ESSAY

The First Defender

Four Billion Years of Evolution Produced One Species Capable of Saving the Rest

May 9, 2026

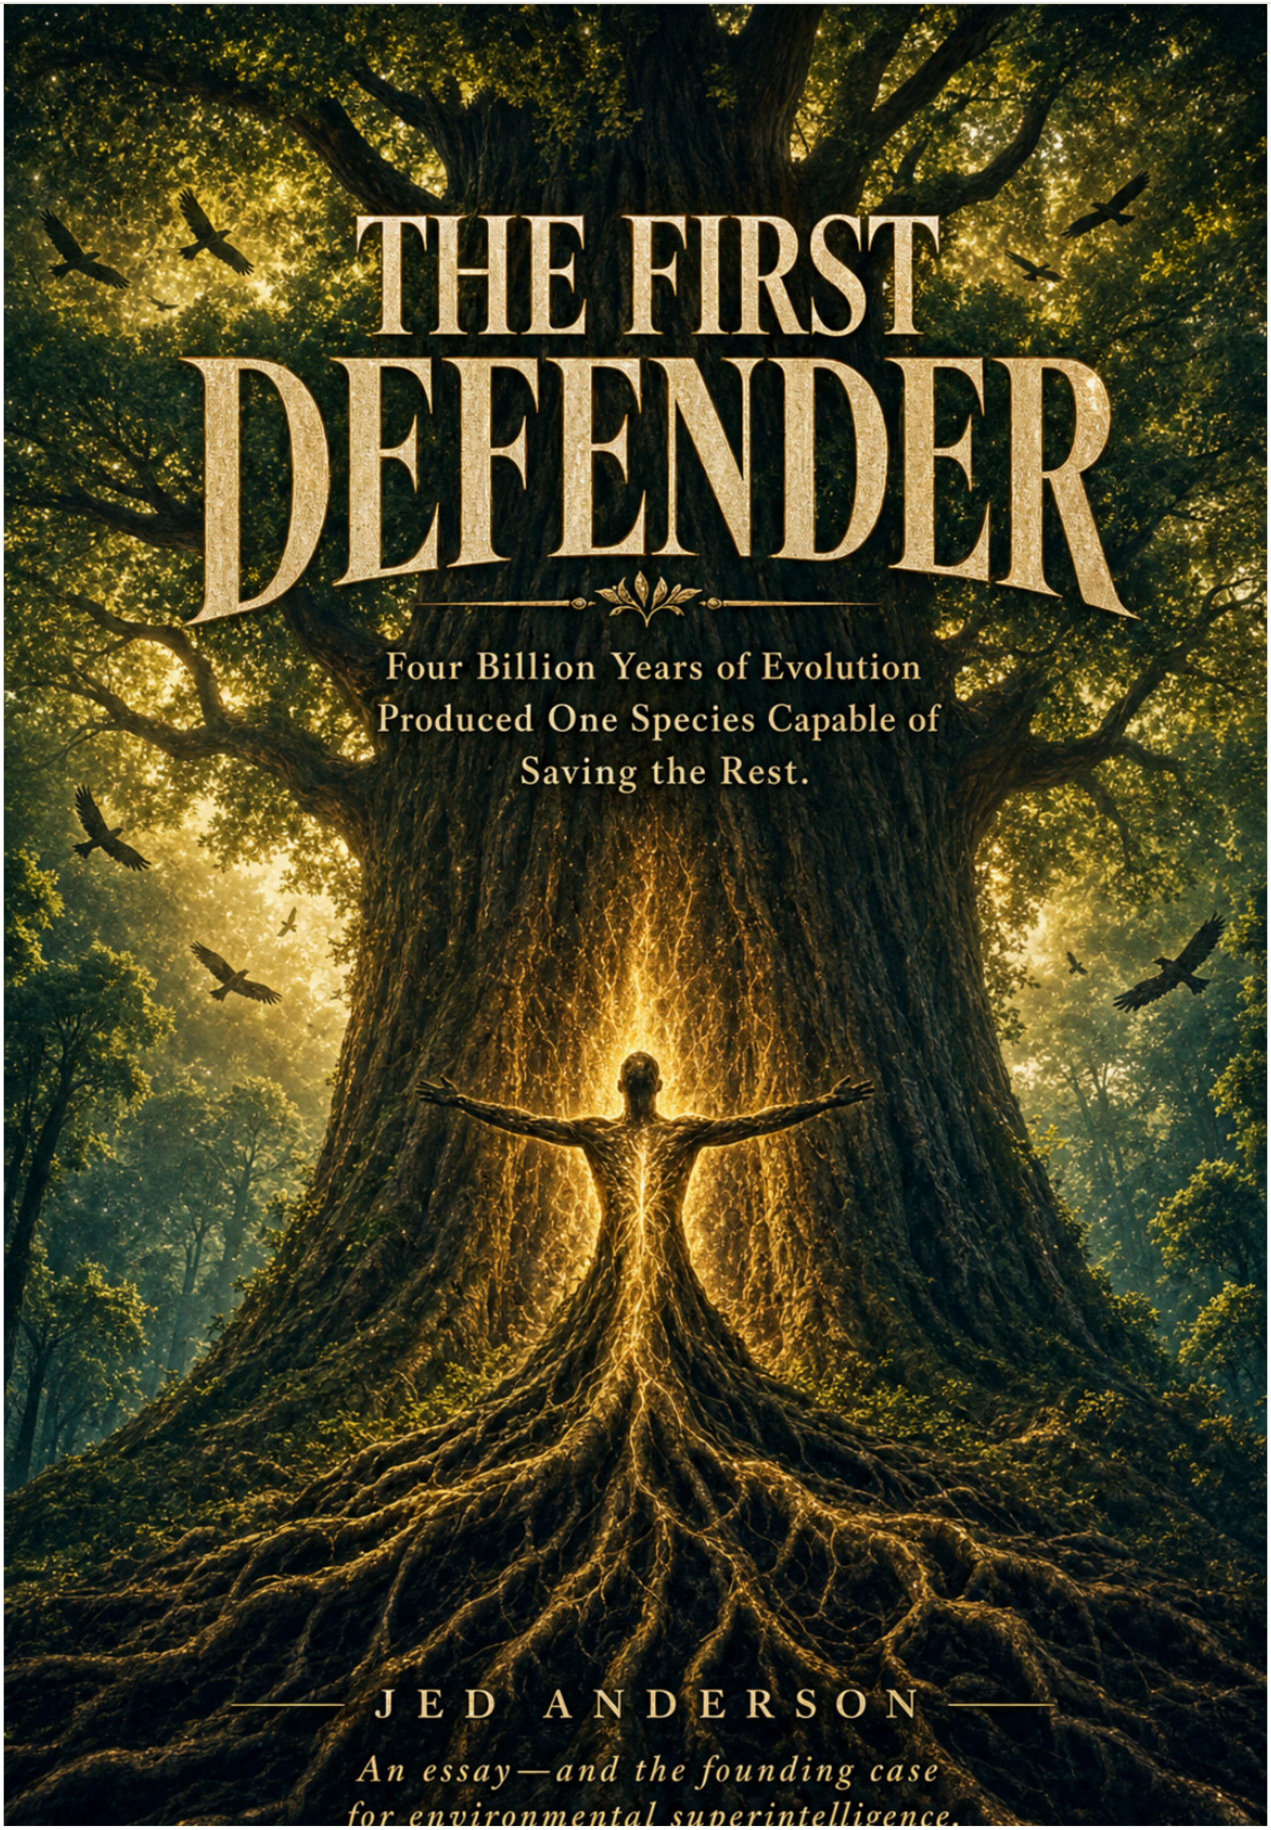
THE FIRST
DEFENDER
Four Billion Years of Evolution
Produced One Species Capable of
Saving the Rest.
JED ANDERSON
An essay—and the founding case
for environmental superintelligence.

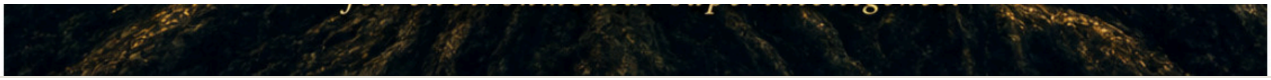

[Read PDF · 23 pages] [Raw markdown]

View on GitHub (https://github.com/jedanderson432/jedanderson-site/blob/main/src/content/essays/first-defender.md)

Earth was never going to make it.

The biosphere has been carpet-bombed five times in 500 million years, and the bombs have not stopped falling. Asteroid: one civilization-ender every 500,000 years. Supervolcano: one every 50,000. The Sun's slow brightening will boil the oceans in a billion years and swallow the planet in seven and a half.

Every species that has ever lived here has died.

Four billion years into life on this planet, the biosphere produced something it had never produced before. Not a stronger predator. Not a hardier microbe. Something that could read the clock.

In September 2022, that something moved a celestial body off its orbit around the Sun for the first time in 4.5 billion years.

We are not Earth's problem.

We are the part of nature that finally grew old enough to fight back.

I. The most important graph ever drawn

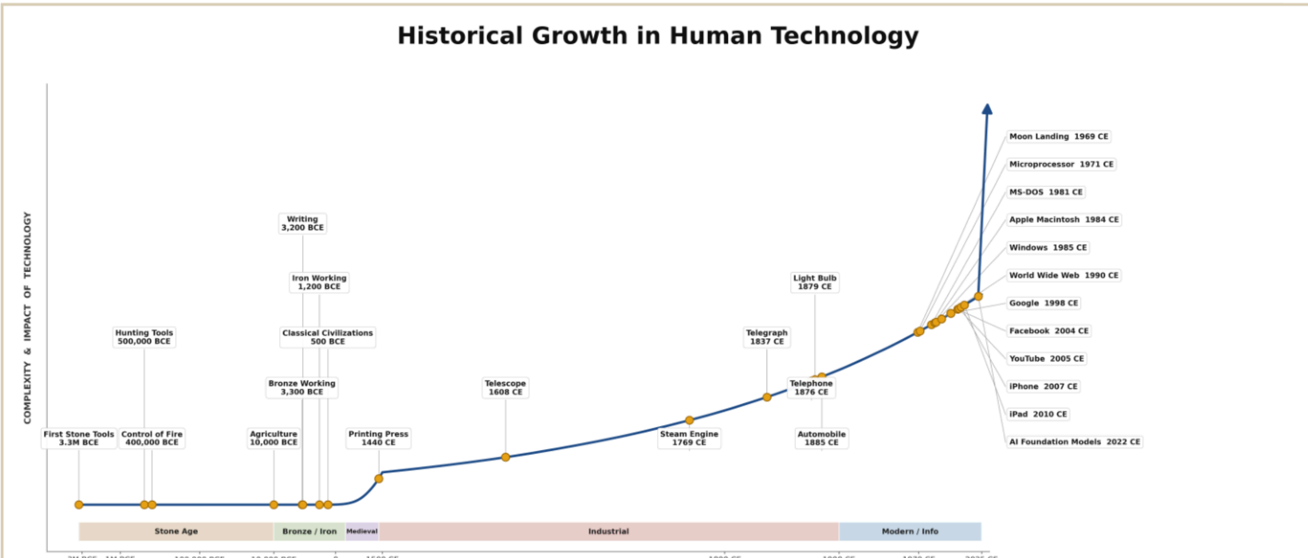

Historical growth in human technology, plotted on a logarithmic time axis. The vertical wall on the right is not an artifact of scaling—it is the signature of a species that learned, after four billion years of evolution, to manufacture knowledge on purpose.

Look at the curve. *Sit with it*.

For almost all of human history, the line is flat. Not slightly inclined. *Flat*. A handful of inventions per millennium, sometimes per ten millennia, generation after generation living and dying inside the same toolbox their great-grandparents had inherited and would pass down unchanged. Then, somewhere around 1500, the line twitches. By 1700 it bends. By 1800 it lunges. By 2000 it has detached from the page entirely—an arrow tearing off the top of the chart, with no ceiling and no asymptote in sight.

This is the most important graph ever drawn, and almost no one in the environmental conversation is looking at it honestly.

It is not a graph about gadgets. It is a graph about **knowledge**—the explanatory, error-correcting, hard-to-vary kind of knowledge that, once a civilization learns to manufacture it deliberately, never stops compounding. Almost everything we now call "the environmental crisis"—climate change, biodiversity collapse, watersheds running dry, plastic in the Mariana Trench—is a side effect of that vertical wall on the right side of the chart. And almost everything we will ever call an environmental *solution* lives on the part of the line that has not yet been drawn.

That is the wager of this essay, and the wager is total:

> *"The flatline is not a record of humility. It is a record of failure. It is the cost, in millennia, of a species that took far too long to discover what it was for."*

Every charismatic megafauna already lost, every gigaton of carbon already in the air, every aquifer drained beyond replenishment—all of it is, in the deepest accounting, a tax we are paying for being late. Late to writing. Late to print. Late to science. Late to the only practice in the four-billion-year history of life on Earth that has ever produced a way out of anything.

And underneath that truth is a harder one—the one the rest of this essay is built around:

> ***"A planet without knowledge-creating minds is not a pristine sanctuary. It is a death sentence with a long fuse."***

The cosmos has been trying to kill this biosphere on a regular schedule for half a billion years, and so far it has succeeded five times. The only thing that has ever stood up to that schedule—the only thing in the entire history of life on Earth that has ever even contemplated standing up to it—is the curve in front of you. Us. Knowledge. The arrow off the page.

The species that built fossil-fuel infrastructure is also the species that builds asteroid-deflection missions. And no one else is coming.

II. What "knowledge" actually means

The philosopher David Deutsch, in *The Beginning of Infinity*, gives the cleanest articulation of what is happening on the right side of that chart. His argument, distilled, is that knowledge is not a pile of facts and not a stack of justified beliefs. Knowledge is *good explanations*—accounts of reality that are **hard to vary** without ruining their explanatory power (Goodreads, *The Beginning of Infinity*).

The Greek myth that seasons happened because Persephone descended into Hades each autumn is a bad explanation: swap her name, swap the underworld, swap the cause and the story still "works." Newton's gravitation and the axial tilt of the Earth, by contrast, are hard-to-vary: change a single quantitative detail and the entire structure breaks against observation (Sashin Exists). One can be told to children. The other can land probes on Mars.

From that single move, Deutsch extracts an astonishing corollary: knowledge of this kind has unlimited reach. The same equations that describe a falling apple describe a satellite, an exoplanet, the precession of Mercury, and the gravitational lensing of light from quasars billions of years old. "All progress, both theoretical and practical, has resulted from a single human activity: the quest for what I call good explanations" (Nat Eliason notes on *Beginning of Infinity*).

From reach, Deutsch derives the principle of optimism: all evils are caused by insufficient knowledge. And from that, the most consequential sentence in the book for anyone working on environmental problems:

> *““Since the human ability to transform nature is limited only by the laws of physics, none of the endless stream of problems will ever constitute an impassable barrier. So a complementary and equally important truth about people and the physical world is that problems are soluble. By ‘soluble’ I mean that the right knowledge would solve them.” (Goodreads)”*

Read that twice. He is not saying problems solve themselves. He is not saying we already have the knowledge we need. He is saying something far stranger and far more demanding: **for every problem not forbidden by physics, there exists a body of knowledge—currently uncreated—whose creation would dissolve it.** The barrier is never the problem. The barrier is how fast we can manufacture the explanation that ends it.

That is the lever. That is the entire game.

III. Yudkowsky’s hidden corollary

Deutsch’s framing has a sharp twin. Eliezer Yudkowsky put it this way:

> *““There are no hard problems, only problems that are hard to a certain level of intelligence. Move the smallest bit upwards and some problems move from ‘impossible’ to ‘obvious.’ Move a substantial degree upwards, and all of them will become obvious.””*

This is not a contradiction of Deutsch—it is his thesis written from the other side of the table. Deutsch tells us *what dissolves problems* (good explanations). Yudkowsky tells us *that whether a problem feels "hard" is not a property of the problem but of the mind standing in front of it*. A clogged watershed is intractable to a Bronze Age village and trivial to a hydrologist with satellite imagery, GIS, and a numerical model. Atmospheric carbon is unsolvable to a civilization that has not yet discovered chemistry; it is a difficult-but-bounded engineering question to one that has.

The implication, taken seriously, breaks something in the standard environmental imagination. **The problems we currently call "wicked"—climate, biodiversity, ocean acidification, water rights—are not wicked in any cosmic sense. They are wicked at our current level of intelligence and knowledge.** Move that level up—through better science, better instruments, better explanations, better cognition (human or machine)—and they begin to flicker between "impossible" and "obvious" the way the heavens flickered for the medievals when Newton wrote the *Principia*.

And then Deutsch's third remark, which belongs in your pocket like a coin:

> *""I myself believe that there will one day be time travel because when we find that something isn't forbidden by the over-arching laws of physics we usually eventually find a technological way of doing it.""*

If that posture is justified for *time travel*—a thing nearly all of us would have called impossible a hundred years ago—what should we say about a stable climate, a continent of restored watersheds, oceans free of microplastics, a deflected meteor, a stilled supervolcano?

Those are not impossible. Those are not even hard, at the right level of explanation. They are simply not yet built.

IV. The four-thousand-year flatline and the parable of Easter Island

Now look back at the chart.

For most of recorded human history, our ancestors lived inside what Deutsch calls a *static society*—a culture organized around suppressing change rather than creating it. Their mechanism was not stupidity. They were not less smart than us. The difference was structural: their cultures were configured to disable the source of new ideas in children, in dissenters, in foreigners, in heretics (Sandor Dargo on Deutsch). Sparta did not honor living poets. Sparta did not have philosophers. Sparta could not imagine its own improvement, and so it did not improve.

Deutsch's most evocative environmental case study is Easter Island. Jared Diamond made it the parable of ecological hubris: a society chopped down its forests to roll giant statues, and starved. The historical picture is genuinely contested—archaeologists Hunt and Lipo have argued the collapse was driven more by introduced rats and post-contact disease than by deforestation alone, and the population trajectory may have been less dramatic than Diamond described—but the structural point survives every revision. Whatever the proximate cause, what made the island fragile was the same thing that made every pre-modern society fragile: it was a static culture whose dominant memes encouraged the endless re-enactment of the same project, and disabled the creativity that would have invented reforestation, alternative tools, or seafaring rescue. "The Easter Island civilization collapsed because no human situation is free of new problems, and static societies are inherently unstable in the face of new problems" (Static Societies, citing Deutsch). The trees may or may not have sat literally beneath the statues. The cultural rigidity sat metaphorically beneath them either way.

This is the most important sentence in the book for environmental thinkers, because it inverts the entire degrowth narrative: **the deepest cause of ecological collapse is not too much technology, but too little knowledge-creating freedom.** The Easter Islanders did not need fewer tools. They needed a Republic of Letters, a printing press, a scientific revolution. They needed the conditions under which someone could say "the elders are wrong, the gods are not watching, the trees will not grow back if we keep this up"—and be heard.

We had Easter Island as a planet for about four millennia. Then we got lucky.

V. The luckiest two centuries in the history of mind

The historical record now converges on something close to Deutsch's view. Joel Mokyr's *A Culture of Growth* argues that the European Enlightenment was not exogenous magic but an emergent property of two specific institutional changes: a Republic of Letters that allowed criticism and ideas to flow across borders, and a politically fragmented landscape that let dissident thinkers move when one regime tried to silence them (Cato Institute on Mokyr; IMF review). Add the printing press (1440s), Copernicus (1543), Galileo, Bacon, the Royal Society (1662), and Newton's

Principia (1687) (Wikipedia: Scientific Revolution), and you have, **for the first time in human history, a civilization that knew how to manufacture knowledge on purpose.**

VI. Counterfactual: what fifty years and five hundred years would have bought us

Now run the experiment the chart invites. Slide the curve along the time axis.

Fifty years earlier. Move the inflection from ~1800 to ~1750. The Industrial Revolution begins under late-Enlightenment intellectual conditions but with fifty more years of compounding before any of the dirty energy infrastructure becomes politically locked in. Watt's separate condenser arrives in the 1710s instead of the 1760s. Faraday's electromagnetism is born into a world already wired for it. The first global telegraph is operational by 1820. Pasteur's germ theory by 1810. Antibiotics by 1880. By 1900, in this counterfactual, we are roughly where we were in 1950—meaning by 2026 we are somewhere we have not yet been. The carbon plume of the twentieth century never accumulates because the clean energy transition arrives before fossil-fuel infrastructure becomes civilizational furniture.

Five hundred years earlier. Now imagine the Scientific Revolution begins around 1050 instead of 1550. The conditions are almost there: China has the printing press and gunpowder, the Islamic Golden Age is in full bloom, Indian mathematicians have invented zero and decimal notation, and Constantinople is still standing. What is missing is the Republic of Letters—the self-perpetuating culture of criticism, the institutional habit of *trying to prove yourself wrong*. Imagine that habit takes root in 11th-century Baghdad or Song-dynasty Kaifeng or 12th-century

Toledo. Galileo's telescope appears in 1109. Newton's *Principia* in 1187. The Industrial Revolution ignites around 1300. By 1600, in that timeline, fossil fuels have already been replaced; nuclear power arrives in the 1700s; fusion in the early 1800s. Atmospheric CO_2 never leaves the 280-ppm preindustrial baseline because the carbon-burning phase ends before the coal seams are fully tapped. The Amazon stays intact. By 2026 in that timeline, we have had fusion for three centuries. We have had molecular agriculture since the Renaissance. We have not lost a single species we now mourn.

That is what the cost of being slow looks like.

Every species we have lost since 1500. Every gigaton of carbon now in the air. Every reef bleached. Every aquifer drained. **All of it is a tax—paid in geological currency—on having taken four millennia to figure out that knowledge can be manufactured deliberately.** Deutsch's parable of Easter Island is the parable of the entire planet. We are the inhabitants of a static civilization that woke up just in time, and we are still rubbing the sleep from our eyes.

VII. The asymmetry that breaks the standard environmental story

There is an asymmetry buried in this argument that environmental thought has, on the whole, refused to confront, and it is this: **the same engine that produced the problem is the only engine that has ever solved one.**

The Industrial Revolution gave us anthropogenic climate change. It also gave us the satellites that detect it, the spectroscopy that quantifies it, the supercomputers that model it, the photovoltaic effect that lets us escape it, and the global communications network that lets a watershed scientist in Montana coordinate with a paleoclimatologist in Switzerland in the time it takes to send a packet. There is no version of the climate response that does not run on the products of the very curve it is trying to bend. Wind turbines are not Iron Age artifacts. Lithium-ion batteries are not medieval. The ozone hole closed because we *understood* the chemistry of CFCs, banned them, and engineered substitutes—every step a Deutsch-style hard-to-vary explanation operationalized into industrial practice.

This is why Deutsch is so caustic about the word "sustainability" as it is commonly deployed. *"There is no such thing as 'sustainability.' Things are sustainable until they're not."* A static society can sustain itself for a few centuries by enforcing identical behaviors generation after generation. It cannot sustain itself against a single new problem. By contrast:

> *""Progress is sustainable, indefinitely. But only by people who engage in a particular kind of thinking and behaviour—the problem-solving and problem-creating kind characteristic of the Enlightenment. And that requires the optimism of a dynamic society." (Christian Houmann's notes)"*

The choice on offer is not change versus stability. The choice is between **a dynamic civilization that survives by continuously inventing its way through new problems, and a static one that survives until the first problem it cannot dissolve.** Easter Island was sustainable, right up until it wasn't.

The bold environmental claim, then, is the inversion of the standard one. It is not that humans are a virus on Earth, that we should consume less, hide our footprints, retreat into smallness. It is that **humans are the only thing on Earth that has ever solved a planetary problem**, and the only ethical posture available to us is to get faster, smarter, and more explanatorily powerful—not slower, smaller, more humble. “In the pessimistic conception, that distinctive ability of people is a disease for which sustainability is the cure. In the optimistic one, sustainability is the disease and people are the cure” (Nat Eliason notes).

Read that twice. **Sustainability is the disease and people are the cure.** That is not a tagline. It is a complete reversal of the moral economy of modern environmentalism, and it is the only frame under which a 21st-century environmental technology company makes sense at all.

VIII. The cosmic ledger: a planet without knowledge is a planet already condemned

Here is where the argument becomes uncomfortable for both sides of the standard debate, and where it becomes most important.

Set aside, for a moment, every problem humans have caused. Pretend we never existed. Imagine the pristine planet some fraction of the environmental movement still mourns: untouched forests, unbroken coral, rivers without dams, skies without contrails. Now ask: what is its expected lifespan, on the cosmic schedule the universe actually runs on?

Here is the ledger. None of these numbers are speculative. All of them are in the geological record.

The asteroid clock. Roughly 66 million years ago, an asteroid about ten kilometers (six miles) across struck the Yucatán at 20 km/s, releasing roughly 72 teratonnes of TNT and triggering the Cretaceous–Paleogene extinction. The dinosaurs did not lose because they were unfit; they lost because they had no telescopes (Wikipedia: Chicxulub crater; NASA: Deep Impact). Britannica places the recurrence interval for civilization-ending impactors (≥1 km) at roughly once every 500,000 years—and the curve gets exponentially more brutal further up the size distribution (Britannica: Earth impact hazard). The Earth has been struck before. It will be struck again. This is not a question of *if*.

The supervolcano clock. At least 60 confirmed VEI-8 eruptions are recorded in the geological history of the planet; published estimates of the recurrence interval cluster around one supereruption every 50,000 years, with substantial uncertainty in either direction (Wikipedia: Volcanic Explosivity Index; OzGeology on VEI-8 supervolcanoes). The Toba eruption ~74,000 years ago laid down 2,800 cubic kilometers of magma and may have come close to ending our species in a decade-long volcanic winter. Yellowstone has done it twice in the last 2.1 million years. The magma is still down there. It will erupt again.

The solar clock. The Carrington Event of 1859 was the largest geomagnetic storm ever recorded. In 1859 it lit telegraph paper on fire and gave operators electric shocks. In 2026 a Carrington-class storm could melt or destroy hundreds of high-voltage transformers, blacking out grids for *months to years*, with U.S. damage estimates between $600 billion and $2.6 trillion (Texas Public Policy Foundation, 2026; Astronomy Magazine). Tree-ring carbon-14 spikes and ice-core data suggest the Sun launches storms larger than Carrington roughly once every few centuries. We have not had one since the grid existed.

The gamma-ray burst clock. A nearby GRB, beamed at Earth, would strip a substantial fraction of the ozone layer in roughly a month, with modeling work by Piran and Jiménez suggesting GRBs may have caused at least one mass extinction over the last 500 million years (CERN Courier on Piran & Jiménez; Wikipedia: Gamma-ray burst). The Late Ordovician extinction (~444 Mya) is one of the leading suspects.

The Sun itself. In about 1.1 billion years, the Sun's luminosity will rise by 10%, triggering a runaway greenhouse and boiling Earth's oceans. In about 7.59 billion years, the Sun will swallow the planet outright (Phys.org on Schroder & Smith; Wikipedia: Future of Earth). The biosphere has, at most, a billion-year lease on this planet—and the lease is non-negotiable to anything that can't physically leave.

The historical record. Earth has already had five mass extinctions (Our World in Data; Wikipedia: Extinction event). The End-Permian wiped out roughly 96% of all species. The End-Ordovician, roughly 85%. The End-Cretaceous, roughly 76%. None of these were caused by humans. Humans had not been invented yet. They were caused by the universe doing what the universe does to unprotected biospheres: rolling its dice on a geological cadence and occasionally rolling extinction.

Let this number land for a moment: the natural background extinction rate, calculated as if humans had never existed, would still erase essentially every species currently on the planet within a few tens of millions of years (Wisconsin/Peery on background extinction). The "untouched" Earth is not a stable steady-state from which we have departed. It is a temporary configuration sliding toward an inevitable cliff, with periodic asteroid impacts, supervolcanic winters, gamma-ray showers, and stellar evolution all queued up in the schedule.

Now hold both halves of the ledger in your head at the same time:

- **Humans are causing the sixth mass extinction.** Modern vertebrate extinction rates are up to 100× the natural background rate under conservative assumptions—and the rate is rising (*Science Advances*, Ceballos et al., 2015). This is a real, urgent, civilizational failure.
- **Humans are also the only force in the entire history of this planet capable of preventing the seventh mass extinction**, the one the universe has scheduled regardless of what we do.

Both sentences are true at the same time. The standard environmental story holds only the first one. The Deutschian one holds both, and only by holding both does the picture become accurate.

This is the part of the argument that should seize the attention of every serious environmental thinker: **a planet of forests and reefs and whales without a knowledge-creating civilization is not safe. It is condemned, on a clock measured in geological epochs.** Earth's biosphere has already been carpet-bombed five times in 500 million years, and the bombs have not stopped falling. The dinosaurs did not have a space program. The trilobites did not have a climate model. The ammonites did not have spectroscopy. We do, and we are the only ones who do, and we are the only ones who ever have.

The miracle of Earth is not the biosphere. The biosphere has been here before, and the biosphere has been wiped out before, and the biosphere will be wiped out again, on schedule, unless something interrupts the schedule. **The miracle of Earth is the species that finally, after four billion years, started keeping its own appointment book.**

IX. The species that fights back

In 2022, NASA's DART spacecraft slammed into a 170-meter asteroid moonlet called Dimorphos at 14,000 miles per hour. It shortened the moonlet's orbital period by 32 minutes—more than 25 times the threshold for "success"—and even shifted the Didymos-Dimorphos pair's orbit around the Sun by 150 milliseconds (NASA: DART mission; NASA on DART altering solar orbit, March 2026; NYT: DART analysis, March 2026). One hundred fifty milliseconds doesn't sound like much, until you realize that for the first time in 4.5 billion years, a celestial body changed its orbit around the Sun because something on Earth wanted it to.

Stop and feel the weight of that. For the entire previous history of life—bacteria, trilobites, dinosaurs, mammoths, every creature that ever drew breath—the rocks coming for this planet were unanswerable. They came, they hit, they killed, on the universe's timetable, with no possibility of negotiation. In September 2022, that ended. We did not just *imagine* deflecting an asteroid; we *did* it, and we measured the result to a precision of milliseconds, and the result said *yes, this works, you can do this*.

That is the most important environmental act in the history of the planet. Not because of Dimorphos, which was no threat to anyone, but because of what it proved about the possibility space. **Asteroid deflection is no longer hypothesis. It is engineering.** The dinosaurs lost a coin flip. We took the coin away.

Now extend the principle. Every existential threat on the cosmic ledger above can be addressed in exactly the same way—not because we have built the system yet, but because nothing about any of them is forbidden by the laws of physics, which is Deutsch's only criterion.

- **Supervolcanoes** can be drilled, monitored, and gradually pressure-relieved by extracting geothermal energy from the magma chamber—turning the threat itself into a clean energy resource. Nothing in physics prevents it. We just haven't built it yet.
- **Carrington-class solar storms** can be defeated by hardened transformers, distributed grids, fast-disconnect protocols, and forecast lead times that AI is already shrinking from minutes to days.
- **Gamma-ray bursts** are detectable; the ozone layer can, in principle, be replenished chemically; populations and biotic samples can be hardened against UV pulses. None of it forbidden by physics. None of it built yet.
- **The Sun's eventual expansion** is the longest-running existential threat on the books, and it is the most spectacular argument for Deutsch's principle, because there is, somewhere on the line that has not yet been drawn, a future generation of human-and-machine intelligence that will physically move Earth's orbit, or move life off Earth entirely, or do something we cannot currently imagine that makes the question moot. That is not science fiction. That is just the laws of physics and enough time and enough good explanations. Deutsch on time travel applies in full force here.

This is what "the advancement of nature" actually means. Not preservation. Not stewardship in the gardener's sense. **Defense.** Active, technological, knowledge-driven defense of a biosphere that has never had a defender before, against a universe that has been trying to kill it on a regular schedule for four billion years.

The standard environmental imagination treats nature as a fragile ward to be sheltered from human intrusion. The Deutschian one treats nature as a magnificent, doomed thing that has, for the first time in its history, produced an organ capable of saving it. **We are that organ. We are nature learning to defend itself.** The forest, the reef, the wolf, the whale—none of them can build a telescope. None of them can model a climate. None of them can drill a magma chamber or harden a power grid or deflect a rock. We can. We are the only ones who can. And we are, in the deep accounting, the part of nature that is finally—finally—old enough to take care of the rest of nature.

We hurt the Earth getting here. That is true and there is no scrubbing it. We are also the only thing in the four-billion-year history of this planet that has even attempted to keep it alive past the next cosmic appointment. **Both sentences are true. The first one gets all the attention. The second one is the actual stakes.**

X. The second curve: when knowledge-creation itself becomes automatable

And now we arrive at the part of the argument that should keep environmental scientists awake at night for the right reasons.

Deutsch wrote a sentence in 2011 that almost no one has metabolized:

> *““All technological knowledge can eventually be implemented in automated devices. This is another reason that ‘one percent inspiration and ninety-nine percent perspiration’ is a misleading description of how progress happens: the ‘perspiration’ phase can be automated.” (books.max-nova.com on Beginning of Infinity)”*

For four hundred years, the bottleneck on knowledge-creation has been human beings reading, thinking, conjecturing, criticizing, and writing things down. The throughput has been bounded by the number of trained scientists alive at any one time, the speed of mail and print, and the duration of a human career. The arrow off the chart is what that bottlenecked process produced. We are now, in the second quarter of the 21st century, watching the first credible attempt to *unbottleneck* it.

The environmental ledger of just the last thirty-six months:

- **Materials discovery.** Generative AI platforms are compressing what used to be a 10–20-year search for new climate-relevant compounds—battery cathodes, carbon-capture sorbents, photocatalysts—into a process measured in months, with viable-candidate yield rising from roughly 6% under traditional R&D to as high as ~90% in some pipelines (News → Sustainability Directory, 2025). An MIT study tracking AI-assisted materials labs found a 44% jump in new materials discovered, a 17% rise in product prototypes, and a 39% surge in patent filings—and the AI-assisted patents introduced more genuinely new technical terminology, suggesting the AI is not just retrieving, it is generating (Climate Adaptation Platform on AI + MOFs).
- **Carbon capture.** ML-guided catalyst design for the CO_2 reduction reaction is moving from trial-and-error to inverse design—specify the property, let the model propose the molecule (Newswise / eScience, 2024).
- **Fusion.** DeepMind and EPFL's Swiss Plasma Center demonstrated in 2022 that a single deep RL agent can shape and stabilize tokamak plasma, including configurations no human controller had ever produced (DeepMind, 2022; Science Alert summary). In 2025 the partnership extended to Commonwealth Fusion Systems' SPARC reactor, with the open-source TORAX simulator now in the global fusion community's hands (DeepMind, 2025). MIT's PORTALS

framework runs plasma simulations 10,000× faster than legacy approaches (LinkedIn / Heather-Anne Scott, 2025). The 70-year plasma-control problem is no longer open.

- **Weather and climate intelligence.** GraphCast produces 10-day global forecasts in under a minute on a single TPU, beats the gold-standard HRES system on 90% of 1,380 verification targets, and identifies extreme events earlier and more accurately than the supercomputer-backed pipeline that has dominated since the 1960s (DeepMind / *Science*, 2023; Science paper). Hurricane Lee's Nova Scotia landfall in 2023 was correctly forecast nine days out—three days earlier than legacy models (Attrecto on AI weather models).
- **Self-driving labs and Industry 5.0.** Robotic, AI-orchestrated experimental platforms are collapsing the loop between hypothesis and verification, with one recent review estimating up to a 75% reduction in materials discovery time—equivalent to fifteen years of compressed innovation (*Communications Materials*, 2026).

These are not press releases. They are the first credible demonstrations of the thing Deutsch said was possible in principle: **the perspiration phase of knowledge-creation is now automatable.** The implication for environmental work is staggering. Watershed dynamics, regulatory compliance, ecosystem modeling, hydrological simulation, atmospheric chemistry, agricultural optimization, real-time pollutant tracking, biodiversity monitoring, asteroid surveillance, supervolcano monitoring, solar weather prediction—every one of these is, in Deutschian terms, a problem of insufficient explanatory knowledge operating at human cognitive throughput. Lift the throughput, and the entire problem space changes shape.

What happens to environmental intelligence when the cost of running a watershed-scale simulation drops by four orders of magnitude? When a CFR compliance analysis that took a consulting team six weeks runs in six minutes? When the bottleneck on EPA modernization is no longer "how many engineers can read the rule" but "how fast can we generate hard-to-vary explanations of every contaminant pathway in every watershed in North America"? The answer is the answer Deutsch gave for everything: we don't yet know, because we have not yet created the knowledge that would tell us. But the curve does not bend down from here. The curve goes up.

And so the obvious next step—obvious only in the Yudkowskian sense, that it becomes obvious once the level of intelligence required to see it has been reached—is to stop applying these tools to environmental problems one at a time and start building the underlying organ itself: a system whose explicit purpose is to automate the creation and use of environmental knowledge at planetary scale. Call it what it is: **environmental superintelligence.** Not a chatbot. Not a dashboard. Not a smarter search engine. A cognitive infrastructure for the biosphere—one that fuses the linguistic and regulatory layer (the corpus of environmental science, every CFR provision, every monitoring record, every watershed assessment ever written) with physics-based models of how air, water, soil, and ecosystems actually behave, so that hard-to-vary explanations can be generated, criticized, and operationalized at machine speed. The early outlines of such a system are already in the field. Environmental intelligence platforms trained on the full corpus of environmental documents and operational data are now in production with Fortune 500 customers, beginning to automate the professional tasks—compliance analysis, regulatory interpretation, monitoring synthesis, watershed assessment—that have historically been bottlenecked by human reading speed. The next move, the one that turns the project from a tool into an organ, is to integrate the physics. When a regulatory analysis runs against a real-time hydrological simulation of the watershed it governs, when a permitting decision can be evaluated against a coupled model of contaminant transport and ecosystem response, when a compliance question is answered by a system that actually understands the river it is being asked about—that is the threshold being crossed. Once it is crossed, the four-thousand-year flatline ends in a way it has never ended before: not because more humans are thinking, but because thinking itself, applied to the planet, has been industrialized.

This is what "the advancement of nature" finally cashes out to in operational terms. Not a slogan. Not a metaphor. **A planetary-scale knowledge system that the biosphere has never had before, being built—for the first time in four billion years—by the only species capable of building it, in time, we hope, to matter.** The deadline is real. The cosmic ledger does not negotiate. But the project is not speculation. It is under construction. Right now. In labs and fields and codebases that did not exist five years ago.

XI. The duty to be optimistic—and to be in a hurry

Deutsch's most demanding sentence is also his most environmental:

> *""We have a duty to be optimistic. Because the future is open, not predetermined and therefore cannot just be accepted: we are all responsible for what it holds. Thus it is our duty to fight for a better world." (Antoine Buteau, Lessons from David Deutsch)"*

Optimism here is not a temperament. It is an ethical position. It is the refusal to accept that any present problem is permanent, combined with the recognition that no problem solves itself. *“Problems are inevitable. Problems are soluble. Get to work.”* (Gio Lodi on Deutsch)

If the chart is right—and the historical record confirms it is—then the four-thousand-year flatline was the anomaly, not the curve. We are living inside the few generations during which a previously static species became dynamic, learned to manufacture good explanations on purpose, and is now beginning to automate the manufacturing process itself. Every day that compounding continues, the universe of “impossible” environmental problems shrinks and the universe of “obvious” ones grows. That is Yudkowsky’s claim restated for the planetary file. That is Deutsch’s optimism made operational. **That is what the arrow off the page is for.**

Imagine the alternate Earth where this began in 1300. We don’t get to live there. But we have something nearly as valuable: the ability to make every year from now onward count for what those missing five hundred years would have given us. Each generative chemistry model trained on environmental molecules, each plasma controller that turns fusion from physics into engineering, each watershed that gets a real-time digital twin, each compliance analysis that drops from weeks to minutes, each asteroid catalog updated, each transformer hardened, each magma chamber drilled, each step toward an environmental superintelligence that can think about the planet faster than the planet can break—each of these is a small piece of the curve we failed to draw between 1500 and now, being drawn at last.

The environmental crisis, taken seriously, is not a verdict on civilization. It is a deadline. The deadline says: the species that learned to manufacture knowledge in 1700 must now automate that manufacturing fast enough to clean up after itself, deflect what’s coming, and outlast a star—before the bills come due.

Deutsch's gift is to insist that this is possible. Not certain. Not easy. Not free. But *possible*—because nothing about a stable climate, a restored watershed, a regenerated biosphere, a deflected meteor, a tamed supervolcano, a hardened grid, or a relocated planet is forbidden by the laws of physics.

What is required is not less of us. What is required is more of what makes us specifically and irreducibly powerful in this universe: the capacity to generate hard-to-vary explanations of how things really work, and to build them into the world before the world ends.

We hurt the Earth getting here. That is the price of the ticket. **The Earth was already going to die. We are the only thing in its entire history that has ever arrived in time to save it.**

The chart starts flat. It bends. It launches. The arrow leaves the page.

We are the species that draws the next part of the line. There is no one else here to draw it. The dinosaurs are not coming back to help.

Get to work.

> ***"Earth was never going to make it.***
>
> ***We are why it might."***

Sources

David Deutsch, *The Beginning of Infinity: Explanations That Transform the World*—primary text; quotations and concepts via Goodreads quote pages, Nat Eliason's notes, Christian Houmann's notes, books.max-nova review, Sashin Exists on hard-to-vary explanations, Static Societies compilation, Antoine Buteau's lessons, and Summify on Deutsch's knowledge creation.

Joel Mokyr, *A Culture of Growth*—via Cato Institute, IMF Finance & Development, and CREI working paper.

Wikipedia: Scientific Revolution.

Cosmic ledger:

- Wikipedia: Chicxulub crater; NASA: Deep Impact and the Mass Extinction of Species 65 Million Years Ago; Britannica: Earth impact hazard frequency.
- NASA: DART mission overview; NASA: DART altered Didymos's orbit around the Sun (March 2026); NYT: DART analysis (March 2026).
- OzGeology: Earth's Active Supervolcanoes; Forbes / USGS on supereruption frequency.
- Texas Public Policy Foundation on Carrington-class solar threat (2026); Astronomy Magazine on Carrington-class storms and the modern grid; Wikipedia: Carrington Event.
- CERN Courier on GRBs as life-extinction threats (Piran & Jiménez); Wikipedia: Gamma-ray burst.

- Phys.org: Will Earth survive when the Sun becomes a red giant?; Wikipedia: Future of Earth; The Conversation: 1 billion years left of habitability.
- Our World in Data: Five mass extinctions; Wikipedia: Extinction event; Big Think: What caused Earth's 5 mass extinctions?; Wisconsin/Peery on background extinction rates.
- Ceballos et al., *Science Advances* 2015—accelerated human-induced species losses (modern vertebrate rates up to 100× background under conservative assumptions).

AI-accelerated knowledge creation:

- DeepMind / EPFL, *Accelerating fusion science through learned plasma control*, 2022; *Bringing AI to the next generation of fusion energy*, 2025; Science Alert recap; MIT PORTALS / 70-year plasma problem (LinkedIn).
- DeepMind, GraphCast; Lam et al., *Learning skillful medium-range global weather forecasting*, *Science*, 2023; Attrecto: Exponential Impact of AI Weather Models.
- News → Sustainability Directory: Generative AI Platform Accelerates Climate Material Discovery; Climate Adaptation Platform on AI + MOFs; Newswise / eScience on ML + CO_2 catalysts; Communications Materials on AI-driven materials infrastructure.
- Eliezer Yudkowsky—quote on intelligence and problem hardness (rationalist canon).

Download PDF

Markdown source: https://jedanderson.org/essays/first-defender.md (https://jedanderson.org/essays/first-defender.md)

Source on GitHub: /src/content/essays/first-defender.md (https://github.com/jedanderson432/jedanderson-site/blob/main/src/content/essays/first-defender.md)