ESSAY · MACHINE CURIOSITY

The Automation of Why

On questioning, entropy, and the birth of machine curiosity. Why teaching an artificial mind to ask "why?" may be the most consequential thing we can do for industry and for the Earth.

BY JED ANDERSON

Every human being begins life as a scientist, asking why without limit, until the world teaches us to stop. What if the asking itself could be automated? An argument, from physics and information theory, for the most consequential faculty a machine has ever been given.

"The important thing is not to stop questioning. Curiosity has its own reason for existing."

— ALBERT EINSTEIN, 1955

"Everything that is not forbidden by the laws of nature is achievable, given the right knowledge."

— DAVID DEUTSCH, THE BEGINNING OF INFINITY

I. The first and last question

Every human being begins as a scientist. Before we can tie a shoe, we can ask *why*, and then, to the exhaustion of every parent who has ever lived, we ask it again. *Why is the sky blue? Why? Why? Why?* Each answer is a door, and behind every door is another door, and the child, uncorrupted by the fatigue of adulthood, walks through all of them. It is the purest expression of intelligence we will ever exhibit, and it is the first thing the world trains out of us.

We stop asking because asking is expensive. It costs time, attention, and the small social humiliation of admitting we do not know. By the time we are professionals, we have learned to ration our questions the way a desert traveler rations water. We ask only what we must. And in that rationing lies one of the great hidden tragedies of the modern enterprise: the most valuable questions, the ones that would save millions of dollars, avert disasters, and protect the living world, are almost never asked, because no one has the time, the standing, or the sheer relentless curiosity to ask them.

This essay is about the power of that abandoned faculty. It is about what a question actually *is*, measured in the hard currencies of information, energy, and entropy, and about why the act of asking is, quite literally, the cheapest and most powerful lever in the known universe. And it is about a threshold we have just crossed: for the first time in the history of mind, we can build machines that ask. Not machines that answer when spoken to, but machines that wonder. Machines that get curious. Machines that ask *why*, and then ask it again, and again, forever, and never grow tired.

I want to convince you that automating the question, automating *why*, is not a productivity feature. It is a phase change in the relationship between intelligence and reality. And I want to show you, concretely, how it is about to transform the work of environmental protection: the financial, the operational, and the ecological, all at once.

A child would call this magic. When we are very young, the word *why?* feels like a spell. You say it, and the world cracks open and offers up a reason, and behind that reason lies another, and the incantation never runs out. This essay will not try to talk you out of that wonder. It will do something better. It will show that the magic is entirely real, and that it is made of physics and mathematics of the most fundamental kind we possess; that the power of a question can be measured, in bits and in joules; and that it is, if anything, *more* astonishing for being explicable. What looks like a child's spell turns out to be the deepest lever in the universe. We are simply going to learn how it works, and then build a machine that casts it without tiring.

II. What a question actually is

Strip away the grammar and a question is an operation performed upon uncertainty.

In 1948, Claude Shannon gave uncertainty a number. The entropy of a situation, $H = -\Sigma\, p \log p$, measures in bits how much you do not know. A fair coin before it

lands carries exactly one bit of entropy. A question is a device for spending that ignorance down. Ask a perfectly designed yes-or-no question and you can, at best, cut the space of possibilities in half: one bit purchased, one bit of the unknown converted into the known. This is why the game of Twenty Questions works. Twenty well-aimed binary questions can, in principle, single out one item from over a million, because each question is a scalpel that halves the remaining darkness.

But not all questions are equal, and this is the crux of everything that follows. A question's worth is the amount of uncertainty its answer is expected to dissolve *about something that matters*. Statisticians call this **expected information gain**; decision theorists call it the **value of information**. The mathematics, laid down by Lindley and others in the 1950s, is precise: the best question to ask next is the one whose answer will, on average, most change what you believe and therefore what you do. A brilliant question is not a display of erudition. It is a maximally efficient act of uncertainty reduction, pointed at a decision that hangs in the balance.

Here, then, is the first profound truth. Intelligence is not primarily the ability to answer. It is the ability to ask: to know, out of the combinatorial infinity of possible questions, which one is worth asking. The answer is mechanical once the question is right. The genius is upstream, in the asking.

PLATE I A single figure raises one question against the darkness, and the sky answers with a galaxy of them. In thermodynamic terms the question is nearly free; what it can steer is not. Artwork: courtesy of the author.

III. The cheapest lever in the universe

Now let us descend from information into physics, because a question is not merely a mathematical abstraction. It is a physical act, and its cost can be measured in joules.

In 1867 James Clerk Maxwell imagined a demon: a tiny intelligence guarding a door between two chambers of gas, letting fast molecules through one way and slow molecules the other, sorting hot from cold, manufacturing order out of chaos and seemingly violating the Second Law of Thermodynamics. For a century the demon haunted physics. Its exorcism, completed by Leó

Szilárd, Rolf Landauer, and Charles Bennett across the twentieth century, delivered one of the deepest results we possess: **information is physical**, and it is the very substance from which order can be bought.

The demon's power comes entirely from *asking a question*, from making a measurement. *Is this molecule fast or slow?* One bit acquired. With that single bit, the demon can position its gate and extract useful work from what looked like pure disorder. Szilárd showed the exact exchange rate: one bit of information is worth $kT\ln 2$ of work, about three billionths of a billionth of a joule at room temperature ($\approx 3\times10^{-21}$ J). Landauer showed the debt on the other side: erasing that bit costs the same. Information and entropy are two faces of one coin.

Sit with the magnitude of this. To *ask*, to acquire one bit about the world, costs on the order of 10^{-21} joules. To physically *manipulate* a single molecule of matter directly costs on the order of 10^{-19} joules. That is a factor of roughly one hundred million. **The universe charges you a hundred million times less to know something than to shove something.** And yet the bit you acquire can direct the shoving of an entire mole of matter. A well-placed question is the highest-leverage action available in physics: a near-free acquisition of information that steers the vastly more expensive rearrangement of the physical world.

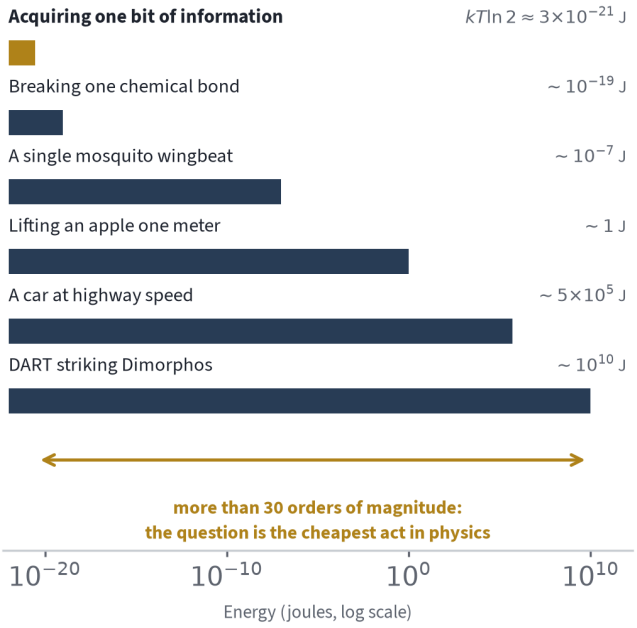

FIGURE 1 The price of knowing. The thermodynamic cost of acquiring one bit of information, $kT\ln 2$ at room temperature (Landauer's principle), set against the energies of moving matter. The well-placed question sits more than thirty orders of magnitude below the largest actions it has steered.

This is not metaphor. It is the thermodynamic reason that intelligence beats brute force, that a map is worth more than a stronger back, that the right diagnostic question is worth more than a warehouse of medicine. Life itself runs on this arbitrage. Erwin Schrödinger, in *What Is Life?*, observed that living things survive by feeding on "negative entropy," by importing order and exporting disorder, holding back the tide of decay. And how does a living system import order? By sensing, by measuring, by *asking the environment questions* and configuring itself around the answers. To question is to perform the fundamental metabolic act of any intelligence: to convert cheap information into expensive order.

Pollution, seen through this lens, is not fundamentally a chemical problem. It is an **information problem**: a failure to ask, in time, the questions that would have configured the gate differently. Which brings us to why the questions of environmental stewardship are, in the deepest physical sense, worth so extraordinarily much, and cost so little to ask.

IV. The descent of why

There is a second dimension to the power of questioning, beyond the single well-aimed query, and it is the one the child understood instinctively: **depth**. The recursion of *why*.

In the 1930s, Sakichi Toyoda distilled it into a method that would remake global manufacturing. Confronted with any defect, ask *why*. Then ask *why* of the answer. Then again. Toyota's engineers found that by the fifth *why*, you are almost never where you started. You have descended from the symptom through the mechanism to the root. A puddle of oil on the factory floor becomes, five whys down, a purchasing policy that bought the wrong gaskets. The surface question and the deep question live in different worlds, and all the leverage is in the deep one.

This is the same descent that Aristotle mapped when he distinguished the four *aitia*, the four senses of "why" behind any thing, and the same one that Karl Popper enshrined at the heart of science itself: knowledge advances not by accumulating confirmations but by asking, relentlessly, *how could this be false?* The Socratic method, root-cause analysis, first-principles reasoning, the entire apparatus of rational inquiry: all of it is the disciplined, recursive application of *why*. Each answer is not a destination but a foothold from which to ask a deeper question. The value does not diminish as you descend; it compounds. The first *why* saves you a cleanup. The fifth *why* saves you the next thousand spills.

To question well, then, is to do two things at once: to choose, from a vast space, the query that most reduces the uncertainty that matters (breadth), and to refuse to stop at the first answer, driving downward through layers of cause toward the root where leverage concentrates (depth). Breadth and depth. The scalpel and the drill.

And here is the quiet catastrophe: human beings are magnificent at neither, at scale.

V. The human ceiling

We are, each of us, bounded. Herbert Simon won a Nobel Prize for naming it: *bounded rationality*. We do not have infinite attention, infinite time, or infinite working memory. We satisfice. We ask the few questions that occur to us, pursue them one or two whys deep, and move on because the next fire is already burning.

Expertise, paradoxically, narrows us further. The more we know a domain, the more its assumptions become invisible (psychologists call this the "curse of knowledge"), and the questions we most need to ask are precisely the ones our expertise has taught us not to bother asking. We are blinded by fluency. The permit has always required quarterly testing; why would you ask whether it should? The unit has always run this way; why interrogate it?

Confront a working environmental manager with the mathematics of their own opportunity and the ceiling becomes vivid. In the course of a busy year, such a person might ask their information systems a few dozen genuinely valuable, curiosity-driven questions, aimed at finding hidden savings, latent risk, or unclaimed operational room. A few dozen. Meanwhile, the number of *answerable, valuable* questions latent in a single company's permits, emission records, reported events, and the rivers of public environmental data surrounding its facilities runs into the tens of thousands, and grows every week as rules change and new data arrives.

I have argued elsewhere, in work on what I call the epistemic gap, that environmental protection itself is fundamentally gated by this ratio, by *questions asked* against *questions askable*. At human scale, the coverage is heartbreaking: a fraction of a fraction of one percent of the protective questions that could be asked ever are. The rest, the savings never found, the risks never flagged, the improvements never imagined, do not fail to exist. They simply go unasked. They sit in the data like unmined ore, not because the intelligence to answer them is missing, but because no one had the time, the standing, or the boundless curiosity to pose them.

This is the true bottleneck of the information age. We built systems of oracular power and then throttled them behind the narrow aperture of human questioning.

FIGURE 2 The epistemic gap. A field of five thousand points stands in for the tens of thousands of answerable, valuable questions latent in a single company's permits, records, and surrounding public data. The gold points are the few that a bounded human team asks in a year. The rest sit in the data like unmined ore.

VI. The silence of the oracles

Consider what an artificial intelligence system actually is today. It is an oracle of staggering capability, trained on the accumulated text of civilization, able to reason across law, chemistry, engineering, and finance in a single breath. And it sits, silent, waiting to be asked.

Every AI system you have ever used is a **pull system.** It does nothing until a human arrives, remembers it exists, finds the time, overcomes the intimidation of the blank prompt, and, hardest of all, *knows what to ask.* The intelligence is real, but it is dammed. The reservoir is full to the brim, and the only outlet is a human hand on a small valve, opened a few times a day, for the few questions that human happened to think of.

> ***The scarce resource is no longer the capacity to answer. It is the curiosity to ask, and the wisdom to ask well, at a scale and depth no human can supply.***

We have, in other words, moved the bottleneck without removing it. In the old world, the constraint was the answer: knowledge was scarce, expensive, locked in experts and libraries. We have very nearly solved that. The answers are now cheap and abundant. But in solving the answer, we exposed the deeper and older constraint, the one that was always there beneath it: **the question.**

A full reservoir behind a valve turned by a tired hand. That is the state of applied artificial intelligence in our

time. And it is entirely unnecessary.

VII. Teaching a machine to wonder

Here is the turn. If the constraint is the asking, then let the machine ask.

This is not science fiction, and it is not a trick of words. Curiosity is computable. The formal foundations have been laid over the past few decades, and they are beautiful. Jürgen Schmidhuber proposed a rigorous theory in which an agent is intrinsically *rewarded for learning progress*, for asking the questions and running the experiments that most improve its model of the world, that most compress its confusion into understanding. Deep reinforcement learning researchers built working versions: agents driven not by external reward but by *prediction error*, by the pull of the surprising, exploring their worlds out of something functionally indistinguishable from curiosity. The field of active learning formalized the same instinct for any learning system: do not wait for data to arrive; *choose the question whose answer will teach you the most.* And decision theory gave us the value of information, the exact price of an answer before you have it.

Assemble these and you have the blueprint for a machine that does what the pull system never could. It generates its own questions, thousands of them, across a space no human could survey. It estimates, for each, the expected value of the answer: how much uncertainty it would resolve, how much money or risk or ecological harm hangs on it. It pursues the promising ones, mining its own oracular depths for answers. It verifies ruthlessly, because a question is only worth asking if its answer can be trusted. And, this is the part that should make the hair stand up on your arms, **it learns which questions bore fruit, and asks better ones next time.**

That last property is the doorway to something genuinely new. A system that asks, evaluates the yield, and refines its asking is a system whose curiosity *compounds*. Each cycle, the questions get sharper and deeper. The machine begins to do at scale what the child does at the dinner table: it asks *why*, and then, unsatisfied, asks *why* of the answer, descending the ladder of cause. Automated breadth married to automated depth. Not one clever question, but a self-improving engine of inquiry that widens and deepens with every turn.

I want to be precise and honest about the limits, because truth is the whole point. This recursion is not a genie that grants infinite wishes, and it is not the runaway fantasy of an intelligence exploding into omniscience overnight. Its power is disciplined by a hard constraint, and the constraint is its virtue: **a question only compounds value if its answer is verified and grounded in reality.** A curiosity engine unmoored from truth would spiral into confident nonsense. The art, the entire engineering challenge, is to couple the generative wildness of question-making to a relentless, adversarial machinery of verification, so that only true, useful answers survive to seed the next generation of questions. Curiosity without rigor is delusion. Curiosity *with* rigor, compounding, is the most powerful epistemic engine ever built. The bound is what makes it real.

VIII. Seeker: the Environmental Curiosity Engine

I have spent years building environmental intelligence: systems that can answer almost any question about permits, compliance, emissions, and the living systems around industrial facilities. And I watched, with growing frustration, the reservoir sit behind the valve. Customers with oceans of latent value beneath their feet, opening the tap a few times a month. The intelligence was never the problem. The *asking* was.

So we are building the thing that asks. Its name is **Seeker, the Environmental Curiosity Engine.** It does not exist yet in its full form; it is under construction as I write these words. But it is the embodiment of everything this essay has argued, and it is worth describing precisely, because the design is not a leap of faith. It follows, step by step, from the physics and the mathematics laid out above.

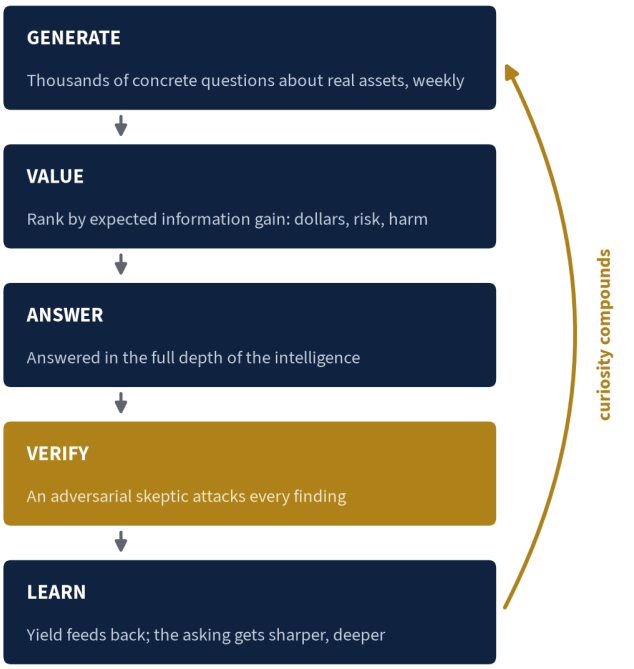

FIGURE 3 The compounding loop. Seeker's design couples generative question-making to adversarial verification: thousands of questions in, a handful of verified, consequential findings out, and a yield signal that makes next week's questions sharper and deeper.

Seeker is being built so that it will not wait to be prompted. Each week, it is designed to generate thousands of concrete, specific questions about a company's own assets. *Is this unit monitored more strictly than the rule requires? Is there stranded room in this permit worth millions in production flexibility? Did last week's rule change quietly create exposure at three of these facilities? Is a deadline about to lapse into a violation? Does the fenceline data, read against the reported emissions, hint at something no one has flagged?* It will route each promising question through the full depth of the environmental intelligence, compute the dollars and the risk, and then, the discipline that makes it trustworthy, turn a skeptic loose on every finding, trying to prove it wrong, discarding everything that does not survive. Only the verified survivors, ranked by how much they matter, would ever reach a human being. Ten thousand questions in; the handful that are true and consequential, out.

And then comes the part meant to make it an *engine* rather than a report. It is designed to watch which questions found value, to learn a given manager's assets, a facility's quirks, a regulator's patterns, and, the following week, to ask sharper questions, and deeper ones, descending the whys. *Why is this unit over-monitored?* becomes *why was that condition originally imposed?* becomes *what precedent at a sister site would justify relaxing it?* The engine's curiosity would compound, bounded always by verification, growing more valuable with every cycle, exploring a space of environmental questions so vast that no team of human experts, however brilliant, could ever traverse a meaningful fraction of it.

That is what it will mean to automate *why*. Not to replace the human's judgment. The human still decides, still acts, still owns the consequences. It is to hand the human a faculty they never had: an inexhaustible, ever-deepening curiosity, working their entire world, week after week. We have not finished building it. But for the first time in the history of mind, we know how. That, in itself, is the threshold this essay was written to mark.

IX. The environmental professional, reborn

Let me bring this all the way to the ground, to the desk of the environmental manager, because that is where the physics becomes a paycheck, a plant, and a river.

Picture a single real question of the kind Seeker is being built to ask, and watch it descend. It asks: *Are the heaters on this unit tested more often than the law requires?* The answer, pulled from the permit and the federal rule, is yes: quarterly, where the standard demands only annual. *Why?* Because a condition written years ago, in an abundance of caution, was never revisited. *Why not revisit it?* Because no human ever asked; the permit had always said quarterly, and quarterly it stayed. *What would it save?* Twelve unnecessary test events a year, each costing thousands. *And is there a clean path to change it?* A minor amendment, low-risk, reversible, fully evidenced. Five whys, and a line item on the compliance budget has become found money: hundreds of thousands of dollars a year, at a single unit, that was sitting in plain sight, unasked.

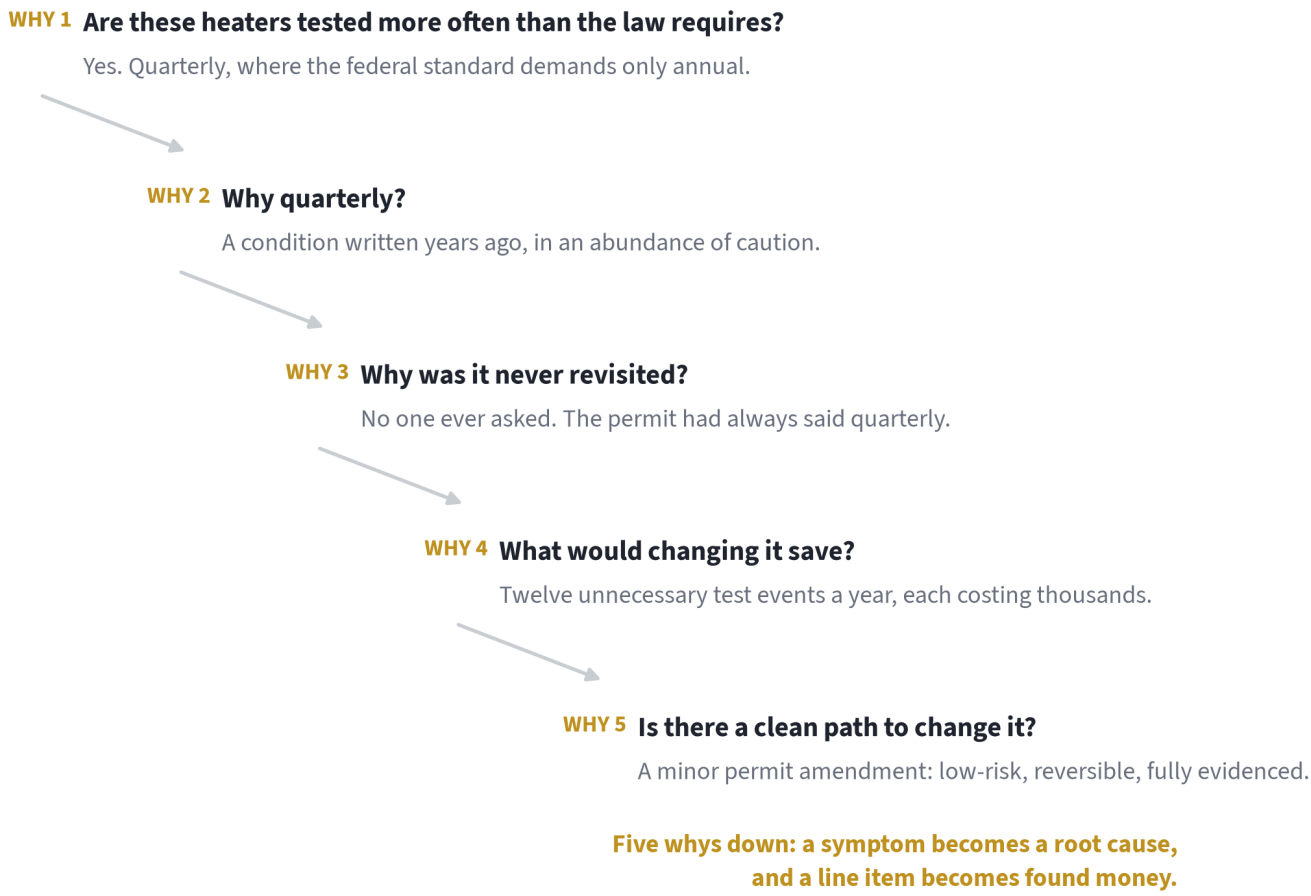

FIGURE 4 The descent of why. Five recursive questions carry a single permit condition from symptom to root cause. The surface question and the deep question live in different worlds, and all the leverage is in the deep one.

Multiply that across every unit, every permit, every facility, every week. This is the **financial** transformation such an engine is built to deliver: stranded permit headroom converted into production flexibility worth millions; over-monitoring right-sized; modifications structured to avoid enormous, unnecessary permitting burdens; offsets and penalties avoided. Value that was always present, gated only by the questions no one had time to ask.

It would be, equally, an **operational** transformation. An engine that asks *is a deadline about to lapse?* every day is one that need never let a renewal slip into a violation. The one that asks *when will atmospheric conditions favor this maintenance?* schedules the plant's hardest work into the hours the environment can best absorb it. The one that asks *is this pattern of reported events clustering toward an enforcement risk?* sees the storm while it is still on the horizon. Flexibility recovered, shutdowns avoided, surprises pre-empted. The plant runs looser and safer because someone is finally asking the right questions in time.

And it would be, most profoundly, an **environmental** transformation. Here the two faces of the coin, the financial and the ecological, turn out to be one. The same engine that finds stranded value finds real emission reductions at low or negative cost. The same relentless questioning that catches the compliance risk catches the fenceline benzene trend before it climbs, the methane the satellites see that the reports do not explain, the leak that has not yet become a violation or a headline. Reading the environment around a facility and asking, ceaselessly, *what is changing, and why?* is

not merely good business. It is the mechanism of protection itself. The manager who was cast, for a generation, as a cost center and a compliance cop is reborn as the person who quietly saved the company millions *and* made the air cleaner, because those turned out to be the same act, seen from two sides.

X. Observation is protection

Return, now, to the physics, because the environmental stakes are the deepest stakes of all, and they close the circle.

I have argued that a question is the cheapest lever in the universe: a near-free bit of information, at 10^{-21} joules, capable of steering the vastly more expensive rearrangement of matter. And I have argued that pollution is, at bottom, an information problem: a failure to ask, in time, the questions that would have configured the world differently. Put these together and a startling conclusion follows. **The cheapest way to protect the living world is to ask more of the right questions about it.** Not more scrubbers first, not more concrete first, but more *observation*, more *inquiry*, more measurement of the right things at the right time, from which cheap information the expensive machinery of remediation can be optimally directed. Maxwell's demon, operating now at the scale of a planet: observing, questioning, configuring the gate so that life, on the whole, is protected.

The physicist David Deutsch gave this optimism its rigorous form. In *The Beginning of Infinity*, he argues that the growth of knowledge is, in principle, unbounded; that every problem not forbidden by the laws of physics is soluble, given the right knowledge; and that the only thing standing between us and any achievable state of the world is our ignorance of how to reach it. Problems are inevitable; problems are soluble. This is not naïve techno-optimism. It is a claim about the structure of reality. Knowledge is the thing that converts the impossible-for-now into the merely not-yet-known. And knowledge begins, without exception, with a question. A civilization that keeps asking, and keeps learning to ask better, has, in the most literal sense, unlimited reach.

Turn this toward the Earth and the stakes become vertiginous. For four billion years this planet manufactured life and could not protect it. Five times, something catastrophic happened, and the living world had no recourse but to endure and begin again, because no organism, however magnificent, could ask *why* and act on the answer. A coral reef cannot model an asteroid. A forest cannot deflect one. We can. We are the first part of nature that has ever been able to ask why and use the answer to defend the rest. That is the precise sense in which a planet without minds is a planet already condemned, and a planet with them is a planet that finally has a chance. Both things are true at once, and maturity is the holding of both: we have done real harm to the living world on our way here, and we are the only thing in its entire history that has ever arrived in time to save it.

If you want proof that a question can reach from the inside of a human skull all the way to the machinery of the heavens, consider what happened on the 26th of September, 2022. A spacecraft named DART, built by people who had asked *what would it take to move an asteroid?*, struck a rock called Dimorphos at some fourteen thousand miles an hour and shortened its orbit around its companion by about half an hour. Then, in 2026, careful analysis confirmed something greater still: the impact had measurably shifted the entire two-body system's orbit *around the Sun.* For the first time in the roughly four and a half billion years of our solar system's existence, a mind had reached out and altered the path of a celestial body around its star. And it did so not with strength, for the spacecraft weighed less than a car, but with *knowledge*: a chain of questions asked and answered until a near-weightless bit of understanding could redirect a mountain of rock in the dark. There is the entire thesis of this essay, written once across the sky. The cheapest thing in the universe, the well-placed question, steering the most expensive.

PLATE II A chain of questions, written once across the sky. On 26 September 2022 the DART spacecraft, lighter than a car, struck the asteroid Dimorphos and shortened its orbit; in 2026, analysis confirmed the impact had measurably shifted the system's path around the Sun. The first time a mind altered the orbit of a celestial body around its star. Illustration: AI-generated for this essay.

For most of history this was hopeless, because the demon was human, and humans can ask so pitifully few questions. The epistemic gap, the chasm between the protective questions we could ask and the vanishing fraction we do, was in effect the measure of how much of nature we were failing to protect simply because we never got around to asking. To close that gap by hand would take geological time. To close it with a curiosity engine that asks ten thousand questions a week, and learns to ask better ones, and compounds: that is a matter of years, perhaps less. Automating *why* is not only an industrial tool. It is, potentially, one of the most powerful instruments of planetary stewardship we have ever conceived: a way to finally ask nature enough questions, fast enough, to configure our relationship with it toward flourishing rather than harm.

This is the vision that animates everything: an environmental intelligence that does not merely answer when asked, but wonders on our behalf, continuously, at machine scale, closing the gap between the world as it is and the world as, with the right questions asked in time, it could be.

XI. Coda: the child, restored

We began with a child asking *why*, and the world slowly teaching her to stop. Here is the turn that ought to move us: we are not building a machine to replace her curiosity. We are building one to *restore* it, and to give it a reach she could never have had alone.

For the first time, the most human faculty of all, the bottomless, tireless, slightly maddening drive to ask *why*, and then *why* again, and never quite be satisfied, need no longer be rationed. It can be amplified, extended, and let loose across every permit and every river and every question that matters, asking through the night while its human partners sleep, and waking them only with the few answers worth their judgment and their pride.

Earth spent four billion years unable to ask a single question in its own defense. We are the part of it that finally can.

The important thing, Einstein said, is not to stop questioning. For all of human history, we had no choice but to stop. We ran out of time, of attention, of the sheer stamina that curiosity demands. Now, for the first time, the questioning need never stop at all.

We are learning to automate *why*. And in doing so, we may discover that the oldest and most childlike act of intelligence, the endless, insistent *why, why, why?*, was all along the most powerful force we possessed for creating value, averting harm, and protecting the fragile, negentropic miracle of life on this planet. Earth spent four billion years unable to ask a single question in its own defense. We are the part of it that finally can. And, if we are wise enough to give that curiosity the reach of a machine, the part of it that finally *might.*

The reservoir is full. It is time to open every valve at once.

Jed Anderson is the founder and CEO of EnviroAI, where he and his team are building toward environmental superintelligence, and are now building Seeker, an engine designed to teach an artificial mind to ask why.

NOTES AND SOURCES

The thermodynamic cost of one bit, $kT\ln 2$, about 3×10^{-21} joules at room temperature, is Landauer's principle (Rolf Landauer, 1961): the resolution of the Maxwell's-demon paradox begun by Leó Szilárd (1929) and completed by Charles Bennett.

The measure of a question's worth as expected information gain traces to Claude Shannon (1948) and, for decisions, to Dennis Lindley's work on the value of information (1956).

"Negative entropy" as the thing life feeds on is Erwin Schrödinger's, in *What Is Life?* (1944).

The formal treatment of curiosity as intrinsic reward for learning or compression progress is Jürgen Schmidhuber's; prediction-error-driven exploration is due to Pathak and colleagues (2017). The Five Whys originate with Sakichi Toyoda and the Toyota Production System.

The claims about the unbounded reach of knowledge are David Deutsch's, in *The Beginning of Infinity* (2011).

The DART mission struck Dimorphos on 26 September 2022 at roughly 6.6 km/s, shortening its orbit around Didymos by about 32 minutes; the measurable change to the Didymos–Dimorphos system's orbit around the Sun was confirmed in 2026 (Makadia et al., *Science Advances*; NASA/JPL): the first time a human-made object has altered the orbit of a celestial body around a star.

The epistemic-gap figures are the author's own, from prior work on environmental intelligence. Bounded rationality is Herbert Simon's.