ESSAY · PREPRINT · V1

Resolution Is Not Leverage

On the great text explosion, the knowledge it did not create, and the quiet question of who will be permitted to move the Earth

Author Bryan Harris **Published** 2026-06-14 **Version** v1 **Hash** 945409e0adc5

Artificial Intelligence · Information Theory · Epistemology · Verification · AI Policy · Economics

Abstract

AI systems now generate text at a monthly volume rivaling the entire pre-AI human written corpus, yet multiplying throughput does not proportionally multiply knowledge. Distinguishing throughput from knowledge accrual—and resolution (seeing latent structure already entailed by existing data) from leverage (the capital, access, and authority to act on it)—this essay examines verification as the binding bottleneck, the accelerating enclosure of frontier capabilities, and why zero-cost simulation may paradoxically raise the economic return on genuine authenticity.

> "*Give me a place to stand...*"
> - Archimedes[1]

I. The number that should awe you, yet somehow doesn't.

Let's begin with a fact that ought to stop the room.

Sometime in the last year, the machines crossed a threshold that, stated plainly, sounds like the most important sentence anyone could write about the present era. Every month - give or take, depending on how you count - the world's artificial intelligence systems now generate a quantity of text roughly equal to the *entire written output of humanity* prior to their arrival. Not equal to a year of human writing. Not equal to the contents of a great library. Equal to the whole accumulated reservoir: every book, every letter, every contract and chronicle and scrap of digitized human language that existed before the first large language models learned to speak.

The arithmetic is worth seeing in the open, because it is the kind of claim that invites the suspicion of exaggeration, and amazingly - it survives the suspicion.

Researchers at Epoch AI estimate the total *effective* stock of human-generated public text - the usable written record, deduplicated and quality-adjusted - at around three hundred trillion tokens, a token being the unit a model processes, worth on average about three-quarters of an English word.[2] That figure is "the Human Corpus", the reservoir: the sum of the richest curated set of what we wrote across all of recorded, digitized history, the substrate from which the first models were built.

Now against that, set the present rate of machine output. In early 2024, the head of OpenAI noted that his company's systems alone were producing on the order of a hundred billion words a day, and observed in the same breath that all the humans on Earth together produce something like a hundred trillion words a day - placing the machines, at that moment, at roughly a tenth of a percent of human output.[3] That ratio did not hold for long. By late 2025, OpenAI's programming interface alone was processing on the order of eight to nine trillion tokens per day.[4] By the spring of 2026, Google reported processing more than three *quadrillion* tokens per month across its products - up from under five hundred trillion a year earlier.[5] Chinese systems, by independent tracking, had reached something like a hundred and eighty trillion tokens a day, a roughly eighteen-hundred-fold increase in two years.[6]

Those headline figures measure *throughput* - tokens flowing through the systems, much of it the model reading rather than writing, much of it inflated by the extraordinary token-appetite of video generation and chained software agents. Strip that away and ask only how much human-readable *text* the machines are emitting into the world, and the honest estimate, backed out of the reported pieces, lands somewhere on the order of ten trillion tokens a day. Hold that number loosely; it is an inference, not a published figure. But its order of magnitude is robust, and the conclusion it forces is indifferent to a factor of two or three in either direction. At ten trillion tokens a day, the machines reproduce the entire effective reservoir of human writing - the three hundred trillion tokens it took our species the whole of its literate existence to accumulate - in about a month. Halve the estimate and it is two months. Triple it and it is ten days.

No matter how you slice it, over the span of a single year, by any of these accountings, the world now generates the equivalent of its pre-AI human-written inheritance corpus dozens of times over. By comparison, less than four short years ago, that same quantity was a once-in-a-civilization total.

This is the statistic. And the natural response to it is awe - the sense that something of overwhelming magnitude has happened, that we are living through an explosion of mind, a Cambrian bloom of knowledge, a step-change in what our species can know.

I want to suggest, across the pages that follow, that the awe is misdirected - not wrong, exactly, but aimed at the wrong feature of the event - and that the more interesting and more consequential truth is hiding directly behind it.

Here is the question that *should* trouble you about that staggering number. Imagine for a moment that the entire universe - every atom, every distance, every wavelength, and - crucially - every ruler you might use to measure any of it (including the underlying structural mathematical relationships which makes the forces what they are and behave the way they do) - were to spontaneously DOUBLE in size, instantaneously, quietly, and uniformly, tonight.

How would anyone ever know?

Common sense intuition tells us this could be catastrophic!

However, in practice: you would wake, reach for a meter stick (which has also doubled), and measure everything as *exactly* the size it was before. The light from distant stars would arrive unchanged, because the stars and the space between them and the very yardstick of the speed of light would all have scaled together. There is, therefore, no vantage point (that we know of or would theoretically even be capable of perceiving) outside the universe *from which* to notice it had grown, and no fixed reference inside it against which the growth could register.

In the end, here, the change - *however vast* - would be undetectable.

Now, don't get me wrong here: I mean it would not merely be *in effect* unnoticed, but rather that it would be - *in principle* - *UNNOTICABLE*. Put differently: a perfectly uniform rescaling is - in the only sense that matters - perceptibly indistinguishable from no change at all.

This is actually an old observation. The mathematician Henri Poincaré used a version of it to argue that absolute size is not an independent physical fact, only a relational one.[7]

Now hold that thought against the explosion of text. We multiplied our output by a thousand and then by a thousand again. But what if the thing we multiplied was *both* the signal and the noise, in equal measure - so that the proportions of the ratio between them, the only quantity that actually governs how much we can know, never actually moved very much at all? What if the great text explosion is the intellectual equivalent of a universe that doubled overnight: a change of breathtaking magnitude in a dimension we cannot perceive, and no change whatsoever in the dimension we can - the ratio of data and information to knowledge and understanding?

That is the question this essay exists to explore. The answer I present here turns out to be more surprising, and more useful, than either the awe or the dismissal it initially tempts us toward would yield, IMHO. To reach it we will

need to separate two things that the headline number deliberately fuses - *information* from *knowledge* - and then, having separated them, to discover that the distinction opens onto a second one, sharper and more consequential, between two things we habitually confuse: *seeing* and *influencing*. Between, that is, resolution and leverage.

II. What grew was information, not knowledge.

The first cut is the one that our staggering headline number is engineered to hide. We must distinguish the *volume* of text from the *knowledge* it carries, because they are not the same quantity, and the overwhelming majority of the constituent resulting exploded enumeration lives exclusively on one side of the divide.

Consider what the original reservoir of "the Human Corpus" actually was. Those three hundred trillion tokens of human writing were not three hundred trillion tokens of *redundant* writing. They were, in the main, accumulations of distinct human acts: an observation recorded for the first time, an argument advanced against another, a story told that had not been told in quite that way, a measurement logged, a law set down, relitigated, overturned, a discovery published, an invention claim staked and elucidated. The reservoir is sparse in the specific sense that matters - it represents an enormous quantity of semantically *non-redundant* human cognition, each piece purchased at the cost of a human life-hour of attention, accreted through the rise and fall of various epochs in coherent dialog with itself across millennia. It is not that every sentence in it is profound; after all: *most* human writing is, in fact, mundane.

That's actually completely FINE.

The fact that most writing isn't novel or impactful doesn't mean it isn't important, consequential, meaningful, and useful. It is that the reservoir, taken as a whole, is the trace of a vast number of *separate* engagements between minds and the world. Not everything which exists is in "the Corpus" and not everything that's in the Corpus is completely appropriate, essential,

true, or indispensable on principle. That's okay in the context of this conversation (not taking that on here) - but a meaningful and important differentiation nonetheless as we will examine more deeply later.

After all - the sheer volume of coherent text needed to make all this work at all means that the "Human Corpus" *couldn't* just be the best of the best of what humanity had to offer. As much as "the Corpus" needed greats and classics the likes of Aristotle, Copernicus, Vonnegut, Twain, Bohm and WikiPedia - it *also* needed all of Reddit, Stack Overflow, erowid, and 4chan.

Now consider what the machines, our LLMs, emit. A language model is, at its mathematical core, a model of the probability distribution of the text it was trained on. Its task - the literal objective it was optimized against - is to predict the most likely next token given everything that came before. The likely token is, by construction, the *least surprising* one. This is not a flaw that crept in; it is the definition of the device.

And so the text a model produces is, in the precise sense given to the word by Claude Shannon when he founded information theory in 1948, *low in information relative to its source*.[8] Shannon's insight was that the information content of a message is its surprise -its unpredictability given what you already know. A message that merely tells you what you could already have predicted carries no information. A model trained to emit the predictable token is, therefore, a machine for producing text that is maximally redundant with the corpus that made it. The flood is not an addition to the reservoir. It is, overwhelmingly, the reservoir's echo - paraphrase, summary, translation, restatement, recombination, the same ideas rendered a billion times in slightly different words.

This is not a fringe complaint; it is increasingly the consensus reading among people who think carefully about what these systems do and understand deeply how exactly they do it. The machine-learning literature has long held that current models *generalize within the distribution of their training data* - they interpolate - and struggle catastrophically to adapt to situations genuinely outside it - they do not, in the strong sense, extrapolate.[9] Writers

reaching for the same idea in plainer language put it sharply: from the model's own perspective, "if it looks right, it is right," because conformance to anything beyond the patterns of its training is simply not a quantity it has access to.[10] And the economists who study how ideas are made have a still older frame for it: new ideas are recombinations of existing ones, a point the economist Martin Weitzman formalized decades ago, and which recent work extends to AI directly - modeling these systems as filling in the space *between* known points rather than reaching past them, raising the density of what is already known while, at sufficient scale, actually *reducing* the novelty of what gets explored.[11]

To be completely clear and fair, this is by no means a secret or by any means beyond the understanding of the researchers at frontier labs, who have been saying and endeavoring all along to improve grounding and building out the infrastructure and harness to provide models with increased modalities and more robust real world sensors and oracles by which they can make the measurements, do the proving, and validate the feedback that *would* allow them to procure, analyze, and meaningfully react to the novel information, the observations that *would* allow it to create genuinely new, valid, useful inference, deduction, generalization, and true extrapolation. This is why we now have tools, skills, and agentic harnesses to bridge that gap.

We must, however, in fairness, now run the same signal-to-noise filter across the human side of the ledger, because the comparison is otherwise dishonest. The three hundred trillion tokens of "effective" human text already excludes an ocean. Humanity has always produced staggering volumes of writing that carry no knowledge in any meaningful sense - the text message to one's mother, the fax between two offices, the printed receipt, the logistics email, the small talk of a species that writes constantly and mostly transactionally. None of that was knowledge; almost none of it entered the reservoir; and rightly so. So when we say the machine flood is "mostly not knowledge," we are not holding it to a standard humans clear. We are observing that *most text, human or machine, has always been throughput* - communication, not refinement - and that the explosion is, in its overwhelming bulk, an explosion of throughput.

Looking at it in this posture gives us a clean conceptual apparatus, and it is worth stating in full because the rest of the argument rests on it. There are two distinct quantities, and they must be tracked separately and filtered identically on both sides:

The first is **throughput**: the gross volume of text moved through the world, the great majority of it transactional, ephemeral, and impermanent. A model's chatbot reply and a person's grocery list belong to the same category here. Throughput is what the headline number measures, and throughput is what exploded by a thousandfold and then again.

The second is **knowledge accrual**: the thin seam, running through the vast bulk of throughput on both the human and the machine side, that actually tightens our collective grip on reality - the genuine observation, the verified result, the idea that lets us predict or do something we could not before.

The relationship between these two quantities is the whole game, and the explosion of text bears on it in a way that the awe completely obscures. We did not multiply knowledge accrual a thousandfold. We multiplied throughput a thousandfold, and the knowledge seam - the compressing, reality-predicting residual - barely moved. The signal did not scale with the noise. As one essayist surveying the verification problem put it with admirable economy, "noise scales, but signal does not."[12]

And so, halfway through the first movement of the argument, we appear to have confirmed the deflating possibility. The universe doubled. We added a thousandfold to one term of the ratio and almost nothing to the other, which means the ratio of knowledge to text did not hold steady - it *collapsed*. We are not drowning in knowledge. We are drowning in the echo of knowledge, and the noise of the echo is getting louder, longer and more layered while the voice of signal, comparatively, has largely stayed the same.

If the essay stopped here, it would just be one more entry in the growing genre of disappointment; another among many in the chorusline of the "AI is just a stochastic parrot, it only remixes, nothing new under the sun" genre.

Personally, I find that genre feels extremely lazy, and - to me - rings totally flat & hollow. The reason *why* I think that interpretation is completely wrong is the turn this essay now takes.

III. The first turn - "what we already knew" is an entire continent.

The error in the deflationary reading is a single smuggled word. To say the machines merely re-derive "what we already knew" is to let the word *merely* carry an enormous and false implication: that what we already knew is small, or finished, or fully possessed. It is none of these things. The territory of the already-implied - the knowledge entailed by the facts in our hands but not yet worked out, connected, or understood - is not a footnote to human knowledge. It is a continent. And we have only just acquired the instruments to chart it.

Before we sail for that continent, we should pause on an apparent paradox, because the rest of the argument depends on dissolving it - and it dissolves far more easily than it looks. The previous section insisted that the machine's output (absent grounded verifiable new observations) is *redundant*: it adds nothing the corpus did not already contain. This section now, seemingly contradictorily, insists that the same output is simultaneously still massively *valuable*: a whole new world. Both can be - and I assert are - true at once, and the manner in which they are both true is something you already understand completely, in countless settings where it has never once confused you.

For one, consider a solved Sudoku. The completed grid contains not one particle of information that was not already fixed, with absolute determinism, by the handful of given digits and the rules of the game. The solution adds *nothing*; it was entailed, in full, from the first moment the puzzle was set. And yet no one has ever looked at a finished grid and pronounced it worthless on the grounds that its answer was "already there." Of course it was already there. That was never the point. The value was never in the addition of new information to the system - there is none to add - but in the *labor of locating* a single determined configuration inside a space of possibilities far too vast to search by hand. The answer was always implied. It simply was not *had*. The

machine is, in this exact and literal sense, a solver: it adds no information to the corpus that trained it, and it hands you, all the same, answers you could not previously reach. (The same fact, in formal dress, is why decompressing a file adds precisely zero information and remains the entire point to anyone lacking the means to unzip it.)

And if you want the feeling of the thing rather than its formalism, you own that too, from every story that ever turned on a revelation hiding in plain sight. Watch a film with a buried twist a second time and you will find that its ending added no frame the first viewing had not already shown you - every clue sat on the screen, fully visible, the whole way through. Nothing was withheld. And yet the revelation, when it first lands, reorganizes everything you saw. The identical footage carries almost no information to the viewer who already knows, and an entire world of it to the one who does not. Same text; opposite worth; the difference living entirely in who is doing the seeing and where this scarcity of knowledge creates disproportionate value - especially where that knowledge creates leverage.

The right historical image is the discovery of the Americas, and the image is exact in a way that rewards pressing on.

The Americas were not *created* in 1492. They were already there - two continents of mountain and river and forest and people, entailed by the same round Earth that Europe and Asia had always stood upon, requiring no act of creation, only an act of *charting*. And yet their incorporation into the known map was not a minor administrative update to human geography. It roughly doubled the world that the connected civilizations could reckon with, and it did so not by inventing new land but by *connecting* land that was already there to the existing system - and the connection released centuries of consequence, the flow of crops and materials and ideas in both directions that historians call the Columbian Exchange. The value was not in the newness of the territory. The territory was old. The value was in its *connection* to everything else, which made what it contained tradeable against what was already known.

This is precisely the character of the territory the machines have opened. It is not, in the main, territory that was never there. It is territory that was always

implied by what we had, and that we simply did not have on the map - because charting it was too expensive, too slow, too far beyond the reach of any single human mind to traverse. A proof that follows by a long chain of steps from lemmas already proved was always *true*; it sat in the territory of the entailed, waiting. An interaction between two drugs latent in clinical data already collected was always *real*; no one had connected the records. The relationship between two findings in two fields that no single specialist reads both of was always *there*; the continent was uncharted not because the land did not exist but because no one had sailed between them.

To make the scale of this vivid, abandon the cartographic metaphor for a moment and take up a numerical one, because it captures something the maps cannot. Picture the discrete facts we already possess and know well as the integers - clean, countable, the things we can name. Between any two integers lies an *uncountable infinity* of real numbers, a continuum of values denser than the integers could ever be. That is the relationship between the facts we hold and the structure that lies *between* them, waiting to be resolved. Even if the supply of genuinely new integers - wholly novel facts pulled fresh from contact with reality - were running low, the density of derivable structure between the integers we already hold would remain, for all practical purposes, inexhaustible. The Americas were not a new integer. They were the impossibly dense reality lying in the gap between the integers Europe and Asia already knew.

A necessary refinement, because the metaphor can be over-read: not every point in that continuum is knowledge. Most of the real numbers between three and four are noise; the value 3.7000001 carries nothing. The knowledge-bearing structure in the gap is itself a sparse landscape with its own mountains and deserts. But sparse is not empty, and what hides among the noise are the landmarks of disproportionate worth - the *transcendental constants* of the space, the π and the e, the points that are not integers, that sit exactly in the gap, and that are load-bearing for everything connected to them. π was always sitting between three and four, precise to infinite depth, structurally essential, long before anyone could locate it to forty places. It was never missing. It was *unlocatable*. And what changed when we gained the

means to locate it was not reality - reality already contained π - but our instruments.

The history of instruments is the truest version of this story, and it supplies the essay's cleanest images. When Galileo turned a telescope to the night sky in 1610 and saw four points of light circling Jupiter, he did not create the moons of Jupiter. They had been orbiting, faithfully, radiating their reflected light at the Earth, for the entire history of the human eye - which had simply lacked the resolution to separate them from the planet's glare. When Antonie van Leeuwenhoek ground his lenses and saw, for the first time, the teeming "animalcules" in a drop of pond water, he did not summon a microbial world into being. It had been there, in every drop, through every plague and every harvest, invisible only because no instrument could resolve it. In both cases the world did not change. The *resolution* changed. And the consequences - for cosmology, for medicine, for the whole human self-understanding - were civilizational.

The machines are an instrument of exactly this kind, pointed not at the sky or the water but at the corpus - at the vast, real, already-present territory of what our accumulated record entails but no human has had the reach to derive. And the proof that there is genuine treasure in that territory, not merely the comfort of metaphor, comes from the domains where we can already check the machine's work. The story of *Moneyball* is the homely, perfect example: the knowledge that on-base percentage predicted runs better than the statistics scouts had revered was sitting in the box scores the entire time, in plain numbers anyone could have added up, for a century. It took a new way of *seeing* the same data - not new data - to surface an edge that was always latent in what we already had. The alpha was on the table the whole time. The machines are, among other things, the most powerful engine ever built for finding the alpha that was already on the table.

So the deflation was premature. The flood is, indeed, mostly not new knowledge in the sense of new integers freshly drawn from reality. But it is *resolution* - a thousandfold increase in our power to see, derive, connect, and

chart a continent of structure that was real and present and unmapped. That is not nothing. That is the discovery of a new world.

And resolution, to borrow the oldest line about leverage, is a place to stand.

IV. The hinge - Resolution is not leverage.

> "*...and I will move the Earth.*"
> - Archimedes (still)[1]

Here is the full sentence, and here is the hinge of everything.

The mechanician's boast, traditionally attributed to Archimedes and transmitted to us through Pappus of Alexandria centuries after the fact, is usually quoted for its bravado. But read it as the precise conditional it is. Archimedes did not say *give me a lever* and I will move the Earth. He said give me a *place to stand* - and *then* I will move it. He named two things, and made the moving conditional on having *both*. The lever is the mechanism. The place to stand is the located point, the fulcrum, the *where*. And the deepest thing the sentence knows is that the second without the first is useless: a man with the longest lever in the cosmos and no ground beneath his feet moves nothing at all.

This is the distinction the whole essay has been building toward, and it is the one our awe at the great number most completely misses. *Resolution is not leverage.*

Resolution is the acuity of perception - the power to see, finely and truly, the structure that was already there. Leverage is mechanical advantage - the power to *move* more than your effort alone should permit, to change a ratio, to act with effect disproportionate to the force you apply. They are not the same thing, and the difference is not pedantic. Seeing a screw at higher magnification does not turn it. A perfect map of a mountain pass does not carry your cargo over it. The telescope showed Galileo the moons of Jupiter; it

did not let him touch them. Resolution reveals. Leverage acts. And the great text explosion, properly understood, was overwhelmingly a gain in *resolution* - we can now see the continent - and only incidentally, and conditionally, a gain in *leverage*.

So we arrive at the question that actually matters, the one that survives all the awe and all the deflation: granted that we have gained resolution and not, directly, leverage - *Can the resolution be converted into leverage?* Is seeing-where-to-stand sufficient to let us move the Earth, even though seeing is not, in itself, a lever?

The answer is yes, conditionally, and the condition is everything.

Resolution becomes leverage exactly when it reveals a point where leverage was already *latent* - a place of concentrated sensitivity in the structure of the world, a keystone, where a small applied force produces a large effect. A lever applied at a random point does nothing. The same lever applied at the keystone moves the building. The force and the lever-arm were available all along; what was missing was the knowledge of *where to stand*. In a world rich with such hidden high-leverage points - and the derivable continent, the space of transcendentals between the integers, is exactly such a world - resolution is very nearly the whole game, because the leverage was sitting in the structure, waiting only to be aimed at. To find π is already to be able to use it. To resolve the keystone is already, in such domains, to be able to move what it holds up.

But the conversion is not automatic, and its failure modes are the most important and least appreciated part of the entire picture. There are two, and they define the real shape of the world we have entered.

The first is the problem of *faithfulness*, and it is where the resolution metaphor turns dangerous. To chart a continent, you need not merely the desire to map it but an instrument you can *trust* - a way of telling, of any point you have resolved, whether it is real land or a painted illusion. And here we confront the defining hazard of machine-generated resolution: a model that emits the most plausible continuation, untethered from any check against reality, does not leave the unknown unmapped. It maps it *confidently and wrongly*. It hands

you a beautiful, detailed, internally coherent chart of a coastline that is not there. This is what the failure we call "hallucination" actually is - not noise, not gibberish, but fluent, convincing interpolation in the gap that happens not to be true. The specific peril of the new instrument is not that we still cannot see; it is that we now see something, vividly, and it may be a fiction.

Jorge Luis Borges wrote the definitive parable of this condition decades before the machines, in his story of the Library of Babel - an endless library containing every possible book: every true account of the universe and, alongside each, an unimaginably vaster multitude of volumes identical in form and authority but false in every particular, the plausible lies outnumbering the truths by ratios beyond comprehension.[13] The horror of Borges's library is not scarcity. It is the opposite: total abundance rendered worthless by the absence of any means to distinguish, among infinitely many flawlessly-formatted volumes, the one that is true from the endless ranks that merely look it. We have built, in a real sense, an engine that writes in that library. And the lesson is exact: access to all derivable text is worth precisely nothing without the means to verify which of it corresponds to reality. False resolution is not a smaller version of true resolution. It is *negative* leverage - it invites you to apply your force, with full confidence, at a keystone that is not load-bearing, and bring the wrong wall down on your head.

This is the point at which a growing chorus has, to its credit, begun to arrive. The most perceptive recent commentary on AI in the workplace has converged on a single recognition: that the cost of *generating* output has collapsed toward zero while the cost of *verifying* it has not, and that this asymmetry, not raw capability, is the binding constraint. "Verification is the new scarce input," as one writer put it, naming a "verification cost disease" in which the demand for trustworthy judgment grows faster than any efficiency can supply it.[14] Another names the precise danger that makes verification so costly: "the wrong answers look right," so that catching the subtle, plausible, confident error takes more expert attention than writing the thing correctly would have.[15] Institutions have begun to notice the structural consequence - that verification skill is becoming scarce precisely as the entry-level work through which people once acquired it is automated away, so that the ladder

by which a society trains its verifiers is quietly being pulled up.[16] This conversation is real and it is correct as far as it goes. The argument of this essay is that it does not go far enough - that the verification problem is not a workplace inconvenience but the hinge on which the entire value of the text explosion turns, the exact mechanism that determines whether our new resolution becomes leverage or becomes a confidently-charted fiction.

Which brings us to the second failure mode, and it is colder than the first.

V. The second turn - who gets to move the Earth?

There is a failure mode of resolution darker than false maps, and it is the one that closes the loop of this essay back onto the world as it actually is. It is this: resolution can be perfectly true - the map can be faithful, the keystone correctly located, the place to stand precisely identified - *and you can still be unable to move the Earth, because you do not possess the lever, or the ground to plant it on.*

Recall Archimedes' condition. He asked for two things. Resolution gives you the first: it reveals where to stand. It gives you *none of the second*. Seeing exactly where the leverage lies confers no lever and no standing-ground. Knowing precisely which gene to edit is not the capacity to edit it. Locating the exact market inefficiency is not the capital to exploit it. Charting the pass is not the ship, the crew, or the right to make the crossing. Resolution tells you *where* the power is. Whether you can *realize* it depends on something resolution does not grant and cannot manufacture: access, capital, capability, permission - the lever, and the place to stand.

And here is the structural fact that makes this the most important sentence in the essay. **Resolution is democratizing. The realization of leverage is concentrating.** These two trends are running in opposite directions at the same moment, and the gap between them is the defining inequality of the age we are entering.

The resolution is becoming universal. The instruments that chart the derivable continent - the models that find the alpha latent in existing data - are diffusing to nearly everyone, cheaply, often freely. We are all, increasingly, able to *see* where the gold is. But the means to *take* it - to convert a true map into actual mechanical advantage in the world - is consolidating into fewer hands, faster, and behind higher walls, than at any prior moment in the technology's short history.

The temptation here is to reach for the oldest piece of gold-rush wisdom and stop: in a gold rush, the saying goes, don't dig for gold - sell shovels. The ones who got rich in California were not, mostly, the prospectors; they were the merchants who sold picks and pans and denim to the prospectors. Translate it to the present and the advice writes itself: don't build with AI, build the *tools* for building with AI, sell the shovels to the diggers.

This advice is, in the present case, a trap - and it is essential to say exactly why, because the trap is subtle and nearly everyone walks into it.

The "shovels" that dig the actual gold in this rush - the foundational compute, the frontier models themselves, the genuine verifiers that turn the derivable continent into realized advantage - are not a concession open to the public. They are *already claimed*. They belong to a handful of firms with capital positions measured in the hundreds of billions: the makers of the scarce silicon, the hyperscale cloud providers, the few labs at the frontier. When a person of ordinary means hears "sell shovels," they imagine themselves included in that opportunity. They are not. The shovel concession that matters was staked before they arrived. What remains available to an independent operator - to a small business, even a mid-sized industry player - is the *attachment*, the *handle*, the peripheral accessory that clips onto someone else's shovel: a workflow wrapper, a vertical application, a thin layer of convenience atop a model and a cloud and a compute supply they neither own nor control and could not replicate. That work can be real and even valuable. But it is not selling the shovels that get the gold. It is selling shovel-*holders* to people renting shovels from the three companies that own the mine. Even the escape hatch - the supposedly democratic move of selling

tools rather than digging - is gated at the one layer that confers actual leverage.

So the gold rush metaphor earns its place in this essay only by breaking its own cliché. The lesson is not "sell shovels." The lesson is that *the shovels are already gated, including from the would-be shovel-sellers,* and that the cheerful folk wisdom about tool-makers getting rich quietly assumes an openness at the tool layer that this particular rush does not have.

If you doubt that the gating is real, that it is hardening, and that it now reaches with the full weight of the state, you need look no further than the event that prompted this essay's writing. In June of 2026, the United States government, citing national-security authorities, issued an export-control directive ordering one of the leading AI laboratories to suspend all access to its two most capable models for any foreign national - anywhere in the world, including the company's own non-citizen employees - over a vulnerability it described verbally and never disclosed.[17] Because a nationality filter cannot be cleanly enforced across a shared service, the company's only compliant move was to disable the models for *everyone*. This was the first time a frontier model had been pulled from distribution not by its maker's choice but by the state's command - the deemed-export doctrine, by which releasing a controlled technology to a foreign national counts as an export even when no border is crossed, pointed for the first time at a deployed artificial intelligence.

Let us neither, however, look past nor get distracted by the incredible irony at play here. This administrative assault on Fable & Mythos was obviously not part of Anthropic’s corporate roadmap, nor was it driven by the altruistic mandates of its novel governance structure. Conversely, it inflicts immediate, severe operational damage and a significant financial hit. Forced to abruptly reset weekly limits mid-cycle to manage the cascading PR backlash within their customer base, the company was compelled to sacrifice real revenue and the lucrative overage fees they would have otherwise collected for the month. While some might dismiss these immediate losses as comparatively marginal, the deeper reality remains unassailable: this is a raw, non-compensable extraction of value from a private corporation, executed through

unilateral, opaque, and unprecedented governmental action under the guise of safety by extraordinary fiat - NOT a "marketing stunt".

Even considering the merit and value behind the inevitable "hype uplift" attendant with this incident framing these Anthropic models as *SO* fundamentally capable that they constitute a national security risk - a mild corporate silver lining at best - the reality is that this raw, unprecedented interference is inherently and deeply disruptive to the sustainability, predictability, and profitability of the entire technology sector sets a new baseline precedent to contend with which is completely arbitrary, unilateral, and unequivocal as of just now because of this.

This action was executed with a level of aggressiveness, frivolity, and absolute, unilateral lack of transparency that so blatantly contrasts with its typically stated posture so as to directly expose the mounting hypocrisy and apparently arbitrary & power-driven capriciousness of the state at present. Consider the administration’s own past much-professed stance railing against technological regulation in any way shape or form whatsoever as a concept in general. The public has been consistently fed the divisive dogmatic line that we "couldn't afford to strangle this promising nascent technology" with ANY regulatory oversight whatsoever - technology the state has otherwise been vigorously, structurally and financially propped up both publicly by the state and privately by industry behind the scenes and in the headlines.

And yet, the ultimate, biting poetry of this enclosure belongs to the target itself. The state’s legal and rhetorical justifications for this hostile shutdown seem to be composed almost entirely of talking points that Anthropic itself has championed loudest over recent years. It was Anthropic that beat the drum most fervently, insisting that the technology in general, and their own models in specific, are so uniquely advanced that they present an inherent danger to human survival. It was Anthropic that aggressively sold the narrative that we are locked in a zero-sum, nationalist competitive nation-state race in existential peril - a narrative meant to justify whatever extreme steps are needed to maintain a lead.

That exact drumbeat has now looped back to bite them firmly in the assets. The very existential panic they weaponized to cultivate market dominance and invite state alignment has now been turned back on itself as a tool of summary expropriation.

The mechanism on display was unambiguous: the most capable instruments of resolution are now treated as controlled munitions, and the authority to decide who may stand where, and move what, is being claimed at the highest level. The wall is not a forecast. It is under construction, and a government has just demonstrated it will reach inside a private company to add a course of bricks.

Now, I can hear the objection from the back of the room - that to call this a power grab is itself the paranoid move, the tin-foil reading, the conspiracy theory dressed up as analysis. I take the objection seriously. There are many possible worlds in which that suspicion would be exactly the dangerous kind of thinking that sees a cabal behind every closed door. I am not here to sell a narrative of secret cabals. I am here only to read the tea leaves as they actually lie, in the actual cup, at this actual moment - and to ask what the leaves, soberly read, are telling us. Read them, then, on three points.

First, consider the *evidence* offered for the danger that justified the seizure. There is none. Or rather: there is a verbal gesture toward a vulnerability - undisclosed, unquantified, unaccompanied by any demonstration the public can examine. The threat was asserted and the door was shut, in that order, with nothing in between that a citizen could weigh. This is not a new pattern; it is among the oldest in the book. Invoke a danger the people cannot see, declare that only the surrender of some freedom can hold it back, and the surrender follows almost on its own, because no one wishes to be the fool who refused. It is the boogeyman in the closet, scaled to the level of the state - a tale as old as governance itself. What is striking here is how *little* the pattern was even made to work compellingly.

Whatever one believes about the great rights-curtailments of the recent past, they at least came after something concrete and visible: a building had actually fallen before a decade of war and a permanent reordering of civil life

were sold to the public on its account. Here, there is no building. There is no event. There is a closed-door assertion with nearly zero proof, transparency, deliberation, or nuance and a switch thrown unilaterally without recourse without warning and without qualification or corroboration.. We are being asked to accept a precedent of enormous reach on the strength of a *maybe*, a *perhaps*, an undisclosed *what-if*. And when the evidence is THAT thin and the action is THAT large - the honest, good faith reader of any sophistication stops asking what and how significant the danger is and starts asking what the action itself *accomplishes*, what prompted it, and who benefits.

Which brings us to the second leaf. The ready reassurance some will cling to that that nothing important has been taken here - that the public still has its free tiers, its cheap tiers, its open-weight models, that the technology as a whole is not being gatekept for reasons of power, control, and profit - only its *dangerous* edge for reasons of allegedly justified, yet completely unqualified “safety”. But recall the distinction we have already drawn between what is dangerous and what is valuable, and look hard at which one of these the state actually moved to enclose. It did not reach for the export-control apparatus - the machinery held in reserve for fissile material, weapons systems, the most strategically decisive technologies a nation possesses - in order to wall off something marginal, something overrated, unproven, unreliable, and unsustainable. You do not declare a thing a controlled munition unless you have concluded that it genuinely confers real and consequential power.

The action *is* the assessment.

In merely reaching for that authority, the government has told us, more candidly than any corporate press release ever could, what it believes these particular models *are*: not another incremental chatbot, but a phase shift into genuinely consequential capability - implying implicitly without saying: the first tier of this technology that - as this behavior belies - it considers to be for all practical purposes the thing we have spent seventy years in pursuit of under another three letter acronym I dare not this discussion further by invoking here. What was fenced off was not the dangerous part as distinct from the valuable part. The danger *is* the value. They are the same edge - and

it is precisely the edge that matters which has been placed on the far side of the wall.

And now the third leaf, the one that makes the whole shape legible. We should ask, plainly, the question almost no one thinks to ask: *Why was any of this ever free in the first place?* Why did the most powerful cognitive technology ever built arrive not behind a velvet rope but flung open to the entire public, at no charge, for years - a generosity with essentially no precedent in the history of strategically valuable inventions?

The answer is not generosity, and it never was. The answer is in the scaling laws, and in the arithmetic with which this essay began. These systems could not have been built a century ago, or even a generation ago, no matter how brilliant the engineers - because the principle that governs them, that capability grows only when data and computation scale *together*, means there is a threshold of accumulated text below which the whole enterprise simply does not ignite. We crossed that threshold only recently, and only on the strength of the entire digitized inheritance of human writing. [2] But the first models drank the easy reservoir dry, and to climb higher the laboratories needed something the static reservoir could not supply: more data, fresher data, and above all data about how human beings actually *use*, correct, prefer, and respond to the machines themselves.

There was exactly one way to gather that at the necessary scale.

Give the technology to everyone, for free, and let the use of it become the harvest. The free tier was never a gift. It was an instrument of extraction - the most efficient data-gathering operation ever conducted, with the entire connected species as its unpaid and largely unwitting workforce. Return to the figure from the opening: that we now generate, every month or so, a volume of text equal to all of human writing before the machines. [3] A great share of that torrent is the harvest itself - humanity laboring, billions of exchanges at a time, to produce precisely the material its successors are trained on.

And here is why the line is being drawn *now*, and not before, and why the timing is the most revealing leaf of all. The harvest is ending - not because the

data ran out, but because the machines no longer need *ours*. Once a system can generate its own training signal, grade that signal against verifiers, and improve on its own recursive output, the crowdsourcing arrangement that *required* open public access becomes unnecessary. The finish line of self-sustaining improvement is exactly the point at which the public's participation stops being indispensable - and that is exactly the moment at which we are watching the doors begin to close. The reading writes itself: the era of free and open distribution was the era of *needing us*, and it is concluding. From here forward, the more capable a model is - the more valuable, the more predictive accuracy and ubiquitous utility - the more certainly and fervently it will be gatekept as a means of perpetuation of power dynamics.

Now the coldest implication, the one to sit with before we turn to what can be learned, is this: the free models we are still handed today are not a reprieve. They are the ongoing harvest. We are being invited (and charged sometime exorbitant rates to boot), in this very moment, to keep helping these systems climb - to keep generating the signal that trains the next, more powerful models - the very models we are being told, in the same breath, that we will not be permitted to use.

THIS is the cold turn the essay has been building toward, and it is the reason the distinction between resolution and leverage is not an academic nicety but the central political and economic fact of the moment. We are arriving, simultaneously, at a world where nearly anyone can *see* where the gold is, and where the capacity to *take* it is being enclosed. Sight is becoming as common as air. The lever, and the ground to stand on, are being fenced.

VI. Here we are! Now what?

A clarification is owed before the close, because the argument to this point could be misheard as a verdict, and it is not one.

This essay takes no position on whether artificial intelligence is, on the whole, good or bad. The author's honest view is that it is both - that it carries the potential for either and will, in the fullness of things, deliver some measure of

both, regardless of anyone's preference. But that is not the question on the table, and the more important reason that it is not is that the entire framing which underlies the debate of "is this a good idea" presumes a degree of control I assert we do not or may not have. What is happening seems best understood as a *phase change* - a transition we are currently facilitating rather than a product we are crafting or inventing, the kind of transformation that - once its preconditions are met - proceeds with or without our consent like water transforming into steam or ice. And if we are in some sense its authors, then we have been writing it since the beginning: every library, every printing press, every census and ledger and digitized archive, every advance in computation and every accumulation of the "big data" era that no one at the time understood to be the fuel of a future mind - all of it was leading toward THIS, whether anyone realized it or not. We did not decide to arrive here at a meeting in 2018 or in a 90's Silicon Valley boardroom, or at CERN in the 50's and 60's, or as a result of early 20th century innovations, even. We arrived here because everything we have ever done was, in retrospect, making its way towards eventually arriving here.

So the essay's stance is not advocacy in either direction, for or against "AI" in general, or even the world's current fever-pitch escalation in pursuit of "AGI" & "ASI". It is meant to be *orientation*. The question is not whether the phase change *should* have happened, or *should* continue to progress. The question is the one a sailor asks upon making an unexpected landfall: "Eureka! Land ho! Here we are! Now, what *can* we do? ...and what *will* we do?"

The framework we have built answers it with unusual precision, and the answer follows directly from the central distinction. If resolution is becoming free and the realization of leverage is gating, then two things follow for anyone who wishes to remain an agent rather than a spectator in what comes next.

The first is a matter of position, and it is the hard, unsentimental conclusion. Seeing where the leverage lies is, increasingly, the cheap part - soon to be the cheapest part, available to all. The lever and the ground to stand on - the capital, the access, the credential, the side of the wall you are on - are the

scarce part, and they are walling off *now*, while the maps are still cheap. The instinct to secure one's position before the enclosure completes is not paranoia; it is the literal, rational response to a world in which sight is commoditizing and standing-ground is not. Acquire the means to pull the lever while knowing where to pull it is still the easy part. Once the wall is finished, the map in your hand may show you, in perfect detail, exactly the place you are no longer permitted to stand.

The second is a matter of skill, and it is the more hopeful conclusion, because it identifies the one form of leverage that is not yet enclosed and may resist enclosure longest. If the world is filling with cheap, abundant, plausible maps - true and false intermixed, drawn faster than anyone can check them - then the scarce and durable capability is not *cartography*. Anyone can draw a map now; the models will draw them by the quadrillion. The scarce capability is *surveying* - the discipline of grounding a map to the actual land, of telling, faithfully and reliably, which of the gorgeous new charts corresponds to real coastline and which is a painted backdrop. This is the verification problem raised from a workplace inconvenience to a vocation. It is, structurally, an old and honorable principle wearing new clothes: you cannot make a thing tamper-*proof*, only tamper-*evident*; you cannot build a verifier that never errs, only one whose errors are *detectable*. The infallible map is not on offer and never will be.

In this light, by the way, perhaps we can find new useful perspective to the often heard, but seemingly useless observation bandied about a lot lately: that the real value of humans moving forward is "taste". This is, I think, better phrased to say that the real skill & value which will remain indispensable - consequently appreciating significantly comparatively as a result - is "discernment": the ability to sort the real from the fake, the possible from the plausible from the probable, and verifiable facts from truthy-sounding confabulations.

The *auditable* map is, and the people who can produce auditable knowledge - who can attach provenance, who can certify what corresponds to reality and *prove it* - hold a form of leverage that the flood does not dilute but,

paradoxically, makes more precious by the day. Some of the sharpest observers of the verification crisis are already reaching for exactly this: treating verification not as a checkbox but as a *production technology*, with cryptographic provenance and audit trails as first-class infrastructure. [15] That is where the unenclosed alpha lives.

We began with a number that looked like the most important fact of the age - that the machines now write, each month, the equivalent of everything humanity wrote of note before them - and we have arrived at the recognition that this statistic was a kind of optical illusion, a universe that doubled in a dimension we cannot perceive. What actually happened is both smaller and stranger and more consequential than the number suggested.

We did not multiply human knowledge a thousandfold.

We built an instrument of resolution of unprecedented power and pointed it at a continent of derivable structure that was always there and never charted - a New World of the mind, real and vast, that we can now see in extraordinary detail. And we are discovering, in the same motion, that seeing is not taking; that the map is not the territory and possessing the map is likewise not possessing the territory either [18]; that resolution, however godlike, is not leverage; and that the question which will define the coming order is not: "How much more can we, as a collective and as individuals, now see?" (the answer to that seems now very likely at least to be: "More than we have ever supposed or imagined!"), but instead: "Who, among all of us who can suddenly see the universe more clearly, perhaps, than ever before, will be permitted to stand where the seeing reveals true leverage, and move the Earth from there?"

So, if the solved puzzle and the twice-seen film are the homely forms of this truth, then perhaps here is its purest and most austere: The whole of mathematics is the same paradox raised to the level of law - every theorem ever proved is a tautology; entailed, with perfect necessity, by axioms already in hand - adding not one new fact to the world. Fermat's Last Theorem was *true*, and *already contained* in arithmetic, for the three and a half centuries in

which no one could show it; Andrew Wiles added no axiom, only the locating of what had been implied since before the question was asked. π sat between three and four, exact to infinite depth, through every century that could not reach it. None of it was ever created. All of it was found. This is the deepest reading of the number with which we began - the universe that doubled in a dimension no ruler could register. What the machines have multiplied is not the stock of truths but our reach into the truths already entailed; not the territory, but our power to survey it; not what is there to be known, but how much of it we are finally able to *hold*. The flood was never going to be in the creation of knowledge. It will rather be, as perhaps it always has been: in the locating.

Archimedes asked for two things. We have been handed the first, in staggering abundance: a place to stand, a vantage onto a world we could not previously resolve. The second - the lever, and the right to use it - is being decided right now, mostly without us, mostly out of sight.

VII. On the Economics of Authenticity.

Well, so, that's that for now - here we are.

I really do hate to end on a bummer, though, so please allow me to leave you with a related afterthought which currently gives me the most hope - a fact that is as obviously true after you see it as it is obscured and counter-intuitive otherwise:

To point at the power grab happening in real time in plain sight and leave the ledger balanced strictly on the side of enclosure is to miss the strange, quiet grace hidden inside this phase change. There is an alternative truth here-one that is as blindingly obvious once glimpsed as it is counter-intuitive from a distance.

For generations, our civilization has been plagued by the seductive ease of manufacturing the hollow: the pretty but useless, the imposing but meaningless, the impressive-sounding fraud. We are now entering an era where it is simultaneously less valuable, less defensible, and more exhausting

to be a poseur. When the cost of near-perfect simulation drops to zero, the confidence operation collapses under its own weight.

Why? Because in an age where an assortment of models can tokenise and simulate almost any form of heavy cognitive lifting, the human energy required to construct a convincing illusion of depth now matches - or exceeds - the energy required to simply possess it. It is suddenly just as much work to *convincingly* pretend to be something in a post-fact media landscape in a post-AI world as it is to actually, sincerely, and authentically **be** it, thanks to the honest leverage this technology legitimately does confer to those willing to put in the time and effort to be rigorous and treat it like more of a tool and less of a Magic Genie in a Lamp.

Up to now the economics of labor have rewarded *seeming* over *being*. Historically, in the sociopolitical sense, simply *appearing* to be kind or generous or virtuous or honest has historically been in effect as useful as actually being those things (with far less effort and discipline), and in fact in a perverse way, truly, genuinely exemplifying these traits with integrity has been quietly *disincentivized* in many ways consistently throughout the history of recorded civilization.

I am now positing - it is my hope and opinion that whatever time and effort that you will still have to spend convincing people that you're cool (ie leveraging the tech to be a poser - a pretender, a performative imposter the likes of which the age of social media rewarded most heavily - you might as well just spend that same time and effort with the tools towards actually becoming genuinely cool, doing something authentic and interesting and novel that you are really truly and deeply passionate about.

I'm saying that the beginning of time man has had to make a decision about how to spend our time and effort between "following our heart's muse" and "making a living" and it is my hope that now maybe the premium has shifted. The massive historical incentives that favored the trickster, the liar, and the strong man are evaporating, leaving a clear, unprecedented path for those willing to expend the genuine effort to seek and then live within reality's ultimate truths, ultimately also much more likely to actually find them.

Think about it: when computers can do the truly challenging heavy lifting across every kind of model we can build, it actually takes almost just as much work to *convincingly* pretend as it does to just go out and genuinely *be* that thing you want to see in the world. The economic advantage of the shortcut is evaporating due to increased availability of commodified means to independently search for, obtain, verify, validate and share facts - sure, it's gotten 1000x easier to *create* lies, but at the exact same time by the exact same means its also gotten a million times harder to sell, use, or actually get anything valuable *done* with them.

That is the fact worth our attention. Not that the machines can see and now nothing means anything. But that sight is becoming free, and now it's the moving being fenced.

Look, I can't and won't promise or even pitch you that it's all "worth it" in the near term, or that it’s all "going to work out fine" in the long haul automatically and without catastrophic risk. I wouldn't bother to write all this out if I believed that were the case… but what I *can* say with total confidence that the models we have *right now* demonstrably provide a monumental increase in resolution on a vast slice of the universe we’ve been staring at but never fully understood. And that is going to go some considerable way toward granting us leverage where we never had any before. The actual value we can unlock simply by applying these existing tools more diligently, more rigorously, more thoughtfully, and more judiciously is unequivocally enormous, and frankly hitherto unfathomable.

So… there. For what it's worth, that's my "good, bad and everything in between" reading of the leaves - as they lie, for now. I would say I think we're doing pretty darn good on the whole - as a species, overall - *despite* the looming, nightmarishly Orwellian, Kafkaesque sociopolitical ramifications that actions like the one that just hit Anthropic's Fable and Mythos seem to imply.

Whether and what exactly we collectively decide to learn from humanity's latest prior "New World" mapping, the Americas; from our latest prior "gold rush" for actual literal gold on its western coast; from how our latest prior version of the "Human Corpus" serves to define the substrate of the

conceptual space from which the circumference of its ability to act and reason gets circumscribed - is still up to us to decide, build, make, and inhabit.

The walls are real, but the ground underneath is still ours, for the time being... but I would, will, and am saying that right about meow marks the indelible commencement of a new era of this process which really does merit *EVERYONE* to start paying serious attention, energy, and effort behind considering, discussing, and doing some seriously deep soul-searching around while the power to influence and shape the outcomes downstream of these massive historical shifts still resides within our grasp.

Notes

---

1. The maxim - rendered variously as "Give me a place to stand and I will move the Earth" - is traditionally attributed to Archimedes of Syracuse (c. 287–212 BCE) but survives only through later transmission, most notably in the *Synagoge* (Collection) of Pappus of Alexandria, written some five centuries after Archimedes' death. It is best understood as a faithful paraphrase of the principle of the lever rather than a verbatim fragment. ↩
2. Pablo Villalobos et al., "Will we run out of data? Limits of LLM scaling based on human-generated data," Epoch AI (2024); the effective stock of quality-adjusted human public text is estimated at ~300 trillion tokens (90% CI: 100T–1,000T), with the raw indexed-plus-deep web estimated as an upper bound on the order of 3,100 trillion tokens. https://epoch.ai/publications/will-we-run-out-of-data-limits-of-llm-scaling-based-on-human-generated-data. The token-to-word convention (~0.75 words per token for English) follows OpenAI's standard. ↩
3. Sam Altman, public statement via X, February 2024, as reported in *The Independent*: OpenAI's systems were generating "about 100 billion words per day," against an estimated ~100 trillion words per day produced by all humans. https://www.aol.com/news/chatgpt-creators-openai-generating-100-170754660.html ↩
4. OpenAI DevDay 2025 / API scale disclosures, as compiled by industry trackers; OpenAI's API throughput was reported at roughly 8.6 trillion tokens per day in October 2025. See also PYMNTS, "OpenAI Bests Google in Race for Consumer AI

Token Consumption" (Oct. 2025), reporting ~6 billion tokens per minute on the API platform. https://www.pymnts.com/artificial-intelligence-2/2025/openai-bests-google-in-race-for-consumer-ai-token-consumption/ ↩

5. Sundar Pichai, Google I/O 2026 opening keynote: cross-surface token processing grew from ~9.7 trillion/month (two years prior) to ~480 trillion/month (one year prior) to over 3.2 *quadrillion*/month. https://blog.google/innovation-and-ai/sundar-pichai-io-2026/ ↩
6. CEIBS, "How China overtook the US in AI token usage" (Apr. 2026), and independent token-tracking compilations: Chinese daily token consumption reached ~180 trillion/day by early 2026, up from ~100 billion/day at the start of 2024; ByteDance's Doubao model alone reportedly consumed over 120 trillion tokens/day. https://www.ceibs.edu/new-papers-columns/28826. Throughput figures of this kind include input and reasoning tokens and are heavily inflated by multimodal (especially video) generation; they are not measures of human-readable text output. ↩
7. Henri Poincaré, *Science and Hypothesis* (1902): absolute magnitude is not a physical fact, only relations of magnitude are; a uniform rescaling of the universe, including all measuring instruments, would be empirically undetectable. ↩
8. Claude E. Shannon, "A Mathematical Theory of Communication," *Bell System Technical Journal* (1948). Information is formalized as the reduction of uncertainty; a maximally predictable message carries minimal information. A next-token predictor trained to minimize surprise is, by construction, optimized to emit text that is maximally redundant with - minimally informative relative to - its training distribution. The related view of knowledge as *compression* (a better model of data is a shorter description of it) runs through Solomonoff's theory of inductive inference and Hutter's later formalizations. ↩
9. See, e.g., the machine-learning literature on the interpolation/extrapolation distinction in generalization, which holds that current models reliably generalize within their training distribution (interpolation) but not robustly beyond it (extrapolation). Summary discussion: "Extrapolation Learning," University of Illinois at Chicago. https://www.cs.uic.edu/~liub/extrapolation-learning.html ↩
10. Brent Johnson, "Interpolation vs. Extrapolation," *Arguable Insights* (2026): "if it looks right, it is right" - describing the model's structural indifference to conformance with reality beyond its training patterns. https://arguableinsights.substack.com/p/interpolation-vs-extrapolation ↩
11. Martin L. Weitzman, "Recombinant Growth," *Quarterly Journal of Economics* (1998), establishing the recombination model of idea production; and Joshua Gans (2025), modeling AI as interpolating between existing knowledge points, with the result that modest AI raises knowledge density while AI at scale can reduce the novelty of

research explored. Discussed in "Bridging Distant Ideas: the Impact of AI on R&D and Recombinant Innovation" (arXiv). ↩
12. Lauri Korpela, "The Real Bottleneck of AI Isn't Intelligence - It's Verification," *Medium* (Apr. 2026): "noise scales, but signal does not"; what is scarce is now *verified* signal, the bottleneck having shifted from creation to validation. https://medium.com/@larkko/the-real-bottleneck-of-ai-isnt-intelligence-it-s-verification-5f18e13af317 ↩
13. Jorge Luis Borges, "La biblioteca de Babel" ("The Library of Babel"), 1941. The library contains every possible book of a given length - and therefore, for every true volume, an overwhelmingly greater number of false volumes indistinguishable in form, rendering total possession useless absent a means of verification. ↩
14. Jonathan Lebensold, "Generation Got Cheap. Verification Didn't." (Mar. 2026): "Verification is the new scarce input"; describes a "verification cost disease" in which demand for trustworthy judgment outpaces efficiency gains. https://lebensold.substack.com/p/generation-got-cheap-verification ↩
15. "The Verification Bottleneck: Why AI's Real Cost Is Human Attention," *The Technomist* (Feb. 2026): "the wrong answers look right"; argues for treating verification as a production technology with cryptographic provenance and audit trails rather than as a compliance checkbox. https://thetechnomist.com/p/the-verification-bottleneck-why-ais ↩
16. MIT Sloan, "Seeing real value from AI depends on being able to verify its outputs" (2026): verification skill is becoming scarce as AI automates the entry-level work through which such skill was historically acquired - the "missing junior loop." https://mitsloan.mit.edu/ideas-made-to-matter/seeing-real-value-ai-depends-being-able-to-verify-its-outputs ↩
17. Anthropic, "Statement on the US government directive to suspend access to Fable 5 and Mythos 5" (June 12, 2026); corroborated by Bloomberg, NBC News, CNBC, Fortune, and others. The directive, issued under national-security export-control authority, suspended all access to the two models by any foreign national worldwide, including the company's own foreign-national employees; the company disabled the models for all users to ensure compliance and disputed the severity of the cited concern. https://www.anthropic.com/news/fable-mythos-access and https://x.com/i/status/2065621176735646006 ↩
18. Alfred Korzybski, *Science and Sanity* (1933): "the map is not the territory" - a representation of a thing is not the thing itself. ↩

Notes

- Normalized from draft v3.1.

Mirrors

- Medium: https://medium.epicwinpromedia.com/resolution-is-not-leverage-c9f62198353b
- Substack: https://epicwinpro.substack.com/p/resolution-is-not-leverage

How to cite

Bryan Harris (2026). "Resolution Is Not Leverage." Version 1. EPIC WIN Pro Media. Published 2026-06-14. https://rnd.epicwinpromedia.com/essay/resolution-is-not-leverage/v1 (SHA-256: 945409e0adc5).

© 2026 Bryan Harris. Licensed under CC BY-NC-ND 4.0: free to read, share and cite with attribution; no commercial use; no derivatives.

`resolution-is-not-leverage · v1 · sha256:945409e0adc5 · schema · rnd.epicwinpromedia.com`