catalogue of scenarios worth building, organized by scale (local, regional, planetary) and by topic domain. Each entry includes a brief description, the key dynamics being modeled, the data sources required, the audience, and an estimate of buildability. This is reference material — a living catalogue, expected to grow and be re-prioritized as scenarios are developed and the tools mature. The narrative case studies that motivate it are in Chapter 4.

Criteria for inclusion

Scenarios in this library meet most of these criteria:

- Specific — a real place or real class of situation, not a generic case study.
- Intersectional — interesting behavior arises from coupled dynamics.
- Counterfactual leverage — the model makes visible something that standard political argument cannot.

- Accessible audience — there is a natural community of people who have a stake in the question.
- Buildable — either the data exists or the model can be illustrative without it.

Local: urban and city planning

Stakeholders: cities, planners, citizens, developers, neighborhood groups.

This is the scale where participation asymmetries are sharpest and the audience is most immediate.

Land use

New suburb approval - Dynamics: infrastructure cost per unit, vehicle-miles-traveled (VMT) increase, fiscal return-on-investment to city over 20 years, habitat conversion at the fringe. - Data: municipal fiscal data, traffic counts, habitat maps (NLCD or local GIS). - Audience: city council, planning departments, neighborhood groups. - Buildability: medium. The fiscal model is straightforward; habitat quantification requires spatial data.

Urban growth boundary change - Dynamics: habitat loss at the edge, housing cost pressure at the core, service extension cost. - Key coupling: growth at the fringe and cost pressure at the core are linked through the same housing market. - Data:

parcel data, regional housing market data, utility extension cost data. - Buildability: medium.

ADU legalization - Dynamics: density gain, water and sewer load, neighborhood character change, absent-stakeholder effect (future residents not in the room). - Data: parcel data, utility capacity data, regional rental market data. - Buildability: low to medium. The structural argument (more units → lower rents) is simple; the spatial distribution is harder.

Stormwater cost-benefit - Dynamics: impervious surface fees versus detention basin construction versus distributed bioswales; peak-flow reduction; water quality. - Key coupling: cost is concentrated (each impervious surface owner); benefit is diffuse (reduced flooding downstream). - Data: DEM, impervious surface map, stormwater infrastructure map. - Buildability: medium. DEM-derived hydrology provides the skeleton.

Green infrastructure

Urban heat island - Dynamics: tree canopy cover, pavement type, building density → surface temperature differential. - Data: land cover (NLCD), DEM, Landsat or MODIS surface temperature. - Audience: city health departments, urban planners, equity advocates (heat disproportionately affects low-income neighborhoods). - Buildability: medium-low. The analytic insolation layer plus land cover provides a fast first version.

Urban tree planting - Dynamics: heat mitigation, air quality, carbon sequestration, root damage to pipes. Multiple weak benefits that are collectively strong. - Data: street tree inventory, canopy cover map, air quality monitoring. - Buildability: low to medium.

Green roof mandate - Dynamics: runoff delay, energy savings, structural load, cost. - Buildability: low.

Air quality - Dynamics: traffic and industrial sources, prevailing wind (from ERA5 climatology), health burden by census tract. - Audience: public health departments, environmental justice advocates. - Data: EPA air quality monitoring, traffic counts, industrial permits. - Buildability: medium.

Streets

Bicycle lane addition - Dynamics: modal shift, parking loss, injury rates, retail impact. - Key coupling: the distributional engagement asymmetry — twelve angry parking-space losers versus three thousand cyclists. - Data: traffic counts, retail sales data, injury records. - Buildability: low. The structural argument is simple; precise calibration requires local data. - This is the canonical introductory scenario for civic simulation.

Road lane reduction - Dynamics: rerouting, delay, induced demand (traffic fills any available capacity, so lane reductions often reduce total traffic). - Key insight: induced demand makes road-widening counterproductive; the model makes this visible. - Buildability: medium. Network flow models exist; the induced demand calibration is harder.

Bus route change - Dynamics: ridership, coverage, transfer burden on low-income riders. - Data: GTFS (General Transit Feed Specification) data — available for most US cities. - Buildability: medium.

Resilience

Power grid failure cascade - Dynamics: N-1 contingency analysis, blackout spread, restoration sequence. - Data: grid topology (partially available from utility regulators). - Buildability: high (network structure is well-characterized).

Water system failure after earthquake - Dynamics: pipe break probability from ground motion, pressure loss, hospital risk. - Data: seismic hazard maps (USGS), water system GIS (available from utilities). - Buildability: medium-high.

Regional: ecology and environment

Stakeholders: environmental NGOs, universities, wildlife agencies, concerned citizens, governments.

Water

California/Central Valley groundwater depletion (Cadillac Desert) - Dynamics: aquifer extraction, recharge rate, land subsidence, agricultural economic pressure. - Historical data available for regression testing; the model can be validated against observed drawdown rates. - Data: USGS NWIS groundwater level data, FAOSTAT agricultural production, DEM for subsidence. - Buildability: medium-high.

California water rights and the water atlas - Dynamics: water-rights allocation under prior appropriation, surface and groundwater accounting, and the visualization of who holds rights to what, made legible as a public, queryable map rather than a buried registry. - Why it belongs here: a publication-layer problem as much as a modeling one. The author contributed to the New California Water Atlas (with the Long Now Foundation), which assembled water-rights data into public-facing form; the scenario couples that data to the groundwater and snowpack dynamics above. - Data: California SWRCB water-rights records, USGS NWIS, snowpack and reservoir data; New California Water Atlas materials. - Buildability: medium. The data wrangling is the hard part; the rights accounting is bookkeeping, not new physics. - See the California water case study in Chapter 4.

River dam removal and salmon recovery - Dynamics: sediment flush, fish passage restoration, downstream flood risk change. - Data: USGS stream gauge data, NOAA fishery surveys. - Buildability: medium. Several dam removals (Elwha River) provide validation data.

Dead zones - Dynamics: agricultural runoff (nitrogen, phosphorus) $\rightarrow$ hypoxia $\rightarrow$ fish kill propagation. - Data: USGS nutrient load data, NOAA dead zone monitoring, FAOSTAT fertilizer use. - Buildability: medium.

Freshwater lens on small islands (Tuvalu) - Dynamics: sea-level intrusion, storm overwash, population viability of freshwater supply. - Data: USGS freshwater lens studies, sea-level data, DEM. - Key property: small domain, well-characterized physics, existing literature. - Buildability: medium-high. Good for demonstrating tipping point dynamics.

Chesapeake Bay nutrient cycle - Dynamics: farm runoff → hypoxia → oyster filtration → submerged aquatic vegetation recovery. - Data: Chesapeake Bay Program monitoring data (publicly available). - Buildability: medium.

Fire

Wildfire spread up canyon terrain - Dynamics: slope, aspect, fuel load (LANDFIRE), wind channeling, spotting. - Data: LANDFIRE national fuel and vegetation data, ERA5 wind, SRTM DEM. - Audience: fire agencies, emergency managers, homeowners in fire-prone areas. - LA County is a topical and high-stakes case as of 2025–2026. - Buildability: high. Standard fire spread algorithms (Rothermel model) are well-documented.

Post-fire debris flow risk - Dynamics: burn scar hydrophobicity, rainfall intensity, downstream community exposure. - Data: USGS post-fire hazard assessments, NOAA rainfall data. - Buildability: medium-high.

Fort McMurray 2016 (historical reconstruction / model validation) - Dynamics: extreme fire-weather onset, ember-driven spread into the wildland-urban interface, mass evacuation of a petro-city, the coupling of a fossil-fuel economy to the fire risk it amplifies. - Why this one: a well-documented, high-stakes historical event with an unusually rich record, which makes it a validation case rather than a hypothetical — run a spread model against what actually happened. John Vaillant's Fire Weather (2023) is the narrative companion; see also the Preface and Chapter 4 further reading. - Data: Canadian fire-weather indices, ECCC weather records, fuel and terrain layers, post-event reconstructions. - Buildability: medium-high. Standard fire-spread algorithms apply; the value is in the historical validation and the human/economic coupling, not new physics.

Biodiversity and land

Rewilding corridors (Banff or similar) - Dynamics: species movement, habitat fragmentation, land acquisition cost. - Data: wildlife tracking data, land ownership records, habitat quality maps. - Buildability: medium. Spatial ABM (GAMA Platform) is the right tool.

Animal-rights policy edge case (Oregon-style ballot measure) - Dynamics: a proposed ban on hunting, fishing, and pest control ripples into wildlife population dynamics, ranching and farming economics, disease vectors, and ecosystem balance. A test bed for representing non-human stakeholders explicitly in a model rather than leaving them as an externality. - Why it belongs here: the value is less prediction than surfacing second-order consequences that neither side foregrounds, and giving the absent stakeholders (the deer herd, the salmon run, the predators) a computable seat at the table (see the absent-stakeholder problem, Chapter

9). - Data: state wildlife population surveys, agricultural and ranching statistics, hunting/fishing license and harvest records. - Buildability: medium. The population dynamics are standard; the contested part is valuing outcomes across deeply opposed value systems.

Amazon tipping point - Dynamics: deforestation fraction → reduced rainfall via forest-moisture feedback → additional deforestation stress. - Key threshold: ~20–40% deforestation may trigger a regional rainfall collapse. - Data: Global Forest Watch deforestation data, ERA5 precipitation. - Buildability: medium. Illustrative model of the feedback is tractable; precise threshold calibration is not.

Coral bleaching threshold and reef restoration - Dynamics: SST anomaly → Degree Heating Weeks (DHW) accumulation → bleaching → mortality → recovery window (10–15 yr) → reef structure collapse → fishery and tourism impact. - Key coupling: fishing pressure reduction and genetic diversity of transplanted corals affect recovery probability; pelagic connectivity (large predators, manta rays) maintains herbivore suppression of algal overgrowth. - Data: NOAA Coral Reef Watch DHW data (real-time, global, public), fishery and tourism statistics, field survey data from active restoration projects. - Live example: emergentreefs.org — artificial reef deployment with structural optimization potential. - Buildability: medium. The thermal stress model (DHW) is analytic and directly computable from public SST data. The restoration ABM (genetic diversity × bleaching threshold) is a clean agent-based problem. The connectivity component requires species movement data (available for mantas from Marine Megafauna Foundation). - See the Chapter 4 case study for the full modeling framing.

Planetary: macro, supply chain, systemic risk

Stakeholders: investors, policymakers, journalists, analysts, anxious citizens.

Supply chain and geopolitical shocks

Strait of Hormuz closure - Dynamics: oil flow disruption → price shock → LNG rerouting → country-level GDP impact cascade. - Data: IEA energy flows, EXIOBASE I-O matrix, OEC trade data. - Key property: a news cycle driver with no public-facing quantitative model. There are currently many speculative articles about Hormuz with no quantitative grounding. - Buildability: medium-high (I-O shock propagation is analytically tractable). - See: Hormuz oil threads; Beginners guide to Hormuz

Suez / Panama blockage - Dynamics: container rerouting cost, delivery delay, inflation pulse. - Data: Comtrade, IHS container ship tracking. - Buildability: medium.

Semiconductor supply chain disruption - Dynamics: fab concentration risk (Taiwan), stockpile depletion timelines, downstream industry impact. - Data: USGS Mineral Commodity Summaries, IHS semiconductor market data. - Buildability: medium.

Agricultural export ban (wheat, rice) - Dynamics: food price cascade, import-dependent country stress. - Data: FAOSTAT, FAO food price index, Comtrade. - Buildability: medium. Food system I-O models exist.

Global agricultural inequality - Dynamics: the agrifood system's double failure, not delivering food security while driving biodiversity loss and emissions; subsidy structures, trade asymmetries, and who bears the externalities. - Why it belongs here: a structural scenario about distribution rather than scarcity. Pairs naturally with the food-price cascade above and with a pre-baked demographic run (the offline-coupling note in Chapter 7). - Data: FAOSTAT, FAO food-security indicators, Comtrade, agricultural subsidy databases. - Buildability: medium. I-O and accounting methods exist; the value is in making the distributional structure legible.

AI data-center water and power demand - Dynamics: rapid growth in compute → electricity and cooling-water demand → competition with agriculture and municipal supply in water-stressed regions; an information-economy boom coupled to a physical-resource constraint. - Why it belongs here: a fast-moving, under-modeled megatrend where the second-order effect (water and grid stress) is diffuse and deferred while the build-out is concentrated and immediate. Clean intersectional scenario across compute, energy, water, and food. - Data: data-center siting and capacity disclosures, utility load data, USGS water-use data, regional drought indices. - Buildability: medium. The accounting is tractable; the contested part is projecting compute demand.

Helium supply disruption - Dynamics: single-source supply chain (US Bureau of Land Management reserves, Qatar, Russia) → MRI, semiconductor fabrication, fiber optic production. - Key property: surprisingly brittle supply chain with broad industrial and medical dependencies. - Buildability: medium.

Megatrends and tipping points

Sea-level rise inundation - Dynamics: which cities, which infrastructure, which timelines, adaptation cost versus managed retreat. - Data: CoastalDEM (improved elevation data for coastal areas), sea-level projections from IPCC AR6, city parcel data. - Buildability: medium-high. The physics is well-characterized; the social response model is harder.

Permafrost thaw - Dynamics: methane release feedback, infrastructure damage (northern Russia, Canada, Alaska), subsidence. - Data: permafrost extent maps, MODIS land surface temperature, infrastructure location data. - Buildability: medium.

MENA warming - Dynamics: surface temperature rise in Middle East and North Africa (amplified because no soil moisture buffers CO_2 forcing); wet-bulb temperature threshold crossings; economic and displacement consequences. - Data: ERA5 historical, CMIP6 projections, WorldPop. - Key papers: AGU press release on MENA warming; JGR Atmospheres paper. - Buildability: medium. Reduced-complexity climate model (FaIR) can reproduce the forcing; regional downscaling adds complexity.

India heatwave mortality - Dynamics: extreme heat frequency under different emissions pathways → excess mortality → economic impact. - Data: India Meteorological Department, CMIP6 temperature projections, census demographics. - Buildability: medium.

Urban demographic divergence and migration - Dynamics: city-level age/sex structure, dependency ratios, and the migration-vs-natural-increase split, and how those shape labor supply, climate vulnerability, and political stability. A macro/structural scenario rather than a physical one. - Why it belongs here: a strong candidate for a pre-baked run (see the offline-coupling note in Chapter 7) — a validated demographic and migration trajectory that other scenarios (food pressure, heat mortality, services demand) can consume as a fixed input rather than re-simulating population dynamics. - Anchoring findings: across >10,000 cities (2000–2020) the dependency ratio fell from 0.87 to 0.59; smaller cities stayed younger than large ones, especially in Africa; pronounced male surpluses in parts of MENA consistent with labor migration; roughly 45% of urban growth from net migration, 55% from natural increase. - Data: WorldPop age–sex grids intersected with GHS-UCDB urban boundaries (the study's method); UN World Population Prospects. - Source: Zimmer, A., Brooks, N., Gaughan, A. E., et al. (2026). "Global divergence in urban demographic change and migration patterns." Nature Cities. https://doi.org/10.1038/s44284-026-00447-7 - Buildability: medium. The dataset method is published; the modeling value is in coupling it to impact scenarios.

AMOC collapse - Dynamics: Greenland melt → North Atlantic freshening → AMOC slowdown/collapse → European temperature and rainfall change. - Data: RAPID array ocean observations, Greenland mass balance (GRACE), CMIP6 projections. - Buildability: medium. Qualitative tipping point model is tractable; quantitative timing is not.

IPCC SSP scenario translation - Dynamics: take the SSP2-4.5 or SSP5-8.5 scenario outputs from IPCC AR6 and translate them into a specific community's terms: what does 2.5°C mean for drought frequency in the Sacramento Valley? - Data: IPCC Atlas (https://interactive-atlas.ipcc.ch/), ERA5, downscaled climate projections. - Buildability: medium.

Smaller-scale / speculative

eDNA freshwater monitoring (Denmark model) - Dynamics: taxonomic-independent community index for watershed health monitoring using environmental DNA. - Not a simulation per se — a data infrastructure problem — but a potential data input to regional models. - See: eDNA initiative.

African spurred tortoise and Sahel vegetation recovery - Dynamics: tortoise grazing compacts soil differently from other browsers, potentially fostering different vegetation recovery in degraded Sahel landscapes. - A smaller-scale example of nonlinear ecosystem interaction. - See: Tortoise Sahel article.

Post-car neighborhood - Dynamics: parking converted to housing and parks; walking and cycling modal share; retail and social outcomes. - A design fiction scenario that is increasingly becoming real policy in European cities.

Further candidates (from project notes)

A backlog drawn from `orbital-thinking`, kept brief until each is worked up or dropped:

- Texas power-grid disruption. A specific instance of the grid-failure cascade above (Winter Storm Uri 2021): generation loss, market-design failure, cascading outages, and mortality. Rich post-event data for validation.
- Spatial poverty mapping. Geographic prediction of poverty and vulnerability. Not a simulation so much as a data layer, but a strong pre-baked input (Chapter 7 offline coupling) for food, heat, and services scenarios.
- Global pension / retirement crisis. Systemic risk in pension funding under aging demographics. Couples naturally to the urban-demographic run above.
- Exposome. The cumulative burden of environmental stressors (air, heat, pollutants, noise) on human health, resolved by place and population.
- Post-growth quality of life. Wellbeing and inequality under aggressive climate mitigation: a scenario where the scored outcome is explicitly not GDP.

Building priority

A rough ordering by buildability and civic impact:

1. Urban heat island — analytic insolation + NLCD land cover. High impact; every city is a potential audience.
2. Bike lane / street reallocation — simple modal shift model. High civic relevance; illustrates distributional asymmetry clearly.
3. Wildfire spread — terrain + LANDFIRE + statistical wind. High stakes; topical; good visual output.
4. Hormuz closure — EXIOBASE I-O shock. No spatial component needed; high media relevance.
5. California groundwater — USGS data; historical regression test available. High stakes; rich data.
6. MENA warming — FaIR + regional downscaling. Important long-range scenario; underserved audience.

Data sources

- Global Forest Watch: https://www.globalforestwatch.org
- LANDFIRE: https://landfire.gov
- IPCC Interactive Atlas: https://interactive-atlas.ipcc.ch/
- EXIOBASE: https://www.exiobase.eu
- Comtrade Plus: https://comtradeplus.un.org

Further reading

- Meadows, D. H. (2008). Thinking in Systems. Chelsea Green.
- Reisner, M. (1986). Cadillac Desert. Viking.
- Davis, M. (2001). Late Victorian Holocausts. Verso.

Open questions

- What scenarios have the highest leverage for changing political outcomes, and how do we prioritize them?
- How do we maintain a living library of civic scenarios with consistent provenance and review?
- What is the right degree of simplification for each scenario tier?

TODO / White-space

- TODO: Add buildability ratings to each scenario entry.
- TODO: Add links to existing published models for each scenario where they exist.
- TODO: Add "worked examples" sub-chapters (12.1 Urban Heat Island, 12.2 Hormuz, etc.) with full implementation notes.
- FIGURE: Scenario map — geographic distribution of the scenarios in this library.

- FIGURE: Scenario comparison matrix (scale, domain, buildability, audience).
- PLACEHOLDER: Community voices on which scenarios matter most in their contexts.

Reference — Tools, Institutions, and People

> A catalogue of the software, organizations, and individuals that recur throughout the book. This is reference material, not argument: it is expected to grow and to go stale as tools and affiliations change. The chapters mention these in context; this is where the full list lives. See Chapter 5 (The Landscape of Simulation) for the paradigms these tools belong to.

Tools by paradigm

Paradigm	Tools	Notes
System dynamics (SD)	Vensim, Stella Architect, PySD	PySD runs Vensim/Stella models in Python. Stella is the accessible/education option.

Paradigm	Tools	Notes
Agent-based (ABM)	NetLogo, Mesa, Agents.jl, GAMA Platform, Repast HPC, AnyLogic, FLAME GPU 2	NetLogo is the research lingua franca; GAMA is strongest when geography is central; FLAME GPU 2 scales to millions of agents.
Input-output (I-O)	EXIOBASE, GTAP, UN Comtrade + custom matrices	EXIOBASE: ~200 countries × ~200 sectors with environmental extensions.
Computable general equilibrium (CGE)	GTAP	The dominant framework for trade-policy analysis.
Integrated assessment (IAM)	DICE, MESSAGE, REMIND, GCAM	Produce the SSP/RCP scenario vocabulary used in IPCC reports.

Paradigm	Tools	Notes
Reduced-complexity climate	FaIR, MAGICC, Hector	Emulate the full ESM/IAM ensemble with a handful of ODEs; the buildable tier.
Earth System Models (ESM/GCM)	CMIP6-class models (NOAA, ECMWF, NASA, etc.)	Supercomputer-scale; inform civic models but do not run on consumer hardware.

Open model repositories and platforms worth knowing: CoMSES / OpenABM (peer-reviewed, DOI-assigned model repository — the highest-yield place to find existing models and practitioners), NetLogo Models Library, Google Earth Engine (planetary-scale remote sensing and environmental data processing).

Institutions

Santa Fe Institute — the complexity science center of gravity. Founded on the premise that economic and ecological systems are complex adaptive systems, not equilibrium machines.

Oxford INET / Macrocosm — J. Doyne Farmer's group built some of the most rigorous agent-based macroeconomic models, then spun out Macrocosm as a company billing itself as an agent-based "digital twin" unifying climate, macro, and geopolitical risk. The closest existing thing to a general-purpose coupled socio-environmental simulator.

PIK — Potsdam Institute for Climate Impact Research — runs several of the major IAMs and does coupled climate-economy scenario work for European policy.

Stockholm Resilience Centre — home of planetary boundaries research and socio-ecological systems (SES) framing.

CASA — Centre for Advanced Spatial Analysis, UCL — Mike Batty's group; the lineage for urban digital twins and geo-coupled simulation.

Google Earth Engine — planetary-scale remote sensing and environmental data processing platform. Arguably the most important infrastructure layer for large-scale environmental civic simulation.

CoMSES / OpenABM — the open research community for computational and agent-based modeling. Peer-reviewed, DOI-assigned model repository.

Several large-scale digital-twin and planetary-data efforts are assembling adjacent pieces: Microsoft Planetary Computer, NASA's digital twins program, ESA/Copernicus, and the WEF/DNV Global Digital Twin ambition.

Notable people

Jay Forrester (system dynamics), Donella Meadows (Thinking in Systems), J. Doyne Farmer (complexity economics), Kate Raworth (doughnut economics), Mike Batty (urban simulation), Stephen Lansing (Balinese water temples), Daniel Christian Wahl (regenerative design), Fei-Fei Li (spatial/physical world models), Bret Victor (explorable explanations; cited in Chapter 8).

TODO

- TODO: Add affiliation/last-verified dates — this list goes stale quickly.
- TODO: Cross-link each tool to the recipe or scenario that uses it (Appendix C, Reference — Scenario Library).

Reference — Media Companion

> The book has a companion media plan: one short video per chapter, published on YouTube, using the chapter headings as video segments. This is a production reference — the media strategy, the visual assets needed, the production approach, and the relationship between the written and video versions of each chapter's argument. (Chapter numbers in the outlines below follow the current structure in `00-chapter-index.md`.)

The companion media strategy

Books are good at depth and at the written argument. Video is good at showing, at emotional resonance, and at reaching audiences who are not book readers. For a book whose thesis is about accessibility and broad participation, video is not optional — it is a core part of the communication strategy.

The plan:

One short video per chapter (target: 5–12 minutes). Each video uses the chapter's main headings as segment markers. The video is not a reading of the chapter but a visual argument for the same ideas — with maps, visualizations, footage, and narration that can stand alone as a self-contained piece.

YouTube as the primary platform — for reach, discoverability, and the ability to embed in the book's web version. The channel should be established early and the first video published before the book is complete, to build an audience.

Cross-links. Each chapter file ends with a link to its companion video (once produced). Each video ends with a reference to the chapter for readers who want depth.

The existing simulate.world talk (https://www.youtube.com/watch?v=ibgt7N is the proof-of-concept. It should be referenced in the book as the author's existing public statement of the thesis.

Visual assets: what exists and what needs to be created

Existing assets (from simulate.world.old)

The original simulate.world site contains a large library of images, organized by topic in `simulate.world.old/data/images.js`. These are mapped to the emotional arc of the site — Earth from space, landscapes, environmental crisis imagery, civic participation — and most are relevant to the book's chapters.

Key existing image categories: - Earth from space (Chapter 1 — Our World) - Landscape and place (Chapter 1 — personal experience) - Environmental crisis imagery (Chapter 2 — emotional arc; Chapter 3 — existential threats) - Case study imagery:

Balinese water temples, California drought, Yellowstone wolves (Chapter 4) - Civic and participation imagery (Chapter 8 — publication layer; Chapter 9 — participation)

Needed assets

Data visualizations — the figures called out as placeholders in each chapter. These need to be created as clean, publication-quality SVG or PNG files. Priority: 1. Planetary boundaries diagram (updated 2023 version) 2. Mass extinction timeline (Phanerozoic eon) 3. Modeling paradigm comparison table 4. Computational hierarchy diagram 5. Interface wireframes for Chapter 8 (the publication layer)

Screen recordings of working civic simulations — when scenarios are built, short screen recordings showing parameter exploration make excellent video content.

Maps — each regional or local scenario needs a base map. OpenStreetMap and Natural Earth provide the data; Mapbox or QGIS provide the rendering.

Video outlines

Preface — The Room (target: 8 minutes)

Segment 1: The room (2 min) - Open with a city council meeting. - Describe the information asymmetry: the consultant has the model; the room has opinions. - Frame the thesis: simulation as civic infrastructure.

Segment 2: The gap (2 min) - Show the existing tools (Vensim, NetLogo, GAMA) and their desktop-bound nature. - Contrast with the publication layer that doesn't exist. - The Bret Victor clip or reference to explorable explanations.

Segment 3: The opportunity (2 min) - AI authorship removing the programming barrier. - Examples of what civic simulation could look like. - The book's arc.

Segment 4: Join us (2 min) - Call to participation. - Link to companion website, GitHub, community.

Chapter 1 — Our World (target: 10 minutes)

Segment 1: The first photographs (3 min) - The 1967 ATS-III full-disk photograph. - Lovelock and Gaia Theory. - Life as a 4-billion-year trajectory.

Segment 2: Individual experience (2 min) - Personal place memories; why we care. - Diversity of experience as data.

Segment 3: Planetary boundaries (3 min) - Stockholm Resilience Centre visualization. - The six crossed boundaries. - Doughnut economics as civic frame.

Segment 4: The feedback dynamics (2 min) - Forest-moisture, ice-albedo, marine carbon pump. - Why simulation earns its keep here.

Chapter 3 — Existential Threats (target: 10 minutes)

Segment 1: A new class of challenge (2 min) - Scale, temporal displacement, nonlinearity. - The hyperobject concept (Morton).

Segment 2: The geological record (3 min) - Mass extinction timeline visualization. - The five events; the Permian-Triassic Great Dying. - The sixth extinction question.

Segment 3: Tipping points (3 min) - Bifurcation diagram. - Amazon, AMOC, permafrost, ice sheets. - The cascade possibility.

Segment 4: Black Swans and honest uncertainty (2 min) - Under a Green Sky / methane scenario. - The range of outcomes matters. - Call to rigorous modeling.

Production notes

Audio. The voice recording can be done in a quiet room with a decent USB condenser microphone (Blue Yeti or equivalent). Good audio matters more than video quality.

Screen recordings. OBS Studio (free, cross-platform) is the standard tool. Record at 1920×1080 minimum.

Graphics. Manim (mathematical animation library, Python) is excellent for the mathematical visualizations (bifurcation diagram, DEM derivatives, ensemble fan charts). For geographic visualizations, QGIS or Mapbox Studio.

Editing. DaVinci Resolve (free tier) is the best free editor. Adobe Premiere or Final Cut for paid options.

Hosting. YouTube as primary; archive on the Internet Archive for permanence.

TODOs

- TODO: Produce the Preface ("The Room") companion video as the first milestone.
- TODO: Create the `assets/` directory with all placeholder figure files.
- TODO: Tag each chapter figure placeholder with the video segment it corresponds to.
- FIGURE: Add figure list here as each chapter's placeholders are resolved.

Reference — Voices and Interviews

> A book about civic simulation should include the voices of the communities and practitioners it is about. This is the placeholder and structure for those voices: researchers, community organizers, indigenous land stewards, civic tech practitioners, and others who have direct experience with the problems this book addresses. As interviews are completed, the material is intended to be woven into the relevant spine chapters; this file holds the targets, templates, and permissions until then. Each interview section includes a permissions metadata block.

About this chapter

This chapter will be built over time as interviews are conducted and permissions are secured. The structure here provides: - A template for interview sections - A list of target interviewees - Permissions metadata format - Placeholder text for each planned interview

---

Interview template

Each interview section follows this structure:

```
# Interview metadata (not published — internal permissions record)
source: [name / organization]
contact: [email or contact method]
permissions:
  granted: [yes / no / pending]
  scope: [full text / excerpt / summary / none]
  date_granted: [YYYY-MM-DD]
  notes: [any restrictions, embargo, anonymity request]
interview_date: [YYYY-MM-DD]
interviewer: [name]
```

Followed by: - A 1–2 paragraph introduction to the subject and why their perspective matters. - The interview transcript or summary (when permissions are granted). - A brief editorial note connecting the interview to the chapter's argument.

Planned interviews

1. Water management / civic modeling practitioner

Who: A planner, hydrologist, or water rights attorney who has worked on a civic watershed decision — ideally one where technical information was contested or inaccessible.

Why: To ground Chapter 4 (case studies) and Chapter 8 (the publication layer) in a practitioner's account of what information asymmetry looks like in practice.

Target questions: - What information did each party in the decision have access to, and what were they missing? - Was there a model involved? Who built it? Was it accessible to the public? - What would have changed if the other party had had access to the same model? - What does "accessible civic simulation" mean to you?

```
source: TBD
permissions:
  granted: pending
```

---

2. Indigenous community land steward

Who: A member of an indigenous community with long-term ecological knowledge of their territory — ideally one that has engaged with or been affected by government environmental modeling.

Why: To address the absent-stakeholder problem concretely, and to explore the relationship between traditional ecological knowledge and computational models. Indigenous communities have been systematically excluded from the decisions that affect their territories; their knowledge has been systematically ignored.

Target questions: - How has your community's knowledge of the land been received or ignored in formal planning processes? - Have government or corporate models been used to make decisions about your territory? Were they accurate? Were they available to your community? - What would a model that genuinely represented your community's knowledge and values look like? - What would you want developers of civic simulation tools to understand?


```yaml
source: TBD
permissions:
  granted: pending
  notes: Approach through community organization or tribal office, not directly.
```


---

3. Civic technology practitioner (participatory democracy)

Who: Someone with direct experience running Pol.is, Decidim, or a Citizens' Assembly process — ideally on an environmental or land-use question.

Why: To connect Chapter 9 (participation and governance) to lived practice, and to understand where the tools work and where they fail.

Target questions: - What is the hardest structural problem in civic deliberation you have encountered? - Did simulation or modeling play any role in the process? Should it have? - What is the relationship between having a shared model and reaching a decision that people accept? - What advice would you give to someone building civic simulation tools for the first time?


```yaml
source: TBD — possible contacts: vTaiwan participants, Decidim Barcelona team, KNOCA
↪  network
permissions:
  granted: pending
```


---

4. Simulation researcher

Who: A researcher at one of the major socio-environmental simulation groups (Oxford INET/Macrocosm, Santa Fe Institute, PIK, Stockholm Resilience Centre, SESYNC) — ideally one who has thought about the gap between their research models and public accessibility.

Why: To connect Chapter 5 (the landscape of simulation) to the research frontier, and to get a practitioner's honest assessment of what civic simulation can and cannot do.

Target questions: - What is the hardest unsolved problem in your modeling work? - Have you ever tried to communicate your model's outputs to a non-specialist audience? What happened? - Is there a version of your model that could be accessible to a community directly affected by the issue? - What do you think about the "illustrative model" tier — models that are not trying to be quantitatively precise but to show structure?

```yaml
source: TBD — possible contacts: J. Doyne Farmer (Macrocosm), researchers at SFI, PIK
permissions:
  granted: pending
```

---

5. Climate scientist / Earth system modeler

Who: A researcher working on IAMs, reduced-complexity models, or climate projections — ideally one who has engaged with IPCC report processes.

Why: To connect Chapters 7–8 (data, uncertainty, implementation) to the state of the art in climate science, and to understand the gap between what GCMs produce and what civic communities need.

Target questions: - What does your model produce that communities could actually use? - What is the right relationship between GCM/ESM outputs and civic-facing simulation tools? - How do you communicate uncertainty in your work? What works and what doesn't? - What would a credible "browser-tier" version of climate scenario analysis look like?

```
source: TBD
permissions:
  granted: pending
```

---

6. Community organizer / environmental justice advocate

Who: An organizer who has worked on environmental issues with communities that have limited access to technical expertise.

Why: To ground the accessibility argument in concrete experience of what "no access to models" means for communities facing industrial siting, pollution, or water rights decisions.

Target questions: - What information did your community need and couldn't get? - Have you ever used a model or simulation in an advocacy campaign? What was the experience? - What would it change if communities like yours had access to the same modeling tools that developers and agencies have?

```yaml
source: TBD – possible contacts: Digital Democracy (digital-democracy.org), local
↪  watershed coalitions
permissions:
  granted: pending
```


7. Marine megafauna biologist — Dr. Andrea Marshall

Who: Dr. Andrea Marshall, founder of the Marine Megafauna Foundation and widely recognized for her foundational research on oceanic manta rays, including the first scientific descriptions of several manta species and the documentation of their population structure, site fidelity, and conservation status. She has satellite-tagged and individually identified thousands of rays across multiple ocean basins, building one of the largest databases of pelagic elasmobranch behavior in existence.

Note on approach: Dr. Marshall is currently recovering from health challenges and should not be approached for a live interview at this time. However, she has an extensive public body of work — published papers, documentary appearances, public lectures, and the ongoing research outputs of the Marine Megafauna Foundation — that can be mined for quotation and context, with appropriate attribution and permissions.

Why this voice matters: Chapter 4 introduces reef connectivity and the dependence of reef health on pelagic megafauna, using manta ray tracking as an example of how a large migratory animal can function as a real-world data collector tracing the invisible structure of an ocean ecosystem. Dr. Marshall's work grounds that argument in decades of field evidence. Her observations on the coupling between open-ocean and reef systems — and on the collapse of that coupling when either component is degraded — are directly relevant to the book's civic simulation argument.

Material to mine: - Marshall, A. D., et al. (2011). "Redescription of the genus Manta with resurrection of Manta alfredi." Zootaxa, 2301, 1–28. - Marshall, A. D., Dudgeon, C. L., & Bennett, M. B. (2011). "Size and structure of a photographically identified population of manta rays Manta alfredi in southern Mozambique." Marine Biology, 158(5), 1111–1124. - Marine Megafauna Foundation public lecture series on oceanic manta behavior and conservation (YouTube, 2019–2024). - MMF project documentation on the Revillagigedo Biosphere Reserve and manta site fidelity at Socorro Island.

Target themes for quotation or synthesis: - What does the spatial behavior of a manta ray reveal about the structure of the ocean ecosystem it inhabits? - How does site fidelity (returning to the same cleaning stations over years and decades) create vulnerability to reef loss? - What does population-level data — rather than individual encounter data — reveal about the health of pelagic-reef coupled systems? - What would a modeling tool look like that could communicate the connectivity argument to reef managers, fishery regulators, and the public?

```
source: Dr. Andrea Marshall / Marine Megafauna Foundation
contact: https://marinemegafauna.org
```

```yaml
permissions:
  granted: pending
  scope: Public work may be cited with attribution; live interview to be arranged
↪  when appropriate
  notes: Do not approach for live interview until health status allows. Use
↪  published work and MMF public materials in the interim.
```

White space for future voices

This chapter will also include: - Voices from communities in the scenarios described in the case studies and scenario library (Bali, Rajasthan, California Central Valley, etc.) - Perspectives from practitioners who have used participatory budgeting or digital democracy tools - International perspectives on civic simulation — how the challenges and tools differ in different regulatory and cultural contexts

---

TODO

- TODO: Begin outreach to target interviewees.
- TODO: Develop consent and permissions process for interview publication.
- TODO: Add interview transcripts as they are completed.
- PLACEHOLDER: Each planned interview above should eventually be replaced with actual interview content.

Appendix A — Notation and Mathematical Conventions

> This appendix collects all mathematical notation used in the book. Symbols are defined here on first global use and at their point of introduction in each chapter. All equations use KaTeX.

General conventions

- Bold lowercase letters denote vectors: $\mathbf{v}$, $\mathbf{x}$.
- Uppercase letters denote matrices or transformations: A, M.
- Scalars are italic lowercase: t, x, r.
- Subscripts denote index or spatial location: $T_{i,j}$ is temperature at grid cell (i, j).
- Superscripts denote time step: $x^{(t)}$ is state at time t.
- The rate of change of a quantity x with respect to time is written $\dot{x}$ or dx/dt.
- Angle brackets $\langle \cdot \rangle$ denote ensemble or spatial averages.

Physical constants and parameters

Symbol	Quantity	Value / Units
I_0	Solar constant	≈ 1361 W/m^2
σ	Stefan-Boltzmann constant	5.67×10^{-8} W m^{-2} K^{-4}
g	Gravitational acceleration	9.81 m/s^2
Γ	Environmental lapse rate	≈ 6.5 °C/km
c_p	Specific heat of air (constant pressure)	1005 J/(kg $\cdot$ K)

Climate and atmosphere

Symbol	Quantity
T	Surface temperature (K or °C)
T_{eq}	Equilibrium surface temperature
θ_z	Solar zenith angle
ϕ	Latitude
δ	Solar declination

Symbol	Quantity
H	Hour angle
τ	Optical depth of atmosphere
α	Surface albedo (dimensionless, 0–1)
F	Radiative forcing (W/m^2)
λ	Climate sensitivity parameter (°C per W/m^2)
ECS	Equilibrium Climate Sensitivity (°C per CO_2 doubling)

Solar zenith angle:

$$\cos(\theta_z) = \sin(\phi)\sin(\delta) + \cos(\phi)\cos(\delta)\cos(H)$$

Solar declination:

$$\delta = -23.45° \cdot \cos\left(\frac{360}{365}(d+10)\right)$$

Direct normal irradiance (clear sky):

$$I_{DNI} = I_0 \cdot e^{-\tau/\cos(\theta_z)}$$

Simple energy balance:

$$C\frac{dT}{dt} = F - \frac{T}{\lambda}$$

where C is the heat capacity of the system.

Hydrology

Symbol	Quantity
z	Elevation (m)
β	Local slope angle
a	Upslope contributing area (m^2)
TWI	Topographic Wetness Index
Q	Flow (m^3/s)
P	Precipitation (mm/day)
ET	Evapotranspiration (mm/day)

Topographic Wetness Index:

$$\text{TWI} = \ln\left(\frac{a}{\tan(\beta)}\right)$$

System dynamics

Symbol	Quantity
S_i	Stock i (level variable)
F_{ij}	Flow from stock i to stock j (rate)

Symbol	Quantity
τ_i	Delay time constant for process i

Stock-and-flow equation:

$$\frac{dS_i}{dt} = \sum_j F_{ji} - \sum_j F_{ij}$$

Tipping points and bifurcation

The canonical fold bifurcation (tipping point model):

$$\frac{dx}{dt} = r - x^2 + \eta(t)$$

where: - x is the system state - r is the slowly-varying forcing parameter - $\eta(t)$ is noise

At $r < 0$: two fixed points (stable and unstable). At $r = 0$: the saddle-node bifurcation. At $r > 0$: no fixed point and rapid transition to a new state.

Critical slowing down near the tipping point: the return rate to equilibrium approaches zero as $r \to 0^-$, meaning the system takes increasingly long to recover from small perturbations. This is an early warning indicator.

Economic modeling

Symbol	Quantity
x_i	Output of sector i
a_{ij}	Technical coefficient (output of j used per unit of i)
A	Technical coefficients matrix
$\mathbf{d}$	Final demand vector
L	Leontief inverse: $(I - A)^{-1}$

Leontief I-O model:

$$\mathbf{x} = L\mathbf{d} = (I - A)^{-1}\mathbf{d}$$

A shock $\Delta\mathbf{d}$ propagates as $\Delta\mathbf{x} = L\,\Delta\mathbf{d}$.

Probability and uncertainty

Symbol	Quantity
$P(A)$	Probability of event A
$\mathbb{E}[X]$	Expected value of random variable X
$\mathrm{Var}(X)$	Variance of X
σ_X	Standard deviation of X

Symbol	Quantity
$\pi(\theta \mid \text{data})$	Posterior distribution over parameters θ

TODO

- TODO: Add notation for agent-based model quantities (agent state, neighborhood, transition rule).
- TODO: Add notation for the FaIR reduced-complexity climate model (carbon cycle, thermal response).
- TODO: Add notation for the Doughnut / planetary boundary indicators as they are defined in the chapters.

Appendix B — Data Sources and Code Links

> A comprehensive reference list of data sources cited or used in the book, organized by domain. Each entry includes the source name, URL, access model (free / subscription / registration), and the chapter(s) where it is used.

Trade and economy

Source	URL	Access	Used in
UN Comtrade	comtradeplus.un.org	Free tier	Ch. 6, Scenario Library
Observatory of Economic Complexity (OEC)	oec.world	Free	Ch. 6, Scenario Library

Source	URL	Access	Used in
EXIOBASE	exiobase.eu	Free (registration)	Ch. 6, Scenario Library
GTAP Database	gtap.agecon.purdue.edu	Subscription	Ch. 5, Ch. 6
OECD I-O tables	stats.oecd.org/Index.aspx?DataSetCode=IOTS_2021	Free	Ch. 6

Energy

Source	URL	Access	Used in
IEA World Energy Statistics	iea.org/data-and-statistics	Free tier / subscription	Ch. 6, Scenario Library
EIA International Energy Data	eia.gov/international/data/world	Free	Ch. 6, Scenario Library
IEA Energy Atlas	energyatlas.iea.org	Free	Scenario Library

Food and agriculture

Source	URL	Access	Used in
FAOSTAT	fao.org/faostat	Free	Ch. 6, Scenario Library
USDA ERS	ers.usda.gov	Free	Ch. 4
Global Agri-Footprint	globagri.org	Free	Ch. 6

Minerals and materials

Source	URL	Access	Used in
USGS Mineral Commodity Summaries	usgs.gov/centers/national-minerals-information-center	Free	Ch. 6, Scenario Library
USGS MRDS	mrdata.usgs.gov	Free	Ch. 6

Climate and atmosphere

Source	URL	Access	Used in
ERA5 / Copernicus CDS	cds.climate.copernicus.eu	Free (registration)	Ch. 6, Ch. 7, Scenario Library
NASA Earthdata	earthdata.nasa.gov	Free (registration)	Ch. 6
NOAA NCEI	ncei.noaa.gov	Free	Ch. 6
CMIP6 / ESGF	esgf-node.llnl.gov	Free (registration)	Ch. 5, Ch. 7
IPCC Atlas	interactive-atlas.ipcc.ch	Free	Scenario Library
FaIR model	github.com/OMS-NetZero/FAIR	Open source	Ch. 7
MAGICC	magicc.org	Free (registration)	Ch. 7

Terrain and geography

Source	URL	Access	Used in
SRTM DEM (30m global)	earthexplorer.usgs.gov	Free	Ch. 7, Scenario Library

Source	URL	Access	Used in
Copernicus DEM (30m global)	spacedata.copernicus.eu	Free	Ch. 7
USGS National Elevation Dataset (3m US)	nationalmap.gov	Free	Ch. 7
GEBCO (ocean bathymetry)	gebco.net	Free	Ch. 6
OpenStreetMap	openstreetmap.org	Free (ODbL)	Ch. 6, Ch. 7
Natural Earth	naturalearthdata.com	Free	Ch. 6, Ch. 7

Ocean and marine

Source	URL	Access	Used in
NOAA Coral Reef Watch	coralreefwatch.noaa.gov	Free	Scenario Library
Global Fishing Watch	globalfishingwatch.org/data-download	Free	Ch. 6

Source	URL	Access	Used in
Copernicus Marine Service	marine.copernicus.eu	Free	Ch. 7
GEBCO	gebco.net	Free	Ch. 6

Land cover and vegetation

Source	URL	Access	Used in
NLCD (National Land Cover Database, US)	mrlc.gov	Free	Ch. 7, Scenario Library
Copernicus Global Land Cover	land.copernicus.eu/global	Free	Ch. 7
Global Forest Watch	globalforestwatch.org	Free	Ch. 6, Ch. 9
LANDFIRE (fuel and vegetation, US)	landfire.gov	Free	Ch. 7, Scenario Library

Source	URL	Access	Used in
Hansen Global Forest Change	earthenginepartners.appspot.com/science-2013-global-forest	Free	Scenario Library

---

Population and society

Source	URL	Access	Use
WorldPop	worldpop.org	Free	Ch.
SEDAC	sedac.ciesin.columbia.edu	Free	Ch.
US Census	data.census.gov	Free	Ch.
UN World Population Prospects	population.un.org/wpp	Free	Ch.

Biodiversity

Source	URL	Access	Used in
IUCN Red List	iucnredlist.org	Free	Ch. 3
iNaturalist	inaturalist.org	Free	Ch. 9
Living Planet Index (WWF)	livingplanetindex.org	Free	Ch. 3
GBIF	gbif.org	Free	Scenario Lib

Source	URL	Access	Used in

Code repositories

Repository	Description	URL
FaIR (Finite Amplitude Impulse Response)	Reduced-complexity climate model	github.com/OMS-NetZero/FAIR
PySD	Runs Vensim/Stella models in Python	pysd.readthedocs.io
Mesa	Agent-based modeling in Python	mesa.readthedocs.io
Agents.jl	ABM in Julia	juliadynamics.github.io/Agents.jl
EMA Workbench	Exploratory modeling under deep uncertainty	emaworkbench.readthedocs.io
NetLogo	Classic ABM environment	ccl.northwestern.edu/netlogo
simulate.world (this project)	Civic simulation platform	github.com/TODO

Source material from this project

The following source files in this repository contain research notes and reference material organized during this project's development:

File	Content
`wiki/theory/20260501-sims-and-tools.md`	Comprehensive landscape of simulation tools, data sources, institutions, notable people, books
`wiki/theory/20260503-participation-and-governance.md`	Participatory democracy, civic tech, computational democracy, regenerative finance
`wiki/technical/20260502-sims-diy-criticisms.md`	Honest critique of building new simulation infrastructure
`wiki/technical/20260518-model-parts.md`	Physical forces taxonomy: geography, climate, hydrology, biosphere, humanosphere
`wiki/technical/20260531-interface-goals.md`	Interface architecture notes, Jupyter/Observable comparison, AI authorship
`wiki/technical/20260531-scenarios-to-model.md`	Scenario catalogue with buildability criteria

File	Content
`simulate.world.old/data/*.js`	Original simulate.world content: thesis, our world, existential threats, digital models, examples, media messages, superpowers
`orbitalbiz/bizlog/20260604-potential-sims.md`	Extended scenario ideas with news links
`incomplete-thoughts.md`	Working notes on the book's framing and argument structure

---

TODO

- TODO: Add DOI links for all academic references.
- TODO: Verify all URLs are current (last checked: 2026-06-04).
- TODO: Add API access notes for sources that have programmatic access.
- TODO: Add example code snippets for loading key datasets in Python.

Appendix C — Implementation Recipes

> Minimal runnable examples for the key implementation patterns described in the book. Each recipe is self-contained, uses open-source tools, and is designed to be a starting point — not a complete implementation. Language: Python unless otherwise noted.

Recipe 1: Analytic solar irradiance (Python)

Compute direct normal irradiance at a given location and time without any external data.

```
import math
from datetime import datetime, timezone


def solar_declination(day_of_year: int) -> float:
    """Solar declination in degrees."""
    return -23.45 * math.cos(math.radians(360 / 365 * (day_of_year + 10)))


def solar_zenith_angle(lat: float, lon: float, dt: datetime) -> float:
    """
```

```
    Solar zenith angle in degrees.
    lat, lon in degrees. dt in UTC.
    """
    day_of_year = dt.timetuple().tm_yday
    hour = dt.hour + dt.minute / 60 + dt.second / 3600
    delta = solar_declination(day_of_year)
    # Hour angle: 0 at solar noon; solar noon at lon/15 hours from UTC midnight
    solar_noon = 12 - lon / 15  # rough approximation (ignores equation of time)
    H = 15 * (hour - solar_noon)  # degrees
    lat_r = math.radians(lat)
    delta_r = math.radians(delta)
    H_r = math.radians(H)
    cos_zenith = (math.sin(lat_r) * math.sin(delta_r)
                  + math.cos(lat_r) * math.cos(delta_r) * math.cos(H_r))
    return math.degrees(math.acos(max(-1.0, min(1.0, cos_zenith))))

def clear_sky_dni(zenith_deg: float, tau: float = 0.3) -> float:
    """
    Direct Normal Irradiance (W/m$^2$) under clear sky.
    tau: atmospheric optical depth (~0.1 very clear, ~0.3 typical, ~0.5 hazy)
    Returns 0 if below horizon.
    """
    if zenith_deg >= 90:
        return 0.0
    I0 = 1361.0  # solar constant, W/m$^2$
    cos_z = math.cos(math.radians(zenith_deg))
    return I0 * math.exp(-tau / cos_z)

# Example: Oakland, California, summer solstice at noon UTC
dt = datetime(2026, 6, 21, 19, 0, tzinfo=timezone.utc)  # 19:00 UTC □ noon Pacific
lat, lon = 37.8, -122.3
zenith = solar_zenith_angle(lat, lon, dt)
```

```
dni = clear_sky_dni(zenith)
print(f"Solar zenith: {zenith:.1f}°")
print(f"Direct Normal Irradiance: {dni:.0f} W/m$^2$")
```

Recipe 2: DEM-derived flow direction (Python / NumPy)

Compute D8 flow direction from a digital elevation model raster.

```
import numpy as np

# D8 neighbor offsets (row_delta, col_delta) for 8 directions
D8_DIRS = [(-1, 0), (-1, 1), (0, 1), (1, 1),
           (1, 0), (1, -1), (0, -1), (-1, -1)]
D8_CODES = [64, 128, 1, 2, 4, 8, 16, 32]  # ArcGIS-style codes

def flow_direction_d8(dem: np.ndarray) -> np.ndarray:
    """
    Compute D8 flow direction from a DEM.
    Returns array of D8 codes (1/2/4/8/16/32/64/128).
    Cells with no downslope neighbor get code 0.
    """
    rows, cols = dem.shape
    fdir = np.zeros_like(dem, dtype=np.int32)
    for r in range(rows):
        for c in range(cols):
            max_drop = 0.0
            best_code = 0
            for (dr, dc), code in zip(D8_DIRS, D8_CODES):
                nr, nc = r + dr, c + dc
                if 0 <= nr < rows and 0 <= nc < cols:
```

```
                    # Diagonal distance = sqrt(2) cell widths
                    dist = math.sqrt(2) if (dr != 0 and dc != 0) else 1.0
                    drop = (dem[r, c] - dem[nr, nc]) / dist
                    if drop > max_drop:
                        max_drop = drop
                        best_code = code
            fdir[r, c] = best_code
    return fdir

def flow_accumulation(fdir: np.ndarray) -> np.ndarray:
    """
    Compute flow accumulation from D8 flow direction.
    Returns array of counts of upstream cells draining through each cell.
    """
    rows, cols = fdir.shape
    accum = np.ones((rows, cols), dtype=np.int64)
    # Build adjacency: which cell does each cell flow to?
    # Simple iterative approach (not optimal for large DEMs — use topological sort
    ↪ for production)
    dir_to_delta = dict(zip(D8_CODES, D8_DIRS))
    for _ in range(rows * cols):  # iterate until stable
        changed = False
        for r in range(rows):
            for c in range(cols):
                code = fdir[r, c]
                if code in dir_to_delta:
                    dr, dc = dir_to_delta[code]
                    nr, nc = r + dr, c + dc
                    if 0 <= nr < rows and 0 <= nc < cols:
                        # This is approximate; proper implementation uses topological
                        ↪ order
                        pass
```

```
    # Note: a proper implementation requires topological sorting of the DEM.
    # See richdem or pysheds libraries for production-quality implementations.
    return accum

# Better approach: use pysheds (pip install pysheds)
# from pysheds.grid import Grid
# grid = Grid.from_raster('dem.tif')
# dem = grid.read_raster('dem.tif')
# pit_filled = grid.fill_pits(dem)
# flooded = grid.fill_depressions(pit_filled)
# inflated = grid.resolve_flats(flooded)
# fdir = grid.flowdir(inflated)
# acc = grid.accumulation(fdir)
```

Recipe 3: Minimal FaIR climate model (Python)

Run a simple emissions scenario through the FaIR reduced-complexity climate model.

```
pip install fair
```

```
import numpy as np
import matplotlib.pyplot as plt
from fair import FAIR
from fair.interface import fill, initialise
from fair.io import read_properties

# Create a minimal FAIR instance
f = FAIR(ch4_method="Thornhill2021")
f.define_time(1750, 2100, 1)
```

```
f.define_scenarios(["ssp245", "ssp585"])
f.define_configs(["default"])

species, properties = read_properties()
f.define_species(species, properties)
f.allocate()

# Fill in SSP emissions from included datasets
from fair.forcing.ghg import meinshausen2020
from fair.energy_balance_model import EnergyBalanceModel
# (fill emissions, forcing, climate configs here - see fair documentation)

# A minimal working example that does not require external data:
# Model a CO2 pulse and show the temperature response
years = np.arange(2000, 2101)
co2_concentrations = 280 + 2.5 * (years - 2000)  # simplified ramp scenario
# Temperature change from simple energy balance:
ECS = 3.0  # equilibrium climate sensitivity, °C per doubling
F = 5.35 * np.log(co2_concentrations / 280)  # simplified radiative forcing
lambda_param = ECS / (5.35 * np.log(2))
T_eq = lambda_param * F
print("Year 2100 equilibrium warming:", round(T_eq[-1], 2), "°C")

# For a full FAIR run with SSP scenarios, see:
# https://docs.fairmodel.net/en/latest/
```

Recipe 4: Simple system dynamics in Python

A stock-and-flow model of aquifer depletion.

```
import numpy as np
import matplotlib.pyplot as plt

def aquifer_model(
    years: int = 50,
    initial_volume_km3: float = 1000.0,  # total aquifer storage
    annual_recharge_km3: float = 5.0,    # natural recharge rate
    initial_extraction_km3: float = 15.0, # current pumping rate
    extraction_growth_rate: float = 0.02, # annual growth in extraction
    critical_fraction: float = 0.3,       # fraction at which collapse occurs
) -> dict:
    """
    Simple stock-flow model of groundwater depletion.
    When the aquifer reaches the critical fraction, extraction pressure
    increases as farmers drill deeper (positive feedback loop).
    """
    volume = [initial_volume_km3]
    extraction = [initial_extraction_km3]
    times = list(range(years + 1))

    for t in range(years):
        current_vol = volume[-1]
        current_ext = extraction[-1]

        # Positive feedback: as volume drops below critical threshold,
        # economic pressure increases extraction
        if current_vol / initial_volume_km3 < critical_fraction:
            # Farmers dig deeper to maintain yield — extraction increases faster
            pressure_multiplier = 1 + 0.05 * (critical_fraction - current_vol /
↪ initial_volume_km3)
            next_ext = current_ext * (1 + extraction_growth_rate *
↪ pressure_multiplier)
```

```
        else:
            next_ext = current_ext * (1 + extraction_growth_rate)

        net_change = annual_recharge_km3 - next_ext
        next_vol = max(0, current_vol + net_change)

        volume.append(next_vol)
        extraction.append(next_ext)

    return {
        "times": times,
        "volume": volume,
        "extraction": extraction,
        "depletion_year": next(
            (t for t, v in zip(times, volume) if v <= 0),
            None
        )
    }

# Run two scenarios
baseline = aquifer_model(extraction_growth_rate=0.02)
conservation = aquifer_model(extraction_growth_rate=-0.01)  # decreasing extraction

plt.figure(figsize=(10, 4))
plt.subplot(1, 2, 1)
plt.plot(baseline["times"], baseline["volume"], label="Baseline (2% growth)")
plt.plot(conservation["times"], conservation["volume"], label="Conservation
↪  (-1%/yr)")
plt.xlabel("Years from now")
plt.ylabel("Aquifer volume (km³)")
plt.title("Aquifer depletion scenarios")
plt.legend()
```

```
plt.subplot(1, 2, 2)
plt.plot(baseline["times"], baseline["extraction"], label="Baseline")
plt.plot(conservation["times"], conservation["extraction"], label="Conservation")
plt.xlabel("Years from now")
plt.ylabel("Annual extraction (km³)")
plt.title("Extraction rate over time")
plt.legend()
plt.tight_layout()
plt.savefig("aquifer_model.png", dpi=150)
print(f"Baseline depletion year: {baseline['depletion_year']}")
print(f"Conservation depletion year: {conservation['depletion_year']}")
```

Recipe 5: Input-output shock propagation (Python)

Trace a supply chain disruption through an I-O matrix.

```
import numpy as np

# Toy 5-sector economy
# Sectors: 0=Agriculture, 1=Energy, 2=Manufacturing, 3=Services, 4=Transport
sectors = ["Agriculture", "Energy", "Manufacturing", "Services", "Transport"]

# Technical coefficients matrix A (fraction of sector i output used by sector j)
# A[i, j] = amount of sector i needed to produce one unit of sector j
A = np.array([
    [0.10, 0.02, 0.05, 0.01, 0.03],  # Agriculture inputs
    [0.05, 0.15, 0.10, 0.02, 0.20],  # Energy inputs
    [0.08, 0.12, 0.20, 0.05, 0.15],  # Manufacturing inputs
    [0.03, 0.04, 0.06, 0.10, 0.05],  # Services inputs
```

```
    [0.04, 0.10, 0.08, 0.03, 0.05],  # Transport inputs
])

# Final demand vector (arbitrary units)
d = np.array([100, 80, 120, 200, 60])

# Leontief inverse: L = (I - A)^{-1}
I = np.eye(len(sectors))
L = np.linalg.inv(I - A)

# Total output to meet final demand
x = L @ d

print("Total output required by sector:")
for s, xi in zip(sectors, x):
    print(f"  {s}: {xi:.1f}")

# Simulate an energy supply shock: 30% reduction in energy output
shock = np.zeros(len(sectors))
shock[1] = -0.30 * x[1]  # 30% reduction in energy sector output

# Propagate through the economy
delta_x = L @ shock
print("\nImpact of 30% energy supply shock:")
for s, di in zip(sectors, delta_x):
    print(f"  {s}: {di:+.1f} ({100*di/x[sectors.index(s)]:.1f}%)")
```

Recipe 6: Minimal civic simulation server (Node.js / socket.io)

A minimal server that runs a simulation and streams results to browser clients.

```
// server.js
const http = require('http');
const socketio = require('socket.io');

const server = http.createServer();
const io = socketio(server, { cors: { origin: '*' } });

// Simple aquifer simulation state
function createSimulation(params) {
  return {
    volume: params.initialVolume ?? 1000,
    recharge: params.recharge ?? 5,
    extraction: params.extraction ?? 15,
    year: 0,
    tick() {
      this.extraction *= (1 + (params.extractionGrowth ?? 0.02));
      this.volume = Math.max(0, this.volume + this.recharge - this.extraction);
      this.year++;
      return { year: this.year, volume: this.volume, extraction: this.extraction };
    }
  };
}

io.on('connection', (socket) => {
  console.log('Client connected:', socket.id);
  let sim = null;
```

```
  let interval = null;

  socket.on('start', (params) => {
    if (interval) clearInterval(interval);
    sim = createSimulation(params);
    interval = setInterval(() => {
      const state = sim.tick();
      socket.emit('tick', state);
      if (state.volume <= 0 || state.year >= 100) {
        clearInterval(interval);
        socket.emit('done', state);
      }
    }, 100); // 100ms between ticks = ~10 years/second
  });

  socket.on('disconnect', () => {
    if (interval) clearInterval(interval);
  });
});

server.listen(3000, () => console.log('Simulation server on port 3000'));
```

```
<!-- client.html -->
<!DOCTYPE html>
<html>
<head><title>Aquifer Simulation</title></head>
<body>
<h2>Aquifer Volume Over Time</h2>
<canvas id="chart" width="600" height="300"></canvas>
<br>
```

```
<label>Extraction growth (%/yr): <input id="growth" type="range" min="-5" max="5"
↪  value="2"></label>
<button onclick="startSim()">Run</button>
<script src="/socket.io/socket.io.js"></script>
<script>
const socket = io('http://localhost:3000');
const canvas = document.getElementById('chart');
const ctx = canvas.getContext('2d');
const data = [];

socket.on('tick', (state) => {
  data.push(state);
  drawChart();
});

function drawChart() {
  ctx.clearRect(0, 0, canvas.width, canvas.height);
  if (data.length < 2) return;
  ctx.beginPath();
  data.forEach((d, i) => {
    const x = i / data.length * canvas.width;
    const y = canvas.height - (d.volume / 1000) * canvas.height;
    i === 0 ? ctx.moveTo(x, y) : ctx.lineTo(x, y);
  });
  ctx.strokeStyle = '#2980b9';
  ctx.lineWidth = 2;
  ctx.stroke();
}

function startSim() {
  data.length = 0;
  const growth = parseFloat(document.getElementById('growth').value) / 100;
```

```
  socket.emit('start', { initialVolume: 1000, recharge: 5, extraction: 15,
↪  extractionGrowth: growth });
}
</script>
</body>
</html>
```

TODO

- TODO: Add recipe for wildfire spread (cellular automaton + LANDFIRE fuel data).
- TODO: Add recipe for population dynamics (Lotka-Volterra predator-prey).
- TODO: Add recipe for I-O model with EXIOBASE data loading.
- TODO: Add Docker compose file for running the full simulation stack locally.
- TODO: Add recipe for reading ERA5 data from the Copernicus CDS API.

References

Armstrong McKay, D. I. et al. 2022. "Exceeding 1.5°C Global Warming Could Trigger Multiple Climate Tipping Points." Science 377 (6611).

Beer, S. 1972. Brain of the Firm. Allen Lane.

Beschta, R. L. et al. 2026. "Does Restoring Apex Predators to Food Webs Restore Ecosystems? Comment." Ecological Monographs. https://doi.org/10.1002/ecm.70050.

Beschta, R. L., and W. J. Ripple. 2016. "Riparian Vegetation Recovery in Yellowstone: The First Two Decades After Wolf Reintroduction." Biological Conservation 198: 93–103.

Boers, N. 2021. "Observation-Based Early-Warning Signals for a Collapse of the Atlantic Meridional Overturning Circulation." Nature Climate Change 11: 680–88.

Bostock, M. 2018. "A Better Way to Code." Observable blog.

Bratton, B. H. n.d. "A General Theory of Simulation." Antikythera. https://theoryofsim.antikythera.org/.

Bringhurst, R. 1992. The Elements of Typographic Style. Hartley & Marks.

Carroll, S. R. et al. 2020. "The CARE Principles for Indigenous Data Governance." Data Science Journal 19 (1): 43.

Carroll, S. R., E. Herczog, M. Hudson, K. Russell, and S. Stall.

2021. "Operationalizing the CARE and FAIR Principles for Indigenous Data Futures." Scientific Data 8: 108.

Clayton, S. et al. 2017. "Mental Health and Our Changing Climate." American Psychological Association; ecoAmerica.

Climate and Traditional Knowledges Workgroup. 2014. "Guidelines for Considering Traditional Knowledges in Climate Change Initiatives."

Darling, E. S. et al. 2019. "Social–Environmental Drivers Inform Strategic Management of Coral Reefs in the Anthropocene." Nature Ecology & Evolution 3: 1341–50.

David-Chavez, D. M., and M. C. Gavin. 2018. "A Global Assessment of Indigenous Community Engagement in Climate Research." Environmental Research Letters 13 (12): 123005.

EcoTippingPoints. n.d. "India — Rajasthan: Rainwater Harvest Restoration." ecotippingpoints.org. https://www.ecotippingpoints.org/our-stories/indepth/india-rajasthan-rainwater-harvest-restoration-groundwater-johad.html.

Epstein, J. M., and R. Axtell. 1996. Growing Artificial Societies. MIT Press.

Forrester, J. W. 1971. World Dynamics. Wright-Allen Press.

Foulks, E. F. 1989. "Misalliances in the Barrow Alcohol Study." American Indian and Alaska Native Mental Health Research 2 (3).

Franks, D. M. et al. 2010. "Closing the Loop on the Resource Cycle." Journal of Industrial Ecology.

Ha, D., and J. Schmidhuber. 2018. "World Models." https://arxiv.org/abs/1803.10122.

Harmon, A. 2010. “Indian Tribe Wins Fight to Limit Research of Its DNA.” The New York Times.

Hickman, C. et al. 2021. “Climate Anxiety in Children and Young People and Their Beliefs about Government Responses to Climate Change: A Global Survey.” The Lancet Planetary Health 5 (12): e863–73.

Hobbs, N. T. et al. 2026. “Does Restoring Apex Predators to Food Webs Restore Ecosystems? Reply.” Ecological Monographs. https://doi.org/10.1002/ecm.70049.

Hobbs, N. T., D. B. Johnston, K. N. Marshall, E. C. Wolf, and D. J. Cooper. 2024. “Does Restoring Apex Predators to Food Webs Restore Ecosystems? Large Carnivores in Yellowstone as a Model System.” Ecological Monographs 94 (2): e1598. https://doi.org/10.1002/ecm.1598.

Hughes, T. P. et al. 2017. “Global Warming and Recurrent Mass Bleaching of Corals.” Nature 543: 373–77.

——— et al. 2019. “Global Warming Impairs Stock–Recruitment Dynamics of Corals.” Nature 568: 387–90.

IPCC. 2021. “Summary for Policymakers.” In Climate Change 2021: The Physical Science Basis.

Kennedy, M., and A. O’Hagan. 2001. “Bayesian Calibration of Computer Models.” Journal of the Royal Statistical Society B 63 (3): 425–64.

Kukutai, T., and J. Taylor, eds. 2016. Indigenous Data Sovereignty: Toward an Agenda. ANU Press.

Kwakkel, J. H. 2017. “The Exploratory Modeling Workbench.” Environmental Modelling & Software 96: 239–50.

Lansing, J. S. 1991. Priests and Programmers: Technologies of Power in the Engineered Landscape of Bali. Princeton University Press.

Lansing, J. S., and J. N. Kremer. 1993. "Emergent Properties of Balinese Water Temple Networks." American Anthropologist 95 (1): 97–114.

Lenton, T. M. et al. 2008. "Tipping Elements in the Earth's Climate System." PNAS 105 (6): 1786–93.

Lenton, T. M. 2016. Earth System Science: A Very Short Introduction. Oxford University Press.

Lovelock, J. E. 1979. Gaia: A New Look at Life on Earth. Oxford University Press.

Lovelock, J. E., and L. Margulis. 1974. "Atmospheric Homeostasis by and for the Biosphere." Tellus 26 (1–2): 2–10.

Marshall, A. D. et al. 2011. "Redescription of the Genus Manta with Resurrection of Manta alfredi." Zootaxa 2301: 1–28.

Marshall, A. D., C. L. Dudgeon, and M. B. Bennett. 2011. "Size and Structure of a Photographically Identified Population of Manta Rays Manta alfredi in Southern Mozambique." Marine Biology 158 (5): 1111–24.

Marshall, K. N., N. T. Hobbs, and D. J. Cooper. 2013. "Stream Hydrology Limits Recovery of Riparian Ecosystems After Wolf Reintroduction." Proceedings of the Royal Society B 280 (1756): 20122977. https://doi.org/10.1098/rspb.2012.2977.

Meadows, D. H. 2008. Thinking in Systems: A Primer. Chelsea Green Publishing.

Meadows, D. H., D. L. Meadows, J. Randers, and W. W. Behrens. 1972. The Limits to Growth. Universe Books.

Medina, E. 2011. Cybernetic Revolutionaries. MIT Press.

Morton, T. 2013. Hyperobjects: Philosophy and Ecology After the End of the World. University of Minnesota Press.

Norman, D. A. 1988. The Design of Everyday Things. Basic Books.

Ostrom, E. 1990. Governing the Commons. Cambridge University Press.

Pihkala, P. 2022. "Toward a Taxonomy of Climate Emotions." Frontiers in Climate 3.

Quaempts, E. J., K. L. Jones, S. J. O'Daniel, and T. J. Beechie. 2018. "Aligning Environmental Management with Ecosystem Resilience: A First Foods Example from the Confederated Tribes of the Umatilla Indian Reservation, Oregon, USA." Ecology and Society 23 (2): 29.

Raup, D. M., and J. J. Sepkoski. 1982. "Mass Extinctions in the Marine Fossil Record." Science 215 (4539): 1501–3.

Richardson, K. et al. 2023. "Earth Beyond Six of Nine Planetary Boundaries." Science Advances 9 (37).

Ripple, W. J., and R. L. Beschta. 2012. "Trophic Cascades in Yellowstone: The First 15 Years After Wolf Reintroduction." Biological Conservation 145 (1): 205–13.

Rockström, J. et al. 2009. "A Safe Operating Space for Humanity." Nature 461: 472–75.

Saltelli, A. et al. 2020. "Five Ways to Ensure That Models Serve Society: A Manifesto." Nature 582: 482–84.

Simpson, A. 2007. "On Ethnographic Refusal: Indigeneity, 'Voice' and Colonial Citizenship." Junctures 9: 67–80.

———. 2014. Mohawk Interruptus: Political Life Across the Borders of Settler States. Duke University Press.

Smith, C. J. et al. 2018. "FAIR V1.3: A Simple Emissions-Based Impulse Response and Carbon Cycle Model." Geoscientific Model Development 11 (6): 2273–97.

Smith, L. T. 1999. Decolonizing Methodologies: Research and Indigenous Peoples. Zed Books.

Snipp, C. M. 2016. "What Does Data Sovereignty Imply: What Does It Look Like?" In Indigenous Data Sovereignty: Toward an Agenda, edited by T. Kukutai and J. Taylor. ANU Press.

Solnit, R. 2016. Hope in the Dark. Haymarket Books.

Taylor, L., L. Floridi, and B. van der Sloot, eds. 2017. Group Privacy: New Challenges of Data Technologies. Springer.

Toffler, A. 1970. Future Shock. Random House.

Tsosie, K. S., J. M. Yracheta, and D. Dickenson. 2019. "Overvaluing Individual Consent Ignores Risks to Tribal Participants." Nature Reviews Genetics 20: 497–98.

Tuck, E., and K. W. Yang. 2014. "R-Words: Refusing Research." In Humanizing Research, edited by D. Paris and M. T. Winn. SAGE.

USDA Economic Research Service. 2009. "California Drought and Water Supply." ERS Report. USDA Economic Research Service.

Victor, B. 2011. "Explorable Explanations." worrydream.com. http://worrydream.com/ExplorableExplanations/.

———. 2013. "Media for Thinking the Unthinkable." MIT Media Lab.

Vidiella, B., S. Duran-Nebreda, and S. Valverde. 2026. "Hypergraphs in Ecology Link Non-Equilibrium Dynamics and Conservation." Communications Physics 9: 198. https://doi.org/10.1038/s42005-026-02707-3.

Walter, M. 2016. "Data Politics and Indigenous Representation in Australian Statistics." In Indigenous Data Sovereignty: Toward an Agenda, edited by T. Kukutai and J. Taylor, 79–97. ANU Press.

Walter, M., and C. Andersen. 2013. Indigenous Statistics: A Quantitative Research Methodology. Left Coast Press.

Wang, L., and H. Liu. 2006. "An Efficient Method for Identifying and Filling Surface Depressions in Digital Elevation Models for Hydrologic Analysis." International Journal of Geographical Information Science 20 (2): 193–213.

Ward, P. D. 2007. Under a Green Sky: Global Warming, the Mass Extinctions of the Past, and What They Can Tell Us about Our Future. Harper.

Westen, R. M. van et al. 2024. "Physics-Based Early Warning Signal Shows That AMOC Is on Tipping Course." Science Advances 10 (6).

Weyl, E. G., and E. A. Posner. 2018. Radical Markets. Princeton University Press.

Wiener, N. 1948. Cybernetics: Or Control and Communication in the Animal and the Machine. MIT Press.

Williamson, B., S. Provost, and C. Price. 2023. "Operationalising Indigenous Data Sovereignty in Environmental Research and Governance." Environment and Planning F.

World Labs. 2024. "Spatial Intelligence and Large World Models." https://www.worldlabs.ai.

WWF. 2024. "Living Planet Report 2024." World Wildlife Fund.

Zalasiewicz, J. et al. 2010. "The New World of the Anthropocene." Environmental Science & Technology 44 (7): 2228–31.

Zhang, J., C. Deng, L. Yang, S. Zhang, D. Wang, J. Gao, and Z. Gao. 2026. "A Scalable and Generic Framework for City-Wide Traffic Prediction with Large Language Model." Nature Communications. https://doi.org/10.1038/s41467-026-73610-2.

Acknowledgments

No book about working together gets written alone, and this one has more debts than most.

Howard Silverman of EcoTrust has been an early supporter of this work and a patient guide to the civic and bioregional networks of the Pacific Northwest — the people who actually convene communities, as opposed to the people who write about convening them.

Tish Shute has helped shape this work in ways the project's own notes failed to record, probably because we spent so many hours on the phone - which is its own small lesson about credit.

Anastasiya Soheil al-Fard built alongside me: engineering, design, and deployment, with a care for trustworthy public-facing systems that this project inherits directly.

Shanee Stopnitzky of Emergent Reefs. The reef restoration work at emergentreefs.org is a startling reminder that the gap between modeling an ecosystem and tending one is closed by people, not software.

SuAnne Fu, Professor Emerita of SCAD and the Georgia Economic Collective, offered encouragement and doors at a moment when both mattered.

Havi Hoffman edited none of these pages and shaped all of them. Whatever is snappy, firm, and full of flavor in this book's prose traces back to her editing at Mozilla Hacks; whatever is not is my own backsliding.

The Long Now Foundation collaborated on the New California Water Atlas, where some of the convictions in this book — that public data wants public form — were first tested in practice - and which I was proud to contribute to.

Futurewei under the grants program managed by Yue Chen supported two years of the underlying engine and application work with a generosity that asked good questions and imposed few answers.

This book also stands on open infrastructure and thinking instigated by a wider community of torch bearers: the data commons of NASA, NOAA, ESA, and the IPCC; OpenStreetMap and its mappers; my early 'geowanking' community, Ben Russell, Steve Coast, Schulyer Earle, Rich Gibson, Amber Case, Joshua Schachter, Tim O'Reilly, FooCamp, WhereCamp attendees, Jeremy Irish, Chacha Sikes, Ben Ceverny, Ross Mayfield, Paige Saez, the open-source simulation, geospatial, and web communities whose philosophy and tools are load-bearing on every page. I'm sure I've left out many voices and for that I apologize; the world would not be the same without you.