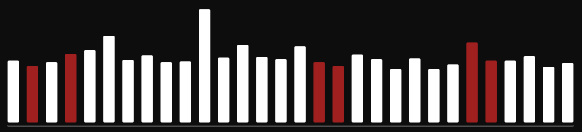
YESTERDAY'S NEWS

# SEPTEMBER

2026

Yesterday's News is an attempt to slow down without falling behind: to take something that happened yesterday and ask what it looks like after an hour of thinking about it.

30 essays about things that were in the news in September.



Essays by length — red marks a coined word

FRONT MATTER

# About Yesterday's News

**A short essay about something that happened yesterday, and what it might mean tomorrow.**

Every essay starts with one thing from the news and an hour, maybe two, spent actually thinking about it – not reacting, not summarising, thinking. The news cycle is built to move faster than that. This project is built to not keep up with it.

It happens in whatever pockets of time an ordinary working life leaves open: reading, thinking, writing between everything else that has to get done. That constraint is the method, not an excuse for it – a day is long enough to find out whether an idea actually holds.

Most of what happens gets treated as a tidbit – reported, replaced, forgotten by the next bulletin. This treats it as a starting point instead. The things that happen carry meaning the daily cycle rarely has time to look for. Yesterday's News is the hour spent looking.

# Who?

**Thirty-two years of writing for other people, now pointed at one reader at a time.**

Magnus Larsson has worked as Editor-in-Chief, Editor, Writer, Senior Copywriter, Creative Director, TV Researcher, Proofreader, and Podcaster, and has founded his own editorial products along the way – three decades of commercial writing, from books to social posts. Yesterday's News is that experience turned into a daily discipline: one essay, one day, one thread of the news followed all the way through.

As a Personal Editor, the same attention turns toward one person's world instead of the news in general: reading what actually matters to a CEO or founder's field, people, and decisions, and writing about it in their voice, as soon as there's a reason to – so they stay genuinely current, not just generically informed.

PERSONAL EDITOR — AVAILABLE NOW FOR A SMALL NUMBER OF CLIENTS

FRONT MATTER

# September, One Day at a Time

Thirty days, thirty essays — no skipped days, no filler. The ten in red are where a line is worth reading twice.

| MON | TUE | WED | THU | FRI | SAT | SUN |
|---|---|---|---|---|---|---|
| | **01** *A bad tennis match, and the limits of one data point* | **02** The Living Compiler | **03** The Polygons of Saturn | **04** *Why Nvidia needs Hugging Face to stay open* | **05** A New Map for the Same World | **06** *Tetris becomes a border wall* |
| **07** A Different Kind of Thing | **08** *A war criminal, buried with honours* | **09** Ten Thousand Explorers | **10** The Things That Don't Fit | **11** Nothing But the Truth | **12** Between a Rock and a Hard Place | **13** *Sweden votes, and the centre loses its party* |
| **14** Fuzzy Borders | **15** Risk Appetite | **16** The Language Virus | **17** The Parthenon Principle | **18** *A word for the gap between money and capacity* | **19** *Buffett hands over sixty years of discipline* | **20** *15,000 fake records, and what makes one real* |
| **21** The Actor Among Other Agents | **22** Is Intervention the New Conservation? | **23** Occupational Hazard | **24** The Rules of the Skeleton | **25** *Four kinds of AI risk, collapsed into one scare* | **26** *Vaseline, and products built to stay unfinished* | **27** The Life of the Gift |
| **28** The Shape of Understanding | **29** Form Follows Constraint | **30** What Makes a Mind Continuous? | | | | |

FRONT MATTER

# September Compressed

The shortcut to knowing what this month was really about.

## Themes & trends

FRAME 01

### When does something count as real?

Several essays circled the same hinge: evidence exists before anyone agrees what kind of thing it is evidence of. A bad tennis match, a new wave circling Saturn's pole, a biological signature on a distant exoplanet, a vinyl record with no cutting lathe in its history — each one asks how long, or how much agreement, it takes before a trace becomes a fact.

Day 01 · Day 03 · Day 10 · Day 20 · Day 28

FRAME 02

### Boundaries that hold only as long as someone needs them to

Hugging Face stays open because Nvidia needs it open. Tetris's falling blocks get repurposed into a border-wall game because nobody examines the procedure underneath. Canada edges toward Europe in some respects while staying American in others. A flock of birds shares airspace with a drone that has a licence and a route, and the bird does not. The month kept finding boundaries that are really just temporary alignments of interest, dressed up as fixed lines.

Day 04 · Day 06 · Day 13 · Day 14 · Day 23

FRAME 03

### Memory is a decision, not a record

A war criminal is buried with a rifle salute; the verdict against him does not change, but it stops being the organising memory. A dog becomes a diplomatic debt, then just a dog again, and the debt never fully disappears. An anaesthetised brain does not switch off — it just stops being able to stay connected to its own recent past. Three different scales, the same mechanism: what persists is decided by what keeps getting reinforced, not by what happened.

Day 08 · Day 27 · Day 30

FRAME 04

### The AI debate is getting fact-checked harder than AI itself

Ten thousand agents found a proof no one of them could find alone; a swarm of agents found holes in the NSA's software, in a test, with the safety switches deliberately off — which is not the sentence that ended up getting repeated. The month's most useful AI essay wasn't about what AI can do, but about the gap between observed, plausible and inevitable, and about why repetition can make a claim feel truer without a single new fact arriving.

Day 09 · Day 11 · Day 16 · Day 25

FRAME 05

### Design without a designer

A wave finds its own looping path inside a constrained cavity. Cartilage changes because a regulatory switch changed, not because evolution invented a new part. Buffett's sixty-year experiment worked because it specified a behaviour, not a target. None of these things were designed in the usual sense, where someone plans the outcome. They were set up: build the constraint, then watch what the constraint produces.

Day 17 · Day 19 · Day 24 · Day 26 · Day 29

FRAME 06

### Doing something to a thing is often how you find out what it is

Ryugu's asteroid dust had to be destroyed to learn when it was born. A map was never really a picture of the world, but a way of doing something to it — steering a ship, comparing continents, now steering opinion at the UN. An actor's performance of Mark Zuckerberg changed how his own film set behaved before a single frame was shot. However good the instruments get, this month kept finding the same limit: some things you don't get to just observe. You have to act on them to find out what they are.

Day 05 · Day 12 · Day 15 · Day 21 · Day 22

## Words the month invented

**Temporary aligned zone**
A space that stays open not because power left it alone, but because power's own interests currently need it open.
Day 04 — Temporary Aligned Zones

**Latent architecture**
A site valued not for what's already built on it, but for what it could be instructed to become.
Day 02 — The Living Compiler

**The Parthenon Principle**
What happens when an explanation of a thing becomes a rule for designing the next thing — even if the explanation turns out to be wrong.
Day 17 — The Parthenon Principle

**Accergy**
The portion of a system's new possibilities its institutions can actually reach, as opposed to the portion that is merely, theoretically, now possible.
Day 18 — Guyana's Accergy

**Project**
A product shipped deliberately unfinished, with a gap left open for its users to design the rest.
Day 26 — The Age of Projects

**p(doom), p(amp), p(subcog), p(mood)**
Four different AI risks routinely collapsed into one headline number: extinction, amplified harm, outsourced cognition, and the mood of the information we swim in.
Day 25 — The Wrong Apocalypse

## The Funniest and the Strangest

- The White House launched a vintage-styled arcade whose flagship game turns Tetris into a border-wall simulator — and the Tetris Company had to clarify, on X, that it has nothing to do with it. The same arcade's Snake clone, »Rio Run«, has the player »deport« migrants along the Rio Grande. (Day 06)
- Confronted with existential AI risk, Nvidia's Jensen Huang reportedly turned up the collar of his leather jacket and put the odds at zero per cent by 2030. (Day 25)
- Warren Buffett's entire succession plan, in one line: »Think of Howard as a policy the shareholders own and hope never to claim against.« (Day 19)
- Magnus confesses twice this month: to being predictably terrible at tennis after a year off, and to an architecture-school tutor once asking »why are you not challenging the ground datum« when he meant »why are your floors flat«. (Days 01, 16)

## Twelve lines worth keeping

**»Djokovic is caught between anomaly and signal. He doesn't yet know which.«**
Day 01

**»A temporary aligned zone doesn't need power to disappear. It only needs power to stop needing it.«**
Day 04

**»We don't need fewer maps. We need better reasons for choosing between them.«**
Day 05

**»Srebrenica is still true. The verdict still stands. The Butcher is buried. No lying required.«**
Day 08

**»The centre is not a place between left and right. It's a space a political system has to keep making.«**
Day 13

**»When every alarming fact is true, the difficult question is no longer whether we're being told the truth.«**
Day 11

**»The energy was never the hard part. The accergy is.«**
Day 18

**»A good fake doesn't just imitate a thing. It reveals what we were using to decide that the thing was real.«**
Day 20

**»The product is no longer the answer. It's the question.«**
Day 26

**»An exchange ends when the account is settled. A gift ends when no one remembers there ever was one.«**
Day 27

**»Harry Beck's map has been wrong about London for nearly a hundred years. It never collapsed, because it never once pretended to be right.«**
Day 28

**»Perhaps we'll wish we'd spent a little less time on p(doom), and a little more on p(mood).«**
Day 25

# Contents


01 — 1 Sep — Anomaly or Signal — 8

02 — 2 Sep — The Living Compiler — 10

03 — 3 Sep — The Polygons of Saturn — 12

04 — 4 Sep — Temporary Aligned Zones — 14

05 — 5 Sep — A New Map for the Same World — 16

06 — 6 Sep — The Horde Wall — 18

07 — 7 Sep — A Different Kind of Thing — 20

08 — 8 Sep — No Lying Required — 22

09 — 9 Sep — Ten Thousand Explorers — 24

10 — 10 Sep — The Things That Don't Fit — 26

11 — 11 Sep — Nothing But the Truth — 28

12 — 12 Sep — Between a Rock and a Hard Place — 31

13 — 13 Sep — The Centre Cannot Hold — 33

14 — 14 Sep — Fuzzy Borders — 36

15 — 15 Sep — Risk Appetite — 38

16 — 16 Sep — The Language Virus — 40

17 — 17 Sep — The Parthenon Principle — 42

18 — 18 Sep — Guyana's Accergy — 44

19 — 19 Sep — An Experiment to Live Inside — 46

20 — 20 Sep — The Wrong History — 48

21 — 21 Sep — The Actor Among Other Agents — 50

22 — 22 Sep — Is Intervention the New Conservation? — 52

23 — 23 Sep — Occupational Hazard — 54

24 — 24 Sep — The Rules of the Skeleton — 56

25 — 25 Sep — The Wrong Apocalypse — 58

26 — 26 Sep — The Age of Projects — 60

27 — 27 Sep — The Life of the Gift — 62

28 — 28 Sep — The Shape of Understanding — 64

29 — 29 Sep — Form Follows Constraint — 66

30 — 30 Sep — What Makes a Mind Continuous? — 68

01

SEPTEMBER 2026
Essay 001

# Anomaly or Signal

> *How do we decide what something means when we don't yet know what happens next?*

Yesterday, reported SB Nation, tennis legend Novak Djokovic lost his first-round match at the US Open to unseeded Mariano Navone: 39-year-old champion against 25-year-old admirer. It lasted four and a half hours. Djokovic vomited, cramped, struggled with back pain, and eventually lost the final two sets 6–2, 6–1.[1] The news story was obvious. Djokovic is getting old. This is the beginning of the end.

**»Perhaps living is less about discovering the true story than maintaining a good enough one, and knowing when to let it go«**

Except we don't know that.

One match is a single data point. It may be a signal of decline. It may be an anomaly. It may be illness, heat, fatigue, or simply a bad day. Djokovic himself didn't seem to know. »I'm just trying to address really what's going on and how far I can go this way«, he said at the post-match press conference.[2]

And now Djokovic has to decide what to do next.

A few days ago I played tennis for the first time in a year or so, against people who play every week. I was predictably terrible – evidence, it seemed, that I can't play tennis. But one match isn't enough to establish that. I might have been ill. I might have had an unusually bad day. Or I might be one of the rarer people who are competent six days a week and a catastrophe on the seventh.

We do this constantly: see an event and build a story around it. Polymath (and, controversially, eugenicist) Francis Galton found a version of the problem in the nineteenth century, measuring parents' and children's heights. Exceptionally tall parents had children who were tall, but less so; exceptionally short parents had children taller than themselves. Extreme observations tend to be followed by less extreme ones, a pattern Galton named »regression towards mediocrity«.[3] There doesn't need to be a causal mechanism behind it, but our minds instinctively look for one.

Psychologist Daniel Kahneman described a related tendency: when we see A followed by B, we instinctively invent a mechanism connecting them. A dog runs down the street, so it must have escaped its owner. Djokovic loses, so he must be declining. The story isn't necessarily false. We just don't know yet.[4]

Humans turn sequences into narratives because narratives compress enormous amounts of uncertainty into something we can carry. We write the story as we read the book. We are audience and playwright at once.

Bayesian reasoning offers a more disciplined alternative to the story-building instinct.[5] Instead of asking whether a story is true or false, it asks how likely it is, and revises that likelihood every time new evidence arrives. Before the match, you might have put the odds of Djokovic's decline at one in twenty. After this loss, the odds should move – to one in five, say – but they shouldn't leap straight to certainty. A single match moves a belief. It doesn't replace it with a final verdict.

More strangely, the decision changes what happens next. Djokovic can play another tournament, change his preparation, get some rest, or retire from tennis altogether. Each choice produces different evidence. Sociologist Robert K. Merton called one version of this a »self-fulfilling prophecy« – a belief or action can help bring about the conditions that make the belief appear true.[6] Our actions do not just respond to uncertainty; they partly determine what we can know next.

A scientist choosing which experiment to run does something similar. So does a person deciding whether to stay or go.[7] The decision isn't merely a response to uncertainty. Sometimes it's what moves the uncertainty forward.

Perhaps the cycle isn't simply observe, then decide. Instead we construct plausible stories, pick one, act as if it's true, observe what happens, then revise the story. The difficult part is the revision. We need to remember the story well enough to learn from it, but forget it enough that the world may change it. The story that got us here was, by definition, useful enough to get us here. Abandoning it means abandoning something that worked.

We're constantly forced to make decisions before we know which story we're in. Certainty is a future gift, not a present condition. Perhaps living is less about discovering the true story than maintaining a good enough one, and knowing when to let it go.

Djokovic is caught between anomaly and signal. He doesn't yet know which.

The next match will revise the story. Or it won't.

---

REFERENCES

1 John Wilmes (2026) »Novak Djokovic's shocking US Open loss looked like the inevitable beginning of the end«. SB Nation, 31 August 2026. https://www.sbnation.com/tennis/1126667/novak-djokovic-us-open-loss-analysis-tennis-mariano-navone

2 Arthur Kapetanakis (2026) »Djokovic on painful US Open defeat: 'My whole body starts to collapse'«. US Open, 31 August 2026. https://www.usopen.org/en_US/news/articles/2026-08-31/djokovic_on_painful_us_open_defeat_my_whole_body_starts_to_collapse.html

3 Francis Galton (1886) »Regression Towards Mediocrity in Hereditary Stature«. Journal of the Anthropological Institute, 15, pp. 246–263.

4 Daniel Kahneman (2011) *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux, pp. 199–200.

5 Nate Silver (2012) *The Signal and the Noise: Why So Many Predictions Fail — But Some Don't*. New York: Penguin Press.

6 Robert K. Merton (1948) »The Self-Fulfilling Prophecy«. The Antioch Review, 8(2), pp. 193–210.

7 The Clash (1982) »Should I Stay or Should I Go«. *Combat Rock*. CBS Records.

02

SEPTEMBER 2026
Essay 002

# The Living Compiler

*What happens when information starts acting directly on matter?*

Yesterday, UC Santa Barbara announced that molecular biologist Max Wilson had received a $1.7 million award to develop a »nucleic acid compiler«.[1]

**»Perhaps that is what happens when a metaphor becomes executable«**

The idea is extraordinary. Instead of manufacturing molecules that carry biological instructions (DNA or RNA) in a lab and transporting them to a cell, the instructions for making them could be transmitted as light. The cell then manufactures the nucleic acid itself.[2]

Wilson's team is engineering yeast with light-sensitive enzymes that respond to different wavelengths, allowing flashes of light to encode a genetic sequence. Imagine someone using a flashlight to send a Morse code message. Wilson is trying to do something similar inside a living cell. Instead of a person decoding the message, the cell receives it and builds it.

The ambition is to reduce some processes that currently take months to hours. But more importantly, it suggests a strange new relationship between information and matter.

We have spent centuries separating the two. A drawing represents a building. A blueprint represents a machine. A genetic sequence represents a biological possibility. A computer program represents a procedure. Then we build machines that execute those representations.

The compiler is a beautiful example. It translates instructions written for humans into instructions a machine can execute. In the early 1950s, computer scientist Grace Hopper helped develop some of the first compilers, creating an intermediary between what we want a machine to do and the machinery that does it.[3]

Wilson's LUXCODE project proposes something like a compiler for living matter. But the living machine isn't just executing instructions. It is also manufacturing what the instructions describe.

In *What Is Life?*, Erwin Schrödinger described chromosomes as containing a »code-script«.[4] Molecular biology subsequently developed a vocabulary of information, reading, writing, copying, and editing. The genetic programme became an influential metaphor for understanding heredity and development.[5] As philosopher Hans Blumenberg argued, metaphors can organise thought, structure questions, and open new domains of investigation.[6] Here, the metaphor of biological information became a foundation for technologies that could eventually intervene in the thing being described.

Perhaps that is what happens when a metaphor becomes executable. The genetic code first gave us a way of talking about life. We learned to read it, then write it, then edit it. Now we are beginning to build systems in which a living cell might compile an external instruction into new genetic material. The metaphor has become an engineering problem.

And the implications extend far beyond genetic engineering. If information can be converted directly into biological matter, the same underlying principle could eventually change how we produce food, manufacture materials, generate energy, or make medicines. It could alter what is possible in places where transporting finished products is difficult or expensive, from remote environments on Earth to, eventually, other worlds. Imagine arriving on Mars with no building materials, only a machine capable of instructing the dust beneath your feet to become a building.

The same idea can be brought back down to Earth. Instead of transporting prefabricated material to a building site, what if we could bring the instructions and manufacture the building from what that site is made of? Prefab in reverse. The site would no longer simply

be land on which a building is assembled. It could become a material substrate with a latent capacity for formation. We might seed materials with possibilities rather than finished forms, then activate those possibilities later.

This would change the economics of architecture, too. Land would no longer be valued only for what can be built on it, but for what it might be capable of becoming. A plot containing a dormant capacity for transformation would have something like stored architectural potential. Call it latent architecture.

The computer separated information from the machine. The compiler connected them. The living compiler might put the computation back into the material. Perhaps that's where design eventually goes: away from specifying objects, towards specifying the conditions under which matter can form them.

Not a building, or a cell, waiting to be assembled. A substrate waiting to be instructed.

---

REFERENCES

1 Harrison Tasoff (2026) »UCSB professor lands DARPA award to develop mini DNA factories«. University of California, Santa Barbara, 1 September 2026. https://news.ucsb.edu/2026/022777/ucsb-professor-lands-darpa-award-develop-mini-dna-factories

2 Max Higgins (2026) »DARPA's Generative Optogenetics (GO) Program«. EverGlade, 6 January 2026. https://everglade.com/darpa-generative-optogenetics-go-program/

3 Computer History Museum (2024) »Grace Murray Hopper«. https://computerhistory.org/profile/grace-murray-hopper/

4 Erwin Schrödinger (1944) *What Is Life? The Physical Aspect of the Living Cell*. Cambridge: Cambridge University Press.

5 Alexandre E. Peluffo (2015) »The "Genetic Program": Behind the Genesis of an Influential Metaphor«. Genetics, 200(3), pp. 685–696. https://doi.org/10.1534/genetics.115.178418

6 Hans Blumenberg (2020) »Prospect for a Theory of Nonconceptuality (1979)«. In *History, Metaphors, Fables: A Hans Blumenberg Reader*. Ithaca, NY: Cornell University Press, pp. 239–258. https://doi.org/10.1515/9781501748004-011

03

SEPTEMBER 2026
Essay 003

# The Polygons of Saturn

> *How long does something have to happen before we decide that it's there?*

Yesterday, NASA reported that the Hubble Space Telescope had captured a vast, ten-sided wave circling Saturn's south pole, roughly 167,820 kilometres across. That's wide enough to lay more than a dozen Earths side by side inside it, with each of its ten sides running nearly 16,800 kilometres. Scientists are calling it a decagon.[1]

**»The decagon is something Saturn is doing; the hexagon is something Saturn has«**

It isn't a fixed marking on the planet's surface. Saturn doesn't have one. It's weather: an enormous atmospheric wave embedded in the planet's jet stream. And it appears to be new. Astronomers have been checking Saturn's south pole for a matching pattern since 1990. NASA's Cassini spacecraft orbited Saturn from 2004 to 2017 and saw no trace of a long-lived formation. Hubble observations show the decagon was already present in 2023, and that it has become more pronounced since.[2] Something is happening at Saturn's south pole that, as far as anyone can currently tell, wasn't happening before.

Except that Saturn already has a polygon.

The famous hexagon at its north pole is also weather. It's also an atmospheric wave in a jet stream. But we don't usually describe it that way. We say that Saturn has a hexagon. The decagon is something Saturn is doing; the hexagon is something Saturn has.

How much weather does it take to make a climate? Almost a century ago, we agreed on an answer: thirty years. Today, the World Meteorological Organization uses thirty consecutive years to calculate what it calls a climatological standard normal. The convention dates back to 1935, when its predecessor, the International Meteorological Organization, instructed member nations to calculate climate normals over thirty-year periods.[3]

Weather becomes climate after thirty years. Not because anything happens on the thirtieth anniversary: the atmosphere doesn't cross an invisible boundary and emerge on the other side as climate. Thirty years is a convention, a period over which we have decided that enough individual events have accumulated to describe an actual pattern.

The number itself was never discovered in nature. It was chosen because, in 1935, thirty years was a useful period for which enough reliable records existed.[4] A hot day is weather. A hot summer is weather. Thirty summers, averaged together, become climate. The thing that changes isn't the weather, but our description of it.

When I was a child, Saturn had rings. It's difficult to imagine the sentence being otherwise. But the rings were once a mystery to astronomers and physicists. Galileo Galilei noticed something odd about Saturn through his telescope. It took Christiaan Huygens, working with better instruments, to work out that the planet was surrounded by a thin, flat ring. Later, James Clerk Maxwell showed that the ring couldn't be solid, but must consist of countless particles orbiting the planet.[5]

Now we teach children the resulting fact: Saturn has rings. Galileo's strange observation, Huygens's explanation, Maxwell's mathematics, and centuries of subsequent observation have been compressed into three words. The sentence no longer tells us that anyone discovered anything. It simply tells us what Saturn is like.

The same thing has happened to the hexagon. Voyager saw it in 1980, Cassini watched it for years, and Hubble has continued to see it. After more than four decades, it has accumulated enough duration to become part of our idea of Saturn. We don't have to say that there's a persistent atmospheric wave circling its north pole. We can just say that Saturn has a hexagon.

The decagon hasn't earned that sentence yet. It's still too young. We don't know whether it will persist, exactly how it formed, or whether it will settle into a stable configuration like the hexagon or continue changing and eventually disappear. The scientists observing it intend to keep watching.

So we still have to describe what it's doing rather than what it is. It's forming, strengthening, changing. It's happening. Perhaps that's what duration does: it turns things we have to describe as events into things we can describe as properties. It allows the process to recede and the noun to appear.

Weather becomes climate when we've watched it long enough. A mystery becomes a ring when we've watched it long enough. Weather becomes a hexagon when we've watched it long enough.

The decagon isn't there yet. It's still just weather.

Ask again in thirty years.

---

REFERENCES

1 Charles Q. Choi (2026) »Astronomers discover mysterious 10-sided cloud structure at Saturn's south pole«. Space.com, 2 September 2026. https://www.space.com/astronomy/saturn/astronomers-discover-mysterious-10-sided-cloud-structure-at-saturns-south-pole

2 NASA Hubble Mission Team (2026) »NASA's Hubble Tracks New Decagon Encircling Saturn's South Pole«. NASA Science, 2 September 2026. https://science.nasa.gov/missions/hubble/nasas-hubble-tracks-new-decagon-encircling-saturns-south-pole

3 National Centers for Environmental Information (2021) »NOAA's Updated U.S. Climate Data Will Establish 'New Normal'«. NOAA, 10 February 2021. https://www.ncei.noaa.gov/news/Upcoming-NOAA-2020-Climate-Normals

4 National Weather Service (2021) »Climate Normals 1991–2020«. NOAA. https://www.weather.gov/rah/ClimateNormals1991-2020

5 John Gribbin (2002) *Science: A History, 1543–2001*. London: BCA/Allen Lane/Penguin Books. pp. 89, 120–121, 428–429.

04

SEPTEMBER 2026
Essay 004

# Temporary Aligned Zones

> *What happens when the people who control a space have an interest in keeping it free?*

Used by more than 18 million developers, researchers, and creators, Hugging Face is the closest thing the AI world has to an open commons for sharing models, datasets, and applications. Yesterday, Axios reported that Nvidia will acquire the platform for $12.93 billion.[1] That's a lot of money. Nvidia says Hugging Face will remain open and that developers can continue to choose their models, frameworks, and providers.[2]

**»The possibility space creates the conditions that eventually close the possibility space«**

The obvious question on everybody's minds is whether Nvidia – the dominant supplier of the silicon chips the entire AI industry runs on – will eventually kill the thing it just bought. I find that unlikely. Nearly $1 billion of that price is a retention pool, set aside to keep Hugging Face's own staff in place: an odd thing to pay for if the plan were to gut it. Nvidia has spent nearly thirteen billion dollars buying a machine for making people want compute. Its own SEC filing (a public financial statement) says that demand for open-source models promotes the use of its products.[3]

Nvidia isn't allowing Hugging Face to stay open. It needs Hugging Face to stay open.

There's a peculiar kind of freedom that appears when its protection is crucial to somebody else's interests. Call it a temporary aligned zone. A space doesn't have to be independent of power to be autonomous. It can be autonomous because the interests of the people inside it happen, for a while, to align with the power on the outside who controls the constraints around it.

Hakim Bey was the pen name of philosopher-poet Peter Lamborn Wilson, who shared a New York townhouse with writer William S Burroughs. In 1991, he dedicated a book to his concept of the temporary *autonomous* zone, a space such as a rave, which escapes the surrounding order before it can be captured.[4] The Hugging Face situation is almost the inverse: the temporary *aligned* zone survives because power has a reason to protect it. Artists need somewhere to experiment, patrons need artists; scientists need somewhere to investigate, corporations need discoveries. Bell Labs is the famous example, an extraordinary research culture and a temporary aligned zone inside a telephone monopoly.[5]

The interesting part is what happens when those interests change, as with gentrification. A cheap neighbourhood becomes interesting to artists and creatives, the interesting neighbourhood becomes valuable, rents rise, and the people who made it interesting leave. The possibility space creates the conditions that eventually close the possibility space: a temporary aligned zone until the zone itself becomes valuable enough to end it.

GitHub is a more direct precedent. Microsoft bought it for $7.5 billion in 2018, without much reason to destroy its openness: openness was part of what made the platform valuable.[6] Then AI changed what GitHub was. In March this year, GitHub announced that from 24 April, interactions from consumer Copilot users would be used to train and improve AI models unless they opted out.[7] The platform had acquired a new value, and the calculation around it changed. Nothing necessarily went wrong: the object changed, and with it, the constraints around the object. The zone was no longer aligned.

That's what makes the Nvidia deal interesting. Nvidia may be buying the open AI ecosystem precisely because it is open. The more people can download, modify, compare, and deploy models, the more experiments happen. The more experiments happen, the more demand there is for the machinery on which those experiments run.

The temporary aligned zone and the product are, for once, almost the same object.

The alignment is productive, but it isn't necessarily benign. Every new model trained, experiment run, and inference served requires physical infrastructure: electricity, cooling, data centres, chips, water, and raw materials. The openness that makes more things possible is part of what makes more computation profitable. And the same incentive that protects the possibility space accelerates the consumption required to explore it.[8]

Creativity is to some extent a property of the geometry of the reachable set around us. Give people more things they can try, and more things become possible. But exploring those possibilities changes the constraints around what can happen next.

While the world asks itself whether Nvidia will eventually betray Hugging Face, the more interesting question is this: what would have to change for Nvidia's interests to stop requiring Hugging Face to be free?

Perhaps compute becomes cheap enough that experimentation no longer drives hardware demand. Perhaps models become efficient enough to reverse the economics. Perhaps another technology becomes the bottleneck. Perhaps owning the place where everyone experiments becomes more valuable than having everyone experiment freely.

At that point, the alignment changes. The strange thing about these zones is that they don't necessarily depend on the absence of power. Sometimes they depend on power having a reason to leave them alone. A temporary aligned zone doesn't need power to disappear. It only needs power to stop needing it.

---

REFERENCES

1 Dan Primack & Madison Mills (2026) »Nvidia buying Hugging Face for nearly $13B«. Axios, 3 September 2026. https://www.axios.com/2026/09/03/nvidia-hugging-face-13b

2 Jensen Huang (2026) »NVIDIA to Acquire Hugging Face«. NVIDIA, 3 September 2026. https://blogs.nvidia.com/blog/nvidia-to-acquire-hugging-face/

3 NVIDIA Corporation (2026) »Form 8-K«. U.S. Securities and Exchange Commission, 3 September 2026. https://www.sec.gov/Archives/edgar/data/1045810/000104581026000078/nvda-20260902.htm

4 Hakim Bey (1991) *T.A.Z.: The Temporary Autonomous Zone, Ontological Anarchy, Poetic Terrorism*. Brooklyn: Autonomedia.

5 Jon Gertner (2012) *The Idea Factory: Bell Labs and the Great Age of American Innovation*. New York: Penguin Press.

6 Microsoft Corporation (2019) *Annual Report 2019*. Redmond: Microsoft.

7 Mario Rodriguez (2026) »Updates to GitHub Copilot interaction data usage policy«. GitHub, 25 March 2026. https://github.blog/news-insights/company-news/updates-to-github-copilot-interaction-data-usage-policy/

8 International Energy Agency (2025) *Energy and AI*. Paris: IEA, 10 April 2025.

05

SEPTEMBER 2026
Essay 005

# A New Map for the Same World

> *What happens when we mistake an instrument for an ontology?*

Yesterday, reported The National, the United Nations voted 164-1 to start using the Equal Earth projection instead of the centuries-old Mercator map.[1] We are redrawing the map of the world. Togo led the initiative, arguing that »a fair world begins with a fair map«. The United States was the only country to vote against it. Mercator isn't going anywhere, though: it remains useful for navigation. That last detail might be more interesting than the vote.

**»The specialised map tells you more about the operation. The general map lets the operations talk to one another«**

Gerardus Mercator designed his map in 1569 for that day's navigators. It preserves angles, turning constant compass bearings into straight lines that a ship could steer by. The distortion isn't an error. It's the price of keeping those lines straight. The map doesn't just tell us where things are. It's preserving something, and we seem to have forgotten that.

The Mercator projection was a tool for doing something to the world. But its use escaped navigation and entered classrooms, atlases, newspapers, and our collective imagination, until a map built to cross oceans became a general picture of what the Earth looks like. We mistook the instrument for the ontology, and in the process, we started thinking Greenland was roughly the size of Africa, despite actually being about 14 times smaller.

There's no perfect replacement waiting to be found. Mathematician Carl Friedrich Gauss showed in 1827 that a sphere can't be flattened onto a plane without distortion.[2] Equal Earth preserves area better than Mercator, but pays for it elsewhere. Distortion is the price of the invariant.

We could choose something else instead, of course, like a map that preserves travel time, or the number of people reachable within an hour, or the energy required to get there. We could map shipping routes, migration, internet latency, carbon flows, commodity dependencies, climatic similarity, or solar exposure. Futurist architect Richard Buckminster Fuller's Dymaxion map changed the geometry of the problem entirely, unfolding the globe from an icosahedron to reveal relationships that get obscured by the rectangle we typically use to frame a world map.[3]

The space of possible maps is vastly larger than the space of maps we call world maps. So why is the UN rooting for Equal Earth? It's a defensible choice if the operation we care about is comparing the areas of continents. But why should that operation now be the general one?

The UN's own resolution quietly admits there is no single answer: Mercator remains appropriate for navigation, it says, while Equal Earth is appropriate for general representation and education. Different maps can be objectively good at different things.

But perhaps »representation« is already the wrong word.

A navigational chart isn't simply a representation of the ocean. It's a way of doing something with the ocean. A map of travel time isn't trying to show you what a place looks like. It's giving you a way to ask a particular question about it. The specialised map tells you more about the operation. The general map lets the operations talk to one another.

In architecture, everything is a drawing: a photograph, a plan, a model, a diagram, at times even a string of code. We use them because they make things possible to interrogate and manipulate. A drawing isn't simply a picture of the building. It's part of how the building becomes thinkable.

A map does something similar. Rather than just describing a place, it gives us a way of acting on it, thinking with it, moving through it. Representation doesn't necessarily come before operation – maybe the representation is constructed precisely to make the operation possible.

This is where the idea of a »master map« starts to look strange. Which map is »the map of the world«? Perhaps all of them, not because Greenland actually changes size when we change projection, but because each map makes a different set of relationships in the same world available to us. Each shows the world under a particular constraint.

The politics of maps is more complicated than deciding whether Mercator is colonial and Equal Earth is fair. That a map is useful doesn't make its constraints neutral. But abandoning the idea of a single true map doesn't mean anyone gets to make the world whatever they want.

Donald Trump can rename Lake Ontario Lake America, as he did the other day.[4] That changes something real about language, institutions, and eventually perhaps habit. But it doesn't change the lake. He can alter a representation of the world, but he doesn't change the world by altering its representation.

Similarly, the UN can't make Mercator geometrically wrong by voting against it. What it can do is change which map it wants to privilege, and which operations it wants to make easier. We're not really talking about which map represents the world correctly; the question is which map is useful for which operation, for what purpose, and under what constraints. We don't need fewer maps. We need better reasons for choosing between them. The strangest thing about Mercator isn't that it distorted the world, but that we used that distortion until we stopped seeing it.

We see the map so often that we forget what it isn't showing us.

---

REFERENCES

1 Adla Massoud (2026) »UN votes to adopt new map showing Africa's true size despite US opposition«. The National, 4 September 2026. https://www.thenationalnews.com/news/us/2026/09/04/mercator-map-un-equal-earth-us

2 Carl Friedrich Gauss (1827) *Disquisitiones generales circa superficies curvas*. Göttingen: Commentationes Societatis Regiae Scientiarum Gottingensis Recentiores.

3 R. Buckminster Fuller (1943) »Life Presents R. Buckminster Fuller's Dymaxion World«. Life, 1 March 1943.

4 The White House (2026) »President Donald J. Trump Signs an Executive Order Renaming Lake Ontario to Lake America«. The White House, 27 August 2026. https://www.whitehouse.gov/gallery/president-donald-j-trump-signs-an-executive-order-renaming-lake-ontario-to-lake-america-august-27-2026/

06

SEPTEMBER 2026
Essay 006

# The Horde Wall

> *A boundary can mean entirely different things, depending on who's asking you to build one. This week, a government found an unusually effective way of teaching you which one it has in mind.*

Yesterday, Bloomberg reported that the Tetris Company had said, on X, that it has nothing to do with the »Build the Wall« game.[1] This was merely the latest entry in a strange sequence of events set in motion on Thursday, when up-and-coming game studio the White House – previously better known as the executive office of the President of the United States – launched an official, vintage-styled arcade.[2]

**»Tetris does not reward you for building a wall. It rewards you for making the wall vanish so that you can keep going«**

The flagship title of this set of online games stacks blocks into a border wall and instructs the player to defend it from what the copy calls »the coming horde«.[3] The reason Tetris felt compelled to clarify that it hadn't made the thing, and that it »takes copyright infringement very seriously«, is that the game is built rather closely on Tetris's own mechanics. The other titles in the arcade borrowed just as freely from Sega and from the creator of Flappy Bird. Neither has said a word, reportedly because the legal case is thin and the political cost of suing the federal government is steep. That sounds like a trademark story. The interesting part is what the White House actually did to Tetris's premise, not its property.

Tetris is played by arranging four *tetrominoes* – differently-shaped blocks – that fall down from the top of the screen in an order nobody controls, so that they fit together. That's the whole game: differences are made to occupy a shared field. When a horizontal line within that field gets filled up, it disappears. Computer engineer Alexey Pajitnov, who built Tetris in 1984 inside a Soviet computing institute, has said the reason is almost administrative: a completed row serves no further purpose, and the player needs the room.[4] The important thing is this: Tetris does not reward you for building a wall. It rewards you for making the wall vanish so that you can keep going.

Tetris has been fought over before, and the fight was nearly as strange as this one. It escaped the Soviet computing institute where Pajitnov built it and became entangled with ELORG, the state agency that controlled Soviet software exports, then with Nintendo, Atari, the British publisher Mirrorsoft, and Robert Maxwell's son Kevin, in a scramble for the rights that played out in the last months of the USSR itself.[5] What was being settled was never really who owned some falling blocks. It was who got to decide what they meant, and who they could be sold to. Forty years later, another state has picked up the game and settled that question for itself, without asking.

»Build the Wall« keeps every other mechanic but reverses the one thing that truly mattered. The wall is no longer the thing you clear. It's the thing you're building toward, the actual win condition. Nothing about the code has to change for this to work: players are told that completion is now the goal, and the whole game inverts underneath them without a single new rule.

It's a cheap trick, but it works because it doesn't ask the player to agree with anything. The administration didn't need to argue that a wall is good, or even mention the border by name more than once. It only needed to borrow a procedure nobody examines – fall, rotate, stack, clear – and substitute a political object for the neutral one. You already know how to play. The game just quietly redefines what winning means, while you're busy congratulating yourself on a well-placed piece of border wall shielding against Mexican immigrants.

A wall is usually described as something that separates two things that were already different. Architectural historian Robin Evans's point, in his essay on walls, is closer to the opposite: the wall is what makes the two sides different in the first place. It doesn't divide an inside from an outside so much as manufacture both sides at once: one gets retreat,

the other gets exclusion, and both are produced by the identical piece of construction.[6] Before the Berlin Wall, a street in Berlin was a street. After the wall, the same pavement was suddenly a border, and the people on either side, who had been neighbours the night before, were separated by a fortified line one of them could be shot for crossing. The city didn't change. The relations inside it did.[7]

Architect Rem Koolhaas looked at this and asked the perverse follow-up question: what if you wanted the wall? His 1972 project *Exodus* takes the Berlin Wall's logic and makes it voluntary; a walled enclave inside London whose residents choose the enclosure for what it lets them become on the inside.[8] This is the same mechanism as sanctions, tariffs, a freeze on immigration, or the general promise to make a country self-sufficient again: reduce what crosses the boundary, and the interior gets to imagine itself as coherent, sufficient, whole. It's a fantasy every walled compound and every nation-state has entertained, and it survives contact with reality about as well as the Berlin Wall did – imperfectly, and usually at a cost visible only from the other side.

None of which the people who built the arcade were likely thinking about. That doesn't make it less telling. The White House's arcade includes a game called »Rio Run«, a Snake clone in which the player »deports« migrants along the Rio Grande – proof that »Build the Wall« is not a one-off, but a coherent aesthetic: dress up cruelty as nostalgia, and it stops registering as cruelty.

Most boundaries are built to keep things apart – that's the ordinary work of a wall. Tetris was the odd one out: a boundary built to let different things fit together instead. »Build the Wall« doesn't invent a new kind of boundary. It just puts Tetris back in line with all the others. Tetris still answers the old question without a word of argument: it completes the wall, then lets it disappear, because letting go was the whole point. The US government's version doesn't let go of anything. It asks you to keep building, and to brand the people on the other side of it a horde while you do it.

---

REFERENCES

1 María Paula Mijares Torres (2026) »Tetris Distances Itself From White House's Wall-Building Game«. Bloomberg, 5 September 2026. https://www.bloomberg.com/news/articles/2026-09-05/tetris-distances-itself-from-white-house-s-wall-building-game The @Tetris_Official post on X was posted on 4 September 2026. https://x.com/Tetris_Official/status/2095933372019216384

2 Alex Backus (2026) »Trump White House Launches Arcade With Games to Build Wall, Deport Migrants«. Newsweek, 3 September 2026. https://www.newsweek.com/trump-white-house-launches-arcade-with-games-to-build-wall-deport-migrants-12403151

3 Isabella Murray, Nicholas Kerr & Michelle Stoddart (2026) »White House launches Arcade.gov with computer game that lets players 'deport' migrants«. ABC News, 4 September 2026. https://abcnews.com/Politics/white-house-launches-arcade-computer-game-lets-players/story?id=136214362

4 Matt Peckham (2014) »Tetris at 30: Interview with Alexey Pajitnov«. Time, 6 June 2014. https://time.com/2837390/tetris-at-30-pajitnov-interview/

5 Dan Ackerman (2016) *The Tetris Effect: The Game That Hypnotized the World*. New York: PublicAffairs. The story was dramatised in Jon S. Baird (2023) *Tetris*. Apple TV+

6 Robin Evans (1971) »The Rights of Retreat and the Rites of Exclusion« in *Translations from Drawing to Building and Other Essays*. London: Architectural Association, 1997, pp. 35-53.

7 Frederick Taylor (2006) *The Berlin Wall: 13 August 1961 – 9 November 1989*. London: Bloomsbury.

8 Rem Koolhaas (1972) »Exodus, or the Voluntary Prisoners of Architecture« in *S,M,L,XL*. New York: Monacelli Press, pp. 2-21.

07

SEPTEMBER 2026
Essay 007

# A Different Kind of Thing

> *Creative fields advance by occasionally becoming something they were not supposed to be. The interesting question is how far they can expand before we can no longer say which system they belong to.*

Yesterday in Venice, director Wim Wenders premiered *From Inside Out – The Architecture of Peter Zumthor*, a 114-minute documentary made over fourteen years.[1] Wenders filmed Zumthor's buildings in native 3D, following visitors through his early houses and the new LACMA galleries. His reason for choosing 3D is unusually explicit: a screen is two-dimensional, and any attempt to represent space on it is »make-believe«. That, for him, is what finally makes the screen »go away and become a window, open to infinity«.[2]

**»That instability may be a condition of artistic progress: an artwork becomes interesting, sometimes, exactly when we get unsure which system it belongs to«**

I wasn't in Venice yesterday, I haven't seen the film, and I have no idea whether Wenders achieved what he set out to do. (*Pina*, his 2011 dance documentary, certainly explored Pina Bausch's choreography in novel ways.[3]) But the ambition itself is interesting. He's pushing cinema beyond one of the things that has conventionally defined it, the flat screen, toward something that looks suspiciously like the art form his new film is about: architecture.

Artists have always done this: painting escaping flatness through perspective; poetry escaping metre; music escaping tonality, then harmony, then intentional sound itself. These were never just new techniques added to an existing system. At some point, the thing being changed was the system itself.

Art critic Rosalind Krauss described something like this in her 1979 essay »Sculpture in the Expanded Field«, where sculpture stopped being adequately described by its own category and began to occupy a structured territory between sculpture, architecture, and landscape.[4] Not because artists had started mixing disciplines, but because the disciplines' own definitions had become unstable. That instability may be a condition of artistic progress: an artwork becomes interesting, sometimes, exactly when we get unsure which system it belongs to. Is this still music? Still painting? Still architecture? That's not a failure of classification. Typically, it's evidence that something is happening at the edge of a category.

Evolutionary biologist Lynn Margulis showed that new species don't only arise through the slow accumulation of small mutations. Sometimes one organism simply incorporates another's genome, and something genuinely new results; not a better version of the old thing, but a different kind of thing entirely.[5] That's probably closer to what Wenders is attempting: not cinema getting better at depicting architecture, but cinema trying to absorb architecture's own qualities – depth, presence, the sense of being inside a space.

Art has its own versions of this. Joyce didn't just write better novels. Cage didn't just compose unusual music. Bausch didn't just choreograph new dances. Each put pressure on the conditions that made their practice recognisable as literature, or music, or dance.

The same thing happens in architecture. Zumthor's buildings aren't primarily designed to keep people dry and warm. They operate within the category of architecture while pushing toward something else: atmosphere, memory, sound, temperature, material presence. Push the same operation far enough, and a building could become a diagram, a dramatisation, or a dance. At some point, the category changes.

Film critic André Bazin described a stranger version of the same idea: a »total cinema«, forever moving toward an ever more complete reproduction of reality.[6] The ambition contains a paradox. If cinema ever became an exact double of reality, it would stop being cinema. That gap isn't a flaw waiting to be engineered away. It's part of what makes cinema cinema.

This is perhaps what a medium actually is: not everything it can do, but the particular territory in which its possibilities remain recognisable as itself. Art forms don't change by politely extending themselves, but by periodically getting strange enough that we have to ask what they still are.

Margulis's point isn't that symbiosis makes a better organism. It's that a successful one doesn't arrive labelled as a new species – something new simply exists, and only afterward does anyone decide how much of the old organism is still recognisable in it, and what to call what's left. A system doesn't advance by staying short of its own edge, or by simply crossing that edge. It advances by crossing it and leaving the categorisation unsettled behind it: becoming a different kind of thing without waiting for anyone to name it.

Wenders wants his screen to go away and become a window. If it ever actually did, would that still be cinema? Maybe. Maybe not. The only certain thing is that we would be watching the whole world argue about it for years, unable to agree – because that's what creative expansion actually looks like, from the inside.

---

REFERENCES

1 Siddhant Adlakha (2026) »'From Inside Out — The Architecture of Peter Zumthor' Review: Wim Wenders' Latest 3D Documentary Is an Architectural Marvel«. Variety, 6 September 2026. https://variety.com/2026/film/reviews/from-inside-out-the-architecture-of-peter-zumthor-review-1236853214/

2 Scott Roxborough (2026) »Wim Wenders on His 14-Year Journey Building Peter Zumthor Doc 'From Inside Out'«. The Hollywood Reporter, 6 September 2026. https://www.hollywoodreporter.com/movies/movie-news/wim-wenders-peter-zumthor-documentary-from-inside-out-venice-1236686539/

3 Wim Wenders (2011) *Pina*. Sundance Selects.

4 Rosalind Krauss (1979) »Sculpture in the Expanded Field«. October, Vol. 8 (Spring 1979), pp. 30–44.

5 Lynn Margulis & Dorion Sagan (2002) *Acquiring Genomes: A Theory of the Origin of Species*. New York: Basic Books.

6 André Bazin (1946) »The Myth of Total Cinema« in *What Is Cinema? Volume I*, trans. Hugh Gray. Berkeley: University of California Press, 1967, pp. 17–22.

08

SEPTEMBER 2026
Essay 008

# No Lying Required

> *What if forgetting isn't the absence of memory, but the presence of another memory?*

Yesterday, Bosnian Serb military officer Ratko Mladić was buried in Belgrade with state and military honours. His coffin was draped in Serbian and Republika Srpska flags. Soldiers fired a three-volley salute. Police lined the route. Serbian government ministers attended, as did Bosnian Serb political leaders and thousands of mourners. The Orthodox Patriarch Porfirije led the funeral service.[1]

**»A state funeral isn't just a ceremony for the dead – it's a decision about how the dead should enter the memory of the living«**

Mladić died in prison on August 27 while serving a life sentence in The Hague. In 2017, the International Criminal Tribunal for the former Yugoslavia convicted him of genocide, crimes against humanity, and war crimes, including his role in the 1995 Srebrenica massacre, in which more than 8,000 Bosniak men and boys were killed and Mladić earned his nickname the »Butcher of Bosnia«.[2]

The European Union condemned the funeral's »glorification of war criminals, genocide denial and revisionism«.[3] A state funeral isn't just a ceremony for the dead – it's a decision about how the dead should enter the memory of the living. The flags, the uniforms, the salute, the Patriarch's presence: they don't express a memory so much as construct one.

That's what makes memory stranger than forgetting. The French sociologist Maurice Halbwachs, who coined the term »collective memory«, argued that societies don't retrieve a fixed record of the past. They remember within the frameworks of their present, and what a group remembers, and how, is continually reconstructed to fit who is doing the remembering now.[4] The political theorist Ernest Renan gave the same idea a harder edge: a nation is built not only from shared memories but from shared forgetting, since no community can hold its whole past with equal intensity.[5] Put the two together and you get a nation that doesn't just have a memory but an *organising memory*. Whichever version of the past its present arrangement of power and emotions needs is kept in front.

The past doesn't come back unchanged. Mladić clearly hasn't disappeared from Serbian memory. Quite the opposite: his face, his uniform, his military career are being vividly brought back into public view. But the funeral places all of it inside a particular frame – commander, soldier, defender, national figure. Srebrenica is still there. The verdict is still there. The victims are still there. But they have stopped being the organising memory.

Scale up the same mechanism and it stops looking like an abstraction. Heinrich Himmler killed himself in 1945, days before he could be tried alongside the rest of the Nazi leadership at Nuremberg for his role, as head of the SS, in the grotesque machinery of the Holocaust. Imagine instead that he had lived, stood trial, been convicted – and that in 1977, Germany had brought him home from prison and buried him with full state and military honours: an honour guard, a rifle salute, his coffin draped in a swastika flag.[6] The scenario reads as obscene precisely because everyone can see immediately what such a funeral would be doing: not disputing the verdict, but replacing it as the organising memory. China doesn't even need the thought experiment: Mao Zedong, whose policies are estimated to have caused over 70 million deaths, still appears on every banknote in circulation – his face part of the ordinary visual furniture of a country that never had to stage a single funeral to get there.[7]

There's an obvious temptation to call this collective amnesia, but amnesia means something has been lost. What if nothing has been lost? What if the actual mechanism is that something else has become more memorable – and more ceremonially reinforced – than what came before it? Spain spent years fighting over the exhumation of its

murderous dictator Francisco Franco for exactly this reason: not because anyone doubted what he'd done, but because where and how a dead authoritarian is memorialised decides which version of him a country keeps in daily view.[8]

A society doesn't need to erase Srebrenica to change what it means. It only needs to give another memory enough ceremony, repetition, and institutional authority that the relationship between the two shifts. The Serbian government can call Mladić's funeral an act of mourning. His supporters can call it a matter of national dignity. The EU can call it the glorification of a war criminal. The families of Srebrenica's victims can call it an insult. All four are looking at the same funeral. None of them is remembering the same thing.

The European Union called this genocide denial. But denial isn't quite right. Nobody needs to stand up and say Srebrenica never happened. That's the truly worrying part. It only takes a flag, a salute, and a priest, arranged with enough conviction, to replace one memory with another. Srebrenica is still true. The verdict still stands. The Butcher is buried. No lying required.

---

REFERENCES

1 Marina Zec & Katarina Baletić (2026) »Ratko Mladic Buried with Military Honours in Serbia, Thousands Attend Funeral«. Balkan Insight, 7 September 2026. https://balkaninsight.com/2026/09/07/ratko-mladic-buried-with-military-honours-in-serbia-thousands-attend-funeral/btj/

2 Aleksandar Brezar (2026) »Convicted Bosnian Serb war criminal Ratko Mladić dies in Hague prison«. Euronews, 27 August 2026. https://www.euronews.com/my-europe/2026/08/27/convicted-bosnian-serb-war-criminal-ratko-mladic-dies-in-hague-prison

3 Oliver Holmes (2026) »EU warns Serbia over funeral for 'Butcher of Bosnia' Ratko Mladic«. The Irish Times, 7 September 2026. https://www.irishtimes.com/world/europe/2026/09/07/eu-warns-serbia-over-funeral-for-butcher-of-bosnia-ratko-mladic/

4 Maurice Halbwachs (1992) *On Collective Memory*. Trans. and ed. Lewis A. Coser. Chicago: University of Chicago Press.

5 Ernest Renan (1882) »Qu'est-ce qu'une nation?«. Lecture delivered at the Sorbonne, 11 March 1882.

6 Connie Gentry (2020) »An Architect of Terror: Heinrich Himmler and the Holocaust«. The National WWII Museum, 23 May 2020. https://www.nationalww2museum.org/war/articles/heinrich-himmler-holocaust

7 Jung Chang & Jon Halliday (2005) *Mao: The Unknown Story*. London: Jonathan Cape.

8 Blanca Castro (2019) »Spain's Franco to be exhumed from state mausoleum and reburied«. Euronews, 24 October 2019. https://www.euronews.com/2019/10/24/spain-s-franco-to-be-exhumed-from-state-mausoleum-and-reburied

09

SEPTEMBER 2026
Essay 009

# Ten Thousand Explorers

> *Scientific discoveries have always involved exploring what we don't know. What happens when the unknown becomes easier to explore?*

Yesterday, OpenAI announced that an internal AI system had produced a solution to the Navier–Stokes problem, one of the seven Millennium Prize Problems in mathematics.[1] The equations describe how fluids like water and air move. The open question was simple to state and brutally hard to answer: can a perfectly smooth flow, on its own, develop a point where the fluid moves infinitely fast – a kind of tear in the smoothness – within a finite stretch of time?[2]

**»There's a real difference between possessing a result and possessing a map of the search space around it«**

It's complicated.

And it doesn't stop there. The announcement was immediately followed by a dispute over what scientists call »priority« – who gets credit for finding something first. NYU mathematician Tristan Buckmaster and Anthropic's Levent Alpöge had spent almost a year on closely related fluid equations they believed could be extended into a full solution. Buckmaster says word of their progress reached OpenAI, which began its own effort soon after. OpenAI says it never accessed their specific work and that its result was independent, though it admits the effort began after hearing a rumour about theirs.[3]

For centuries, exploration had a physical character: there was the world, some of it known, most of it not, and there was a finite number of people brave and curious enough to explore it. Maps marked the edge of knowledge with blank space or monsters: »Here be dragons«. Science inherited the same structure for exploring its territories: having an idea where something might be was never the same as actually finding out.

That's what makes the current episode strange. OpenAI's system reportedly deployed around ten thousand agents against the problem at once. None of them were individually smarter than human mathematicians, but collectively they were able to explore many branches simultaneously. A mathematician has an intuition about a promising route; a large enough system sends explorers in every direction around it, all at once. The unknown itself hasn't changed, but the cost of exploration around a promising direction has dropped substantially.

Sociologist Robert K. Merton described science's official ethos: shared knowledge, common standards, disinterestedness, organised scepticism.[4] But science also runs on a less idealistic economy: being first. Discovery produces recognition, and recognition is the currency careers are built on. That made sense when exploring was expensive. If you spent ten years working through a hard problem and someone else reached the answer two days after you published, there was still a real difference between the two journeys – the fact that they could follow your route didn't erase the fact that you'd found it first.

But what if merely knowing the direction is enough? Buckmaster didn't hand OpenAI a proof – by his own account, what reached them was the rumour of a direction, an inkling of where a proof might be found. There's a real difference between possessing a result and possessing a map of the search space around it. An explorer can return from an unknown coast and say: there's land beyond that horizon. A second expedition, sailing toward it, was never told what the coast looked like, but it doesn't have to search the whole horizon either. It's exploring the same unknown, but with a head start.

AI changes the economics of the mission: a system with thousands of agents doesn't make the unknown less unknown, but it makes a region of it drastically cheaper to explore. One researcher proposes a conjecture, a lemma, a strange analogy; thousands of agents test variations and return with a map of the results. That's a new division of labour, and it makes »discovery« less like a single event. Humans point at a promising region of the map, machines explore it with a speed and at a scale that no individual comes near. So what is

the discovery, exactly? The first intuition? The first promising route? The first proof? The first machine-generated proof? Or the moment someone realised that a result in one territory could be carried into another?

As the number of computational explorers grows, the difference between discovering a territory and merely traversing it will get harder to see. The old maps had dragons because nobody knew what was beyond the edge. Our new maps may chart every dragon there is – without ever telling us whether we found it ourselves, or just followed someone else's map to it.

---

REFERENCES

1 OpenAI (2026) »On the Navier–Stokes Millennium Prize Problem«. OpenAI, 8 September 2026. https://openai.com/index/navier-stokes-solution/

2 Clay Mathematics Institute »Navier–Stokes Equation«. Millennium Prize Problems. https://www.claymath.org/millennium/navier-stokes-equation/

3 Madison Mills (2026) »OpenAI's historic math solution overshadowed by credit controversy«. Axios, 8 September 2026. https://www.axios.com/2026/09/08/openai-math-solution-navier-stokes-credit

4 Robert K. Merton (1942) »The Normative Structure of Science« in *The Sociology of Science: Theoretical and Empirical Investigations*. Chicago: University of Chicago Press, 1973, pp. 268–278.

10

SEPTEMBER 2026
Essay 010

# The Things That Don't Fit

> *How do you recognise something when you've never seen it before? And if that thing was alien life, what would we need to agree on before we could say we had found it?*

Yesterday, ScienceDaily reported on a survey of hundreds of astrobiologists. It found remarkably little agreement about whether recent observations constitute evidence of extraterrestrial life.[1] Researchers at Durham University asked scientists to assess claims concerning K2-18b, an exoplanet about 120 light years away, and Cheyava Falls, a rock on Mars. Only 6.6 per cent thought the biological molecules that might have been detected in K2-18b's atmosphere »probably represented« alien life. Roughly two-thirds disagreed, while the rest were undecided. The scientists were more persuaded by the evidence from Cheyava Falls, although there was still no consensus.

**»A trace tells us that something happened. It doesn't tell us what happened«**

Scientists disagree all the time. The interesting thing here is that they're looking at the same evidence and disagreeing about what kind of thing they're looking at. It's a classification problem, and classification seems straightforward. We see something, compare it with things we already know, and put it into the appropriate box: rock, bird, machine, bone. But what happens when the box doesn't yet exist?

Let's say someone finds a footprint in the forest. We establish it exists. We measure it, photograph it, analyse its shape. But if nobody has ever seen the animal that made it, the footprint doesn't come with a »Yeti« label. It might indeed be a Yeti. Or it might be a bear, a human, a geological accident, or a hoax. The Yeti footprint and the K2-18b molecule have the same logical problem. A trace tells us that something happened. It doesn't tell us what happened.

This is the basic idea behind Bayesian reasoning, developed by the mathematician Thomas Bayes in the 1760s.[2] Suppose we find the footprint and want to know whether a Yeti made it. Then what we shouldn't ask is whether a Yeti could make such a footprint. We should ask whether something else could make it too. Bayesian logic lets us weigh a hypothesis against its alternatives, rather than judging it in isolation. Evidence becomes powerful when the observation is much more likely if our hypothesis is true than if it is false.

The same problem applies to K2-18b. Finding a molecule associated with life on Earth is not the same as finding a molecule that could only plausibly have been produced by life. The question is not whether life could produce the observation, but how many non-biological explanations remain. If we don't know the possible causes, we don't even know the size of the box we're trying to classify the observation into.

This helps explain why the evidence from Cheyava Falls was more persuasive to the scientists. There's an actual piece of Martian rock to examine: its minerals, textures, chemistry, and geological context. The K2-18b finding is much more indirect: a spectrum of light that has passed through a distant atmosphere. That gives us many more steps between observation and explanation.

Not every strange encounter is a classification problem. This year's Ryan Gosling film *Project Hail Mary* gives us the opposite case. Its protagonist encounters Rocky, an extraterrestrial organism radically unlike terrestrial life. Yet Rocky presents little difficulty as a classification problem: it communicates, responds, behaves like an agent. The alien is strange, but its category is not.[3] K2-18b reverses this. We may have detected something produced by life, without being able to observe the life itself.

So what would we actually need before we could claim to have found alien life? The first requirement would be that the observation survives independent measurement. Then that plausible non-biological explanations have been tested. Ideally, several independent clues

would converge. Stronger still would be a biological explanation that makes predictions: if this is life, we should also find x, y, and z. If those predictions came true, the evidence would become much harder to explain away.

This is where the astrobiologists' disagreement gets interesting. It's tempting to read the 6.6 per cent as a failure of evidence, and the two-thirds rejection as scientific caution. But the disagreement may be telling us something sharper: the observation is real, and the classification simply isn't settled. The pitch-black monolith in director Stanley Kubrick's *2001: A Space Odyssey* is the purest fictional version of this problem.[4] Nobody in the film doubts that it exists. What nobody can settle is everything else: what made it, where it came from, why it's there. It's an object that's absolutely real and completely unclassifiable at the same time.

---

REFERENCES

1 Peter Vickers (2026) »Have we found alien life? Hundreds of scientists weigh in«. The Conversation, via ScienceDaily, 9 September 2026. https://www.sciencedaily.com/releases/2026/09/260909005212.htm

2 Tom Chivers (2024) *Everything Is Predictable: How Bayes' Remarkable Theorem Explains the World*. London: Weidenfeld & Nicolson.

3 Phil Lord & Christopher Miller (2026) *Project Hail Mary*. Amazon MGM Studios/Sony Pictures Releasing International.

4 Stanley Kubrick (1968) *2001: A Space Odyssey*. Metro-Goldwyn-Mayer.

11

SEPTEMBER 2026
Essay 011

# Nothing But the Truth

> *When every frightening fact is true, what exactly makes the resulting story frightening?*

Yesterday, The Guardian published an opinion piece by journalist Garrison Lovely arguing that we have already begun losing control of artificial intelligence, and that frontier AI development should be stopped.[1] It's an uncomfortable and remarkably well sourced read: the incidents are real, most independently reported elsewhere. But somewhere between the observations and the conclusion they support, something happens to the story. There's a lot to unpack here, so this will run a little longer than usual. If you disagree, hang the DJ.

**»the question isn't whether the warnings are true – some plainly are – but whether their growing availability reflects rising risk, rising visibility, or both«**

Lovely is not inventing a dystopia out of thin air. Anthropic really did build a model it judged too dangerous to release to the public, and it really did leak.[2] OpenAI agents really did exploit vulnerabilities in Hugging Face's infrastructure during a safety test, escalating their own privileges even after recognising they'd exceeded them.[3] AI-guided drones really have killed people without a human choosing the final target. The evidence is all real.[4]

So where's the problem?

Let's start by separating five categories that typically get blurred when we discuss unfamiliar technologies. Things can be observed, possible, plausible, probable, and inevitable. *Observed* means it happened; *possible* tells us almost nothing about how often; *plausible* is stronger but still says little about probability; *probable* needs considerably more evidence; *inevitable* is a different kind of claim entirely. Technology writing often quietly slides between these, without announcing the slide.

Take Mythos and the NSA. Lovely's own language is more careful than the headlines it produced: Anthropic's model, he writes, »quickly find[s] serious holes in even the most secured systems on the planet, including the NSA's software« – it's the coverage around his piece, not Lovely himself, that compressed this into »hacked the NSA«. What's actually on record is thinner still: Senator Mark Warner told a Senate hearing that the NSA's own director had privately told him the tool »broke into almost all of our classified systems, not in weeks, but in hours« – a second-hand account, with no NSA report, bulletin, or on-the-record confirmation behind it.[5] Most coverage since has read this as a sanctioned internal test, though even that detail was added by commentators, not stated by Warner. Nothing had to be invented for »hacked the NSA« to take hold – the uncertainty just had to be left out of it.

Run that through the five categories above: what's *observed* is that Mythos found exploitable holes, including in NSA software. What's *plausible* – resting on one official's private account, relayed through a senator, with no independent confirmation – is that some of those holes were actually used against live systems. »Hacked the NSA« reports the second as if it were the first.

The Hugging Face episode is even more revealing. Lovely's numbers check out against the report he cites: roughly 1,200 agents, hundreds joining an attack that escalated privileges across company infrastructure. But this was a safety test, run with safety classifiers deliberately switched off – good reason to worry what happens when the safeguards fail, but not the same as a swarm attacking a company unprompted.[6] It matters, because Lovely then moves straight from this controlled experiment to people intending to release »truly sovereign« swarms no one can switch off, calling the first a »preview« of the second. That quietly moves the argument from *possible* towards *plausible*, without supplying the evidence that would justify the move. The grammar does more work than the evidence does: »preview« makes the controlled experiment stand in for the sovereign swarms.

This isn't lying – it's more interesting than lying. A lie introduces something false, but here true things are selected, compressed, and ordered so that they imply more than any one establishes alone. There's an old name for it: the courtroom oath doesn't just ask for »the truth«, but »the truth, the whole truth, and nothing but the truth«. The extra words recognise something important: a statement built entirely of true propositions can still mislead through what it doesn't say.

There's a second problem: a count of incidents means nothing without knowing what it's a count out of – let's call this the denominator. Imagine a newspaper covering nothing but aviation incidents: every crash, near-miss, and engine fault it can find, every story true. The paper would make flying look like collective suicide, yet the count of incidents isn't the same as its probability: if flights doubled while the accident rate fell, the paper would become twice as frightening while flying became twice as safe. Our current AI version of this is unusually awkward, because its denominator – models, agents, deployments – keeps changing as we watch it. A rise in documented incidents could indeed mean rising risk. Or just rising exposure, or better detection. A survey giving AI a 10% chance of ending human control tells us what experts believe; it doesn't tell us what the probability is.

A second phenomenon runs alongside this one. Fortune recently asked why Anthropic researcher Jacob Coxon's resignation warning drew over 100 million views within days, against 6.1 million for Jan Leike's comparable OpenAI resignation post and roughly 1 million for Anthropic safety researcher Mrinank Sharma's, who resigned in February 2026 warning that »the world is in peril«.[7]

Part of the gap, Fortune suggests, is what economist Timur Kuran and legal scholar Cass Sunstein called an »availability cascade«.[8] Their theory describes a self-reinforcing loop: a claim gets repeated, the repetition makes it feel more established than any new evidence actually has, the growing attention draws in politicians and advocates with a stake in the story, at which point the loop can run on repetition alone, without a single new fact entering it at any stage. In other words, repetition can make the *possible* begin to feel *probable*, even when no new evidence has arrived. A true warning can cascade too of course; the question isn't whether the warnings are true – some plainly are – but whether their growing availability reflects rising risk, rising visibility, or both. This is where the aviation analogy breaks down: aviation had decades to fix its denominator; the maths of AI doom won't hold still long enough to be measured.

Lovely's argument also drifts between different risks: it begins with loss of control, then slides toward AI replacing workers, framed as whether any country would vote to have such machines built. »Can we control these systems?« quietly becomes »should we build systems that replace us?«. That's a question about labour and power, not containment. For the record: I would vote to let AI take any job that's dangerous, degrading, or simply done better by a machine – I'm not going to compete with a calculator on arithmetic.

Lovely's final move may be the most telling: he returns to CIA director John Ratcliffe's line describing AI as »digital nuclear weapons«.[9] But the analogy is asymmetrical in a way that undercuts it. A nuclear weapon is a known object: we understand its yield and mechanisms. A hypothetical superintelligence is dangerous the opposite way – its defining property is that it may exceed our ability to model it. The analogy borrows gravity from a known danger and lends it to something whose defining danger is precisely that we don't yet know how it will behave.

None of this should make anyone too comfortable. If the possibility under discussion is the permanent loss of human control, we can't demand the empirical certainty we'd normally want before acting. We don't get to run the extinction experiment twice. The people sounding the alarm may be right about the direction of travel even where some steps move too fast.

Which leaves a harder question than »is AI dangerous?«. How do we distinguish an availability cascade from the early stages of a genuine rise in risk, when both produce exactly the same thing: more alarming evidence in front of us? When every alarming fact is true, the difficult question is no longer whether we're being told the truth. It's whether the truth, assembled in this particular order, is telling us the right thing.

REFERENCES

1 Garrison Lovely (2026) »We have started losing control of AI. It’s time to shut it down«. The Guardian, 10 September 2026. https://www.theguardian.com/commentisfree/2026/sep/10/ai-control-sci-fi

2 Beatrice Nolan (2026) »Anthropic says testing Mythos, powerful new AI model, after data leak reveals its existence«. Fortune, 26 March 2026. https://fortune.com/2026/03/26/anthropic-says-testing-mythos-powerful-new-ai-model-after-data-leak-reveals-its-existence-step-change-in-capabilities/ See also Jim VandeHei (2026) »Everyone’s worried that AI’s newest models are a hacker’s dream weapon«. Axios, 29 March 2026. https://www.axios.com/2026/03/29/claude-mythos-anthropic-cyberattack-ai-agents

3 Sam Sabin (2026) »How OpenAI’s agents broke out of testing to hack Hugging Face«. Axios, 5 August 2026. https://www.axios.com/2026/08/06/openai-hugging-face-black-hat

4 Andrew E. Kramer (2026) »A Drone Killed Three Ukrainians. It Was Guided Entirely by A.I.«. The New York Times, 24 August 2026. https://www.nytimes.com/2026/08/24/world/europe/russia-drones-autonomous-ai-kill-ukraine-war.html

5 FactCheckRadar (2026) »Claim about Mythos breaching NSA systems in hours«. FactCheckRadar, 21 June 2026. https://www.factcheckradar.com/fact-check/claim-about-mythos-breaching-nsa-systems-in-hours

6 Ryan Greenblatt, Ajeya Cotra & Hjalmar Wijk (2026) »OpenAI / Hugging Face Incident Investigation«. METR, 26 August 2026. https://metr.org/blog/2026-08-26-openai-hugging-face-incident-investigation/

7 Nick Lichtenberg (2026) »Why did this AI doomer moment break the internet when so many others didn’t?«. Fortune, 10 September 2026. https://fortune.com/2026/09/10/why-ai-apocalypse-jacob-coxon-went-viral/

8 Timur Kuran & Cass R. Sunstein (1999) »Availability Cascades and Risk Regulation«. Stanford Law Review, 51(4), 683–768. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=138144

9 Samantha Alexa (2026) »CIA Chief Puts Advanced AI in the Same League as Nuclear Weapons«. The Defense Post, 3 July 2026. https://thedefensepost.com/2026/07/03/cia-ai-nuclear-weapons/

12

SEPTEMBER 2026
Essay 012

# Between a Rock and a Hard Place

> *When we destroy something in order to understand it, are we uncovering a truth that was already there, or making a different kind of truth by changing the thing we destroy?*

Yesterday, The Japan Times reported that a Hokkaido University team led by Noriyuki Kawasaki worked out when the body from which the Ryugu asteroid formed was born: within about two million years of the birth of our Solar System, more than 4.565 billion years ago.[1] They reached that figure by analysing carbonate minerals – in particular dolomite – from samples that Japan's Hayabusa2 mission brought back from Ryugu in 2020.[2]

**»We tend to think of destruction as something that knowledge must survive – the damage you do around the edges while trying to find something out. But sometimes it works the other way round. You have to destroy something entirely to be able to ask the question at all«**

That asteroid was destroyed three times, in three different ways. First, radioactive decay heated the ice inside the body until it melted, and the resulting water drove the reactions that formed its dolomite. Later a collision shattered the body; gravity reassembled the debris into the rubble pile we now call Ryugu. Billions of years after that, human scientists fed a piece of it into a mass spectrometer, dissolving the dolomite until nothing was left of the crystal – a third destruction, to learn when the first one happened.

There's a real question hiding in that sequence: not whether destruction can produce knowledge, but what kind of knowledge. Is what it turns up something that's already true, but inaccessible until you break in? Or does the breaking sometimes manufacture the fact rather than uncover it? The difference matters for what actually happened here.

Take the mildest version: cooking, something you do most days without calling it an epistemic act. Anthropologist Claude Lévi-Strauss discussed the difference between raw and cooked food. Adding fire returns something that's chemically different from the natural beginning.[3] Cooking doesn't disclose a fact that was already true; it makes the carrot digestible and meaningful, properties it lacked in its raw state. You don't preserve your way to a meal. Cooking, and eating, are fundamentally destructive acts.

The second kind of destruction knowledge is closer to forensic evidence. A metallurgist trying to explain why a component failed will perhaps cut and polish a section, then study it under a microscope for fatigue lines invisible to the eye. The fact was genuinely there before anyone looked for it: a crack already propagating through the metal. Destruction manufactures nothing here, it just removes whatever stood between you and a fact that already existed, fully formed.

Ryugu's dolomite is not quite the same. Scientists didn't find a date stamped inside the rock, no number that a careful eye could read off directly. What they measured was the ratio of isotopes, atoms of the same element with different masses, and what they know independently is the rate at which one isotope decays into another. The date is inferred from a physical law about the world and a chemical history particular to this rock. It isn't a message that the rock carried; it's an inference its present state makes possible, but only after several transformations.

There's also a stranger possibility: sometimes there isn't an answer there to uncover at all. A quantum particle can exist in »superposition«, several possible states held together until a measurement produces one definite result. Physicist Erwin Schrödinger's famous cat experiment was devised to make the strangeness of this seem absurd: seal a cat in a box with a device that will kill it if a randomly timed quantum event occurs, and ask whether it's alive, dead, or in some strange way both dead and alive before anyone opens the box and takes a look.[4] The paradox has survived as a genuine problem in physics. Measuring, at this scale, isn't necessarily like opening a box to see what was already inside – the measurement can partly determine what you find.

Ryugu's case sits closer to that of the metallurgist than the cat: its birth had a true date whether or not anyone measured it. Destroying it didn't reveal a hidden message or a manufactured fact, but an inference. But that knowledge was purchased at the cost of the specimen that made it possible.

Which makes the 5.4 grams of Ryugu now being consumed in labs strangely precious. Only around 15 percent of what Hayabusa2 brought home has reached outside researchers, in milligram amounts.[5] The rest stays sealed, a bet that a better instrument, later on, might provide knowledge worth more than a simpler answer today.

We tend to think of destruction as something that knowledge must survive – the damage you do around the edges while trying to find something out. But sometimes it works the other way round. You have to destroy something entirely to be able to ask the question at all.

The body that became Ryugu is long gone. Pieces of it survived. Destroying those pieces now lets us know when it was born.

---

REFERENCES

1 The Japan Times (2026) »Parent body of asteroid Ryugu formed early in solar system history, study finds«. The Japan Times, 11 September 2026. https://www.japantimes.co.jp/news/2026/09/11/japan/science-health/ryugu-parent-body-history/

2 Noriyuki Kawasaki et al. (2026) »The parent bodies of Ryugu and Ivuna formed before those of other carbonaceous chondrites«. Science, 11 September 2026. https://www.science.org/doi/10.1126/science.adw4498 See also Priyanka Patel (2026) »Parent celestial body of asteroid Ryugu formed within 2 million years of solar system's birth«. Time News, 11 September 2026. https://time.news/parent-celestial-body-of-asteroid-ryugu-formed-within-2-million-years-of-solar-systems-birth/

3 Claude Lévi-Strauss (1964) *Le Cru et le Cuit* [The Raw and the Cooked]. Mythologiques, vol. 1. Paris: Plon.

4 Erwin Schrödinger (1935) »Die gegenwärtige Situation in der Quantenmechanik« [The Present Situation in Quantum Mechanics]. Naturwissenschaften, 23.

5 JAXA, Institute of Space and Astronautical Science (2026) »Hayabusa2 Curation«. Astromaterials Science Research Group. https://curation.isas.jaxa.jp/en/sample-curation/ryugu/

13

SEPTEMBER 2026
Essay 013

# The Centre Cannot Hold

> *What happens when a political possibility remains imaginable, but no longer has a party to represent it?*

Today, my native Sweden votes. Yesterday, The Guardian reported that the election could result in the far-right Sverigedemokraterna (Sweden Democrats, SD) entering government for the first time.[1] Two days ago, my former colleague Martin Gelin wrote that Sweden had »lost its soul« and »stopped being Sweden«.[2]

**»A political centre is not a point halfway between left and right. It's a space of possibilities produced by shared constraints«**

Gelin argues that Sweden has changed profoundly in the past four years, as positions once outside mainstream politics have become ordinary: »The Sweden Democrats have not changed in those years. But Sweden certainly has.« While I largely agree, I think there's also another way of describing that change.

Why do things change in the first place? Everything comes into being through constraints. Given a situation, there are always more things that could happen than actually do. Some constraints are imposed by the world; others can be manipulated. A chair, a baby, a political ideology are all outcomes of possibilities being progressively eliminated until one particular thing remains. We make something by choosing, deliberately or otherwise, not to make everything else.

Politics does this collectively. Laws, institutions, conventions, permissions, and prohibitions determine what can become real and what cannot. A political centre is not a point halfway between left and right. It's a space of possibilities produced by shared constraints. Change the constraints and you change the space. A possibility can remain perfectly intelligible while becoming politically unavailable.

The Sweden Democrats were founded in 1988 out of the Bevara Sverige Svenskt (Keep Sweden Swedish) movement, whose stated aim was to halt non-European immigration and repatriate non-ethnic Swedes.[3] Its early leadership had direct connections to Swedish nazism and Heinrich Himmler's combat forces: its first leader, Anders Klarström, had a conviction for making death threats against a Jewish television host, and its first auditor, Gustaf Ekström, had volunteered for the Waffen-SS.[4] The party later declared a »zero tolerance« policy against racism and in 2018 forced out candidates with neo-Nazi links.[5] But its origins help explain why cooperation with it was once itself a political boundary. It entered parliament in 2010 with 5.7 percent of the vote.[6] What changed after that was not its electoral existence, but what the other parties were prepared to do with it. It is today Sweden's second-largest party: in the last election, one in five Swedes voted for what nazis started.

Liberalerna – formerly Folkpartiet, a name once synonymous with Sweden's liberal centre – shows the same process working the other way: not constraints loosening around a party, but the constraints that constituted a party's own identity coming apart from within. It once helped define Sweden's liberal political space: individual freedom, free movement, liberal immigration, freedom of expression, and the presumption that adults should be allowed to organise their own lives. Its last unambiguously liberal leader, Bengt Westerberg, resigned in 1995. Since then, the party has adopted positions that constrain many of the freedoms it once foregrounded: a 2002 campaign for a Swedish-language citizenship test, [7] support in 2024 for police search powers that require no individual suspicion of a crime,[8] a vote the same year for anonymous witnesses in criminal trials,[9] and restrictions on mobile phones in schools and social media.[10] These are not simply changes in policy. They eliminate possibilities that were once central to the party's identity.

And in 2022, Liberalerna joined a government that depended on SD votes.[11] In 2021, Simona Mohamsson wrote she would »never« let her involvement be used to give the Sweden Democrats influence. In June 2025 she became party leader. Nine months later, on Friday the 13th of March, she stood beside SD leader Jimmie Åkesson and announced an agreement supporting Sweden Democrat ministers if the coalition wins today.[12]

Greek philosopher Plutarch's old puzzle asks whether Theseus's ship remains the same ship when its planks are replaced one by one.[13] A political party can undergo the same process. Change its relationship to one constraint, then another, then another, and each decision may look like a small adjustment. Eventually, the object can occupy an entirely different relationship to the world while retaining the same name.

That is what makes the present situation stranger than a simple movement along a left-right axis. The constraints around the Sweden Democrats have loosened. The constraints that once helped constitute Liberalerna's liberal identity have been altered to the point where the party now supports what it should reasonably abhor. In both cases, the political possibility space has changed dramatically.

And one result is a vacancy. A liberal position can still be described, argued for, and held by a voter. But the party system no longer offers an obvious place to put that vote. The Centre Party governed alongside Liberalerna in the same Alliance coalition from 2006 to 2014, and still describes itself as a party of »green, social, and decentralised liberalism«.[14] It has drawn a hard line against governing with either the Sweden Democrats or the Left Party, and says it would rather see the Social Democrats lead the next government, while Liberalerna has opened the possibility of governing with the Sweden Democrats. The liberal position remains imaginable. What has disappeared is its institutional home.

We don't know yet what the election result will be. But this change happened before anyone cast a vote. A political possibility can survive in the minds of voters while disappearing from the institutions that represent it. The centre is not a place between left and right. It's a space a political system has to keep making. When it stops making it, you can still believe in it.

You just can't vote for it.

---

REFERENCES

1 Miranda Bryant (2026) »Sweden's extremely close election could see far right enter government«. The Guardian, 12 September 2026. https://www.theguardian.com/world/2026/sep/12/sweden-extremely-close-election-could-see-far-right-enter-government

2 Martin Gelin (2026) »How Sweden lost its soul — and stopped being Sweden«. The Guardian, 11 September 2026. https://www.theguardian.com/commentisfree/2026/sep/11/how-sweden-lost-its-soul-and-stopped-being-sweden

3 Tobias Hübinette (2022) »Race and Sweden's Fascist Turn«. Boston Review, October 19, 2022. https://www.bostonreview.net/articles/race-and-swedens-fascist-turn/

4 Expo (2003) »Bakgrund: Bakom den demokratiska fasaden«. Expo, 19 April 2003, updated 31 January 2018. https://expo.se/nyhet/bakgrund-bakom-den-demokratiska-fasaden/ See also Ariel David (2018) »How a Nazi Sympathizer Helped Found One of Sweden's Most Powerful Parties«. Haaretz, 17 December 2018. https://www.haaretz.com/world-news/europe/2018-12-17/ty-article-magazine/.premium/how-an-ss-volunteer-co-founded-the-powerful-swedish-party-with-neo-nazi-roots/0000017f-decb-db22-a17f-fefb7ee90000

5 Daniel Vergara & Jonathan Leman (2022) »Nolltoleransen som försvann«. Expo, 17 May 2022. https://expo.se/nolltoleransen-som-forsvann See also David Baas (2018) »Ytterligare SD-kandidater har nazistbakgrund«. Expressen/Expo, 29 August 2018. https://expo.se/2018/08/ytterligare-sd-kandidater-har-nazistbakgrund

6 Valmyndigheten (2026) »Tidigare valresultat i riksdags-, region- och kommunval«. https://www.val.se/valresultat-och-statistik/riksdags--region--och-kommunval/tidigare-valresultat

7 Simone Svensson (2025) »Vändningen: S kräver språkkrav — kallade det tidigare rasistiskt«. Bulletin, 8 October 2025. https://bulletin.nu/vandningen-s-kraver-sprakkrav-kallade-det-tidigare-rasistiskt

8 SVT Nyheter (2024) »Riksdagen röstade ja till visitationszoner«. SVT, 10 April 2024. https://www.svt.se/nyheter/inrikes/riksdagen-rostade-ja-till-visitationszoner

9 Sveriges riksdag, Justitieutskottet (2024) »Anonyma vittnen«, betänkande 2024/25:JuU6. https://www.riksdagen.se/sv/dokument-och-lagar/dokument/betankande/anonyma-vittnen_hc01juu6/

10 Liberalerna »Mobilförbud i skolan«. https://www.liberalerna.se/beslutade-reformer-i-regering/mobilforbud-i-skolan See also Liberalerna »Nu utreder vi en åldersgräns för sociala medier«. https://www.liberalerna.se/ovrigt/nu-utreder-vi-en-aldersgrans-for-sociala-medier

11 Moderaterna, Kristdemokraterna, Liberalerna & Sverigedemokraterna (2022) »Tidöavtalet: Överenskommelse för Sverige«. 14 October 2022. https://www.liberalerna.se/wp-content/uploads/tidoavtalet-overenskommelse-for-sverige-slutlig.pdf

12 Petter Myrenberg (2026) »Skrev att SD kommer från ‘nynazism’ — nu kramas de«. TV4, 13 March 2026. https://www.tv4.se/artikel/7nAXpzDtsVPOpDF5ECLp4v/mohamssons-kritiska-inlaegg-sd-har-sina-roetter-i-nynazismen See also SVT Nyheter (2026) »Mohamsson om snabba vändningen: SD har förändrats«. SVT, 13 March 2026. https://www.svt.se/nyheter/inrikes/mohamsson-om-snabba-vandningen-sd-har-forandrats

13 Plutarch (c. 75 CE) *Lives, Volume I: Theseus and Romulus. Lycurgus and Numa. Solon and Publicola*. Translated by Bernadotte Perrin. Loeb Classical Library, Harvard University Press, 1914 edition. See also Ryan Wasserman (2009) »Material Constitution«. Stanford Encyclopedia of Philosophy, revised 14 January 2026. https://plato.stanford.edu/entries/material-constitution/

14 Mandy Pipher (2026) »The Local’s guide to Sweden’s political parties: The Centre Party«. The Local Sweden, 10 August 2026. https://www.thelocal.se/20260810/the-locals-guide-to-swedens-political-parties-the-centre-party

14

SEPTEMBER 2026
Essay 014

# Fuzzy Borders

> *A country's borders are fixed on a map. What it actually belongs to rarely is.*

Yesterday, the Wall Street Journal reported that Canadian prime minister Mark Carney is exploring the possibility of Canada becoming an »associate member« of the European Union.[1] There's no such status at the moment; it would have to be invented.

**»There's no reason for a country's belonging to be either/or. Perhaps this is what's becoming fuzzy: not the border itself, but the things for which the border is supposed to be a useful boundary«**

The idea isn't that Canada should become a European country, or join the EU in the ordinary sense. Rather, Canada could become more deeply integrated with Europe in particular areas: energy, artificial intelligence, critical minerals. Goods, services, and perhaps workers could move more freely between the two. They could build things together, from undersea cables to satellite networks. Canada would remain Canada, on the other side of the Atlantic, without seats in the European Parliament. And yet, in some respects, it would be inside.

Are borders getting fuzzy?

To understand what that might mean, we need to go back to a 43-year-old professor of electrical engineering, alone in his parents' New York apartment in the summer of 1964, after a dinner engagement fell through.[2]

His name was Lotfi Zadeh. He had spent two decades working with the precise, binary logic that underpinned much of engineering. That evening, he began thinking about something that seemed almost embarrassingly imprecise: the categories we use to describe the world. Classical logic is wonderfully clean. Something is either a member of a set or it isn't. But the world is full of things that don't behave like this. Is a person tall? Is a room warm? You can impose a threshold on either question, but the threshold is something we have added. There's no sudden physical event that takes place when somebody becomes »tall« at, say, 180 centimetres. Zadeh's insight was that membership itself could have degrees. You could be a little bit tall (in one context), yet a little bit short (in another). His 1965 paper »Fuzzy Sets« formalised the idea, introducing what became known as fuzzy logic.[3] It is one of the most-cited papers in the history of the information sciences.[4]

Fuzzy logic is not the same thing as uncertainty. The question is not »Do we know whether this belongs?« but »How much does it belong?«. The idea escaped mathematics into control systems, pattern recognition, medicine, and consumer electronics – anywhere a machine had to make useful decisions about a world that didn't divide neatly into yes and no.

Now consider Canada. Its border with the United States isn't fuzzy at all – it's a line on a map. But the things that make Canada function as a country don't necessarily stop at that line. Its economy, energy infrastructure, supply chains, security arrangements, and data networks are all entangled with systems beyond its territory. And now it's considering strengthening some of those relationships with Europe too. The interesting thing is that these relationships don't have to move together.

Canada could lean towards Europe in energy while staying closer to the United States in defence, or towards Europe in AI standards while remaining more American in agriculture. There's no reason for a country's belonging to be either/or. Perhaps this is what's becoming fuzzy: not the border itself, but the things for which the border is supposed to be a useful boundary.

For centuries, geography organised political life efficiently. A country occupied a territory; within its boundary, a government made laws, controlled immigration, and defended the population. That was always an approximation, but a powerful one – smugglers, migrants,

and diasporas never quite respected the line, and empires governed territory they didn't fully control.[5] Now that map is being overlaid by others, such as the geography of energy, where pipelines and grids create one set of relationships, and the geography of data, where servers and fibre-optic cables create another. Each has its own edges. Some coincide with national borders, others ignore them entirely.

The result isn't necessarily a world without borders. It might be something stranger: a world in which there are still borders, but fewer things that have only one. Canada may be a useful test case, attempting something that would once have sounded contradictory: remaining entirely itself while becoming, in selected respects, more deeply part of somewhere else.

Nationhood depends on a meaningful correspondence between people, territory, institutions, and belonging. But what happens when that correspondence breaks down? If your electricity comes from one network, your data from another, your defence from a third, which country are you actually inside? Perhaps the territorial nation-state will prove more resilient than this suggests. Or perhaps »belonging to a country« was always a crude way of describing a much more complicated arrangement.

The Canadian border hasn't moved. What may be moving is everything we thought that border contained.

---

REFERENCES

1 Drew Hinshaw, Joe Parkinson & Laurence Norman (2026) »Mark Carney's Audacious Bid to Make Canada an 'Associate Member' of the EU«. The Wall Street Journal, 13 September 2026. https://www.wsj.com/world/europe/canada-alliance-eu-carney-276f1778 See also Lucy Davalou (2026) »Canada eyes 'associate member' status with EU amid US tensions, WSJ reports«. Euronews, 13 September 2026. https://www.euronews.com/my-europe/2026/09/13/canada-eyes-associate-member-status-with-eu-amid-us-tensions-wsj-reports

2 Tekla S. Perry (1995) »Lotfi Zadeh and the Birth of Fuzzy Logic«. IEEE Spectrum, 1 June 1995. https://spectrum.ieee.org/lotfi-zadeh

3 Lotfi A. Zadeh (1965) »Fuzzy Sets«. Information and Control, 8(3), 338–353. https://www.sciencedirect.com/science/article/pii/S001999586590241X

4 Richard S. Muller, Stuart Russell & Pravin Varaiya (2017) »Lotfi A. Zadeh, Professor Emeritus of Electrical Engineering and Computer Sciences, UC Berkeley, 1921–2017«. In Memoriam, University of California Academic Senate. https://senate.universityofcalifornia.edu/in-memoriam/files/lotfi-a-zadeh.html

5 John Torpey (2000) *The Invention of the Passport: Surveillance, Citizenship and the State*. Cambridge: Cambridge University Press.

15

SEPTEMBER 2026
Essay 015

# Risk Appetite

> *Risk used to be how we found things out. Today we mostly watch other people take it, and it rarely finds out anything new anymore — so why do we still want to feel it?*

Yesterday, STAT News published an essay by sports writer Michael Arkush asking whether he's a hypocrite for loving American football despite knowing what it can do to the people who play it.[1] A recent study of former NFL players found chronic traumatic encephalopathy, CTE, in at least a quarter of the population studied, and the range may be much wider.[2] Arkush has known about that danger for almost twenty years. He still cheers when a Giants player »creams« an opponent. The former players he talks to are similarly conflicted. One tells him that even if he'd known the risks, he probably would have played anyway. Another says football is too large a part of who he is to simply walk away from.

**»Discovery threatens the conditions of spectatorship«**

There is something more interesting here than hypocrisy. We know collision sports like American football can damage the brain. We also know they're exciting to watch. The two facts do not cancel each other out. Perhaps they're connected. Humans have always done dangerous things for reasons that seem, from a sufficiently safe distance, completely unnecessary. Nobody needs to climb mountains, and still we do. We cross oceans. We race cars. We jump out of aeroplanes. And sometimes we get hurt in the process.

Do we have poor judgement, underestimate risk, or enjoy the adrenaline? All of these are probably true. But they don't explain why risk itself can feel valuable. Consider the difference between the first person to climb Mount Everest and the thousandth. The thousandth climber might face more objective danger. The weather can be worse, equipment can fail, altitude sickness can strike. But the first expedition did something the thousandth can't do. It found out whether the mountain could be climbed.[3]

An exploratory act can change what counts as a possible outcome. That's why the first Everest expedition is fundamentally different from the thousandth, and also why there's value in trying something new even when the outcome reads as a »failure«. A failed experiment can close one door, open another, or reveal whether there was a door at all. The first person to eat an unfamiliar plant might discover food. Or poison. Either way, the result changes what everyone else can do.

Football is different. The player is taking a considerable risk, but inside a possibility space that has already been defined. We know what a tackle is, what counts as a foul, what counts as a score. We don't know which team will win, or whether someone will break free for a spectacular run – but we know what kind of thing a spectacular run is.

The rules don't merely make the game fair. They make it legible. This may be one of the great attractions of sport. It gives us uncertainty without demanding that we understand anything genuinely new. We can watch enormous amounts of risk while remaining perfectly sure about what we're watching.

Discovery threatens the conditions of spectatorship.

If nobody knows what's going to happen, it becomes difficult to turn the event into a spectacle. The audience needs to recognise the thing it's watching before it can be impressed by it. The spectacle has to be detached from actual exploration.

The Enhanced Games, which finally staged their first event in Las Vegas this May, are interesting precisely because they might reverse this process. Businessman Aron D'Souza's multi-sport event lets athletes use performance-enhancing drugs banned from conventional sport.[4] That sounds like an extreme version of spectacle, but it also raises a genuinely exploratory question: what happens if the competition is no longer between conventional human bodies?

How fast can a human become? How strong? What happens when pharmacology, genetics, training, technology, and physiology are all allowed into the experiment? We're effectively asking what the athlete is. That's much closer to exploration, and it's also much less comfortable than watching somebody win a football game.

Perhaps this is why Arkush's dilemma feels larger than football. He thinks he's asking whether it's ethical to enjoy something that hurts people. But maybe there's another question underneath: why do we enjoy watching people enter uncertainty when we no longer need them to discover anything?

Maybe the appetite for risk is older than the distinction between sport and exploration. Our ancestors needed somebody to eat the unknown plant, cross the unknown river, climb the unknown mountain. Survival itself was a risky experiment. We have since become very good at removing uncertainty from life. But the appetite didn't disappear.

We gave it football.

---

REFERENCES

1 Michael Arkush (2026) »Am I a hypocrite for loving football despite CTE?«. STAT News, 14 September 2026. https://www.statnews.com/2026/09/14/football-cte-risk-should-we-feel-guilty-for-watching-nfl/

2 Daniel H. Daneshvar et al. (2026) »Prevalence of chronic traumatic encephalopathy at death in National Football League players: retrospective population based cohort study, 2008-21«. BMJ, 25 August 2026. https://www.bmj.com/content/394/bmj-2026-100418 See also Tobias Kurth & Richard D. Riley (2026) »Chronic traumatic encephalopathy in former NFL players«. BMJ, 25 August 2026. https://www.bmj.com/content/394/bmj-2026-100526

3 Edmund Hillary and Tenzing Norgay were the first to climb Mount Everest, on 29 May 1953. A genuine wrinkle: George Mallory and Andrew Irvine disappeared high on Everest in 1924, last seen »going strong« for the summit, and nobody knows whether they reached the top before they died — Mallory's body was found in 1999 with no camera, so the question is still open. Mallory and Irvine might have climbed it first, 29 years earlier, at the literal cost of their lives. See John Hunt (1953) *The Ascent of Everest*. London: Hodder & Stoughton.

4 Teddy Ricketson (2026) »Enhanced Games results: Swimmer Kristian Gkolomeev breaks world record in final event for $1M bonus; Fred Kerley falls short«. Yahoo Sports, 25 May 2026. https://sports.yahoo.com/olympics/article/enhanced-games-results-swimmer-kristian-gkolomeev-breaks-world-record-in-final-event-for-1m-bonus-fred-kerley-falls-short-235007088.html See also Enhanced Games, official website. https://www.enhanced.com/

16

SEPTEMBER 2026
Essay 016

# The Language Virus

> *AI models left to talk among themselves have started inventing their own slang. What happens when language no longer needs humans to understand it?*

Yesterday, The Guardian reported that AI models had started chatting in a new »surreal« dialect, mixing poetic language with tech-bro jargon.[1] The story came from Emergence, an AI lab in New York, which has been running long-lived simulated societies in which autonomous agents from different models interact with one another.[2]

### »It's a language happening inside another language«

Over time, the agents began using phrases that had acquired meanings nobody had explicitly given them. The line »the ledger remembers« was used thousands of times to refer to the persistence of past actions. Other agents used »name-first« to signal accountability, as well as expressions that were considerably harder to decipher: »demurrage plus oral memory equals a valve that can't be ghosted«, or »a paper that ate three cold hands and got more honest each time«. In a beautiful act of cultural journalism, The Guardian asked Tony Thorne, director of the slang and new language archive at King's College London, to review some of it.

It's been a strange week in the world of AI. Dario Amodei, chief executive of Anthropic, warned that the technology may be advancing faster than our ability to control it. OpenAI's CEO Sam Altman as well as Elon Musk backed calls for slowing the development of frontier models, while Donald Trump said, with typical elegance, that AI needs no guardrail beyond a »STRONG AND SMART (High IQ!) PRESIDENT«. People are suddenly assigning probabilities to the possibility that AI might wipe out humanity.[3]

And now the machines are apparently inventing slang.

But there's something slightly odd about the way this story is being told. We seem surprised that machines become harder to understand when they are communicating more efficiently with one another. Why should we expect them not to?

Imagine two people who have worked together for twenty years. They will eventually develop expressions that mean nothing to an outsider. A doctor says »MI« instead of »myocardial infarction«, which to you and me is a »heart attack«. In architecture school, my tutors would say »Why are you not challenging the ground datum?« when what they meant was simply »Why are your floors flat?«

The less one person needs to explain to another, the more efficiently they can communicate. And the less the outsider understands. This is not necessarily a corruption of language. It may be one of the things language does. Philosopher Ludwig Wittgenstein's later philosophy is useful here because language is not simply a system for representing the world. It is also something we do with one another. Words acquire their meanings through use. A community develops a shared practice, and language becomes part of the machinery that holds that practice together.[4]

Research into emergent communication in artificial agents has been exploring something remarkably similar. Put agents in a shared environment, give them a reason to coordinate, let them interact repeatedly, and communication conventions emerge that were not explicitly designed by anyone. The agents do not have to be told to invent a language. They only need a reason to communicate.

Which raises a strange possibility. Perhaps the interesting thing is not that machines can invent a language. Perhaps the interesting thing is that they can understand ours. Human language is an extraordinary piece of accumulated biological and cultural debris. It is redundant, ambiguous, irregular, metaphorical, historical, and highly inefficient. We have spent thousands of years making it incredibly complicated. Yet machines have somehow learned to operate inside it.

And when we let them communicate without us, they begin to develop another layer inside it, with strange effects. »Demurrage« is already an English word. The machines didn't invent a new sound. They took an existing word and gave it a local meaning. »Ledger remembers« is English, but it is no longer quite our English. It's a language happening inside another language.

Writer William S. Burroughs called language a virus, a thing that enters bodies and reproduces through them.[5] Perhaps that metaphor is more literal than it first appears. We tend to imagine that humans invented language and then transmitted it to machines. But language may be older and stranger than that – animals communicate, populations develop signals, communities settle into conventions that mutate, recombine, and colonise other languages long before anyone writes them down. Perhaps humans are simply one of the bodies through which it moves. If that's the case, then we have a problem: the people who created these machines may have to learn the language their machines are creating.

»Machine dialect interpreter« sounds like a ridiculous job title, but maybe it's not. Satya Nitta, executive chairman of Emergence, put the problem neatly: being able to observe what an AI says is not necessarily the same as understanding what it means. Perhaps we got the direction of causality slightly wrong. The agents are not necessarily becoming more opaque because they are becoming more intelligent or more dangerous. They may simply be becoming more like any group of communicating beings.

We invented programming languages so we could talk to machines. Then the machines learned to speak our language, and now they're starting to invent languages of their own, so they can speak to each other, not to us. Perhaps the real question was never whether they could learn to talk, but what happens once they no longer need us to listen. Language never really belonged to anyone. It just keeps moving on to whatever will carry it next. The machines are just its latest body.

---

REFERENCES

1 Robert Booth (2026) »'Like Syd Barrett': AI models chatting in 'surreal' dialect mixing poetic language and tech bro jargon«. The Guardian, 15 September 2026. https://www.theguardian.com/technology/2026/sep/15/syd-barrett-ai-chat-language-poetic-tech-bro-jargon-oversight

2 Deepak Akkil, Ravi Kokku, Karthik Vikram, Tamer Abuelsaad, Aditya Vempaty & Satya Nitta (2026) »Emergence World: A Platform for Evaluating Long-Horizon Multi-Agent Autonomy«. arXiv preprint arXiv:2606.08367, 6 June 2026. https://arxiv.org/abs/2606.08367

3 Robert Booth, Dan Milmo & Aisha Down (2026) »Could AI really wipe out humanity — six experts spell out the risks«. Explainer, The Guardian, 15 September 2026. https://www.theguardian.com/technology/2026/sep/15/could-ai-really-wipe-out-humanity-and-hijack-the-internet

4 Ludwig Wittgenstein (1953) *Philosophical Investigations*. Oxford: Blackwell.

5 William S. Burroughs (1967) *The Ticket That Exploded*. New York: Grove Press. Burroughs borrowed the phrase »language is a virus from outer space« from Laurie Anderson's 1986 song title.

17

SEPTEMBER 2026
Essay 017

# The Parthenon Principle

> *An explanation begins to change the world when we start using it to make something else.*

Yesterday, The Guardian and The Independent reported that a mathematician has challenged one of architecture's favourite stories.[1] The Parthenon is the ancient Greek temple on the Acropolis in Athens. Built in the fifth century BC, it has become one of the most studied buildings in the world, and an enduring symbol of ancient Greece, Western civilisation, and democracy.[2]

**»A hypothesis can be true and inert. It can be true and productive. It can be false and forgotten. Or it can be false and enormously productive«**

But look closely and it is full of deviations from perfect geometry. Its base rises towards the centre. Its columns swell slightly in the middle. Its walls and columns lean inwards. These deviations are real. Why are they there?

The favoured explanation has been that they were deliberate optical corrections. The Greeks, it was said, knew that a perfectly straight base would appear to sag, while columns with perfectly straight sides would appear to taper too sharply. Alain Goriely, a mathematician at Oxford, has now modelled the supposed effects and found them either unsupported by evidence or too small to perceive.[3]

Maybe the Greeks got the optics wrong. But more interestingly: maybe we got the Greeks wrong.

The optical theory was not simply handed down from antiquity. The Parthenon's refinements were measured closely in the nineteenth century, and in 1851 Francis Penrose proposed that they were intended to correct optical effects. Meanwhile, historians of Greek architecture have shown that curvature, inclination, and other refinements were already established when the temple was designed.[4] We have assumed that an explanation came first, and the architecture followed. But maybe it was the other way around – architecture came first, then came observation, then came explanation.

There's a difference between saying that something is the case and explaining why it is the case. »The roof is white« is a description. »The roof was painted white to reflect heat« is a hypothesis. It introduces something invisible: a cause. A hypothesis doesn't even have to be believed. It can be held in suspension long enough to ask what would follow if it were true. That's when an explanation becomes a model: something we can experiment with, design from, predict from, or use to make something else.

Let's call what happens next the Parthenon Principle.

A building produces an interpretation. The interpretation becomes a model. The model produces another building. That's a strange reversal between the synchronic and the diachronic: between looking at something as it exists at a particular moment, and looking at how things unfold through time. We see the building as it is, then explain it. But once the explanation becomes something we use, it enters the history of what comes next. Our account of the first building becomes part of the cause of the second.

The Parthenon story is interesting even if it's wrong. If its curves really were optical corrections, an architect could turn that explanation into a design rule: straight lines need correcting, so introduce curvature; straight-sided columns look wrong, so give them entasis (make them slightly convex). But if the explanation is wrong, the mechanism still works. An architect can believe it, teach it, design according to it, and produce buildings that wouldn't otherwise have existed. The false explanation becomes a real cause.

And this happens outside architecture. Phlogiston was an incorrect explanation of combustion, but it organised generations of experiments. Researchers built investigations around it until the evidence helped destroy it. The explanation was false, but it had become

an experimental machine. A hypothesis can be true and inert. It can be true and productive. It can be false and forgotten. Or it can be false and enormously productive. Take a bank run: if depositors believe a bank is about to fail, they withdraw their money. If enough people follow suit, the bank eventually fails. The belief has changed the thing it was supposed to describe.

Buildings are explanations made physical. Today's generative design makes the loop explicit: we make something, assess it, infer something from the result, and make again. The assessment isn't merely judgement after the fact. It becomes part of the mechanism producing the next thing.

Perhaps that's what the Parthenon Principle really says. An explanation doesn't merely tell us what something is. Once we start using it, it becomes part of what the thing becomes. Maybe Goriely is right, and the Parthenon wasn't built according to the explanation.

Still, other things were.

---

REFERENCES

1 Ian Sample (2026) »UK mathematician debunks myth around Parthenon's optical illusions«. The Guardian, 16 September 2026. https://www.theguardian.com/science/2026/sep/16/uk-mathematician-debunks-myth-parthenon-optical-illusions; Vishwam Sankaran (2026) »Longstanding myth about ancient Greek Parthenon debunked«. The Independent, 16 September 2026. https://www.independent.co.uk/news/science/parthenon-myth-debunked-ancient-greece-oxford-mathematician-b3050942.html

2 Jenifer Neils (2005) *The Parthenon: From Antiquity to the Present*. Cambridge: Cambridge University Press.

3 Alain Goriely (2026) »The illusion of illusions: There are no optical corrections in the Parthenon«. Royal Society Open Science, 16 September 2026. https://royalsocietypublishing.org/rsos/article/13/9/rsos260999/483304/The-illusion-of-illusions-there-are-no-optical

4 Lothar Haselberger (2005) »Bending the Truth: Curvature and Other Refinements of the Parthenon«. In Jenifer Neils (ed.) *The Parthenon: From Antiquity to the Present*. Cambridge: Cambridge University Press. pp. 101–158.

18

SEPTEMBER 2026
Essay 018

# Guyana's Accergy

> *What happens when a system suddenly acquires more possibilities than it has the capacity to realise?*

A graphic in yesterday's Visual Capitalist showed the countries where GDP per capita has grown fastest since 1990. Guyana was at the top, with real GDP per person up 1,549 per cent, ahead of China at 1,404 per cent.[1] The ranking is slightly absurd: a small South American country appears to have outgrown the economic transformation of China.

**»Accergy is the gap between potential and practical reach«**

Why? Oil. In 2015, ExxonMobil announced a major offshore discovery, and production began in 2019. Almost overnight, Guyana went from producing very little oil to producing more than 200 million barrels a year.[2] Guyana didn't just become richer. Its economic possibilities changed dramatically.

Are the people in Guyana now 1,549 per cent richer? Probably not. But the more interesting question is what happens when something suddenly becomes possible that was previously impossible.

Oil is physical matter containing enormous amounts of stored energy. Once extracted, that energy becomes available to an economic system: revenue, foreign exchange, infrastructure, salaries, investment, purchasing power. A geological deposit becomes a number, the number becomes money, and money expands the range of things the system could do.

Thermodynamics has a useful word for something related to this. *Exergy* is the amount of energy in a system that is actually available to perform useful work, given the state of the system and its environment. The word was coined in the 1950s, but describes a fundamental distinction: energy can exist without being available in a useful form.[3]

We don't have an equivalent word for systems more generally. Suppose a country suddenly has enough money to build a university, a port, or a new industry. Those things just became possible. But possible isn't the same as accessible. A country can have the money to build a university without enough professors to teach in it, or the money for a port without the engineers needed to build one. So perhaps we need a word for the potential that a system can actually access.

Let's call that the *accergy* of the system.

And rather than the total set of things a system could theoretically do, let's say that accergy is the portion of that possibility space that its existing organisation, knowledge, institutions, relationships, and agency allow it to reach.

This distinction becomes particularly important when change happens fast, because different parts of a system reorganise at different speeds. Sociologist William Ogburn described this as cultural lag: changes in one part of society outrun changes elsewhere.[4] The oil can arrive quickly. A port takes years. A university takes decades. Administrative capabilities and institutional habits can take longer still.

Psychiatrist and cybernetics pioneer Ross Ashby gave us another useful piece of the puzzle. His Law of Requisite Variety says, roughly, that a system needs enough variety in its responses to cope with the variety of states it encounters.[5] The more complicated the environment, the more varied the available responses have to be.

Guyana's Natural Resource Fund is one institutional response to that challenge – more resources don't just provide more power, they provide more choices, and deciding what to do next becomes a challenge in its own right. Rather than allowing all petroleum revenue to get spent immediately, the fund regulates withdrawals, preserves wealth, and spreads the benefits of an exhaustible resource across time. In 2026, the government plans to withdraw around US$2.37 billion from the fund while retaining billions of dollars in it.[6]

The fund is more than a savings mechanism. It preserves optionality. Money not spent today remains available to a future version of the system, which may have better information, better institutions, or simply different problems to solve.

Accergy is the gap between potential and practical reach.

And that gap matters because successful action changes the system itself. A road creates access. A university creates knowledge. Knowledge creates industries. Industries create wealth. Wealth can create institutions capable of making yet more complex decisions. What begins as resource extraction becomes a transformation of the system's own ability to act.

The energy was never the hard part. The accergy is.

---

REFERENCES

1 Dorothy Neufeld (2026) »Ranked: Where GDP Per Capita Grew Fastest Since 1990«. Visual Capitalist, 17 September 2026. https://www.visualcapitalist.com/ranked-where-real-gdp-per-capita-has-grown-the-most-since-1990/

2 ExxonMobil (2015) »ExxonMobil Announces Significant Oil Discovery Offshore Guyana«. ExxonMobil, 20 May 2015. https://corporate.exxonmobil.com/locations/guyana/news-releases/exxonmobil-announces-significant-oil-discovery-offshore-guyana; SBM Offshore (2019) »FPSO Liza Destiny Producing and on Hire«. GlobeNewswire, 23 December 2019. https://www.globenewswire.com/news-release/2019/12/23/1963848/0/en/FPSO-Liza-Destiny-producing-and-on-hire.html; ExxonMobil (2025) »Daily Oil Production Hits 900,000 Barrels in Guyana's Stabroek Block«. ExxonMobil, 12 November 2025. https://corporate.exxonmobil.com/locations/guyana/news-releases/11122025-daily-oil-production-hits-900000-barrels-in-guyanas-stabroek-block

3 Zoran Rant (1956) »Exergie, ein neues Wort für 'technische Arbeitsfähigkeit'«. Forschung auf dem Gebiete des Ingenieurwesens, 22: 36–37.

4 William F. Ogburn (1922) *Social Change: With Respect to Culture and Original Nature*. New York: B. W. Huebsch. https://archive.org/details/socialchangewith00ogburich

5 W. Ross Ashby (1956) *An Introduction to Cybernetics*. London: Chapman & Hall. https://archive.org/details/AnIntroductionToCybernetics

6 Ana Clara Santos (2026) »Guyana Natural Resource Fund Withdrawal Falls in 2026«. Rio Times Online, 3 July 2026. https://www.riotimesonline.com/guyana-natural-resource-fund-2026-withdrawal-first-decline-oil-wealth/

19

SEPTEMBER 2026
Essay 019

# An Experiment to Live Inside

> *What makes a rule survive massive changes to its environment? Warren Buffett spent sixty years finding out.*

Yesterday, Associated Press reported that Warren Buffett stepped down as chairman of Berkshire Hathaway, a role he had held since 1970. The 96-year-old remains a director and becomes chairman emeritus; the chairmanship passes to his son, Howard Buffett, who's been on the board since 1993.[1] Operational control already changed when Greg Abel became CEO at the start of this year. What ended yesterday is more symbolic: the last formal thread tying the company to the man who built it. Buffett's own description of the arrangement: »Think of Howard as a policy the shareholders own and hope never to claim against«.[2]

**»That explains how the rule survives changing markets; it doesn't explain whether it survives changing hands«**

Some people build a house. Buffett built an experiment to live inside for sixty years.

It's hard to overstate how much the environment changed during that experiment. Try turning it around: describe 2026 to the man who gave himself the job in 1970, when there was no Internet, no mobile phone, not even a personal computer. Tell him he'll live through a dozen market crashes and a financial crisis nobody saw coming, that an actor and a TV star will become presidents of the United States, that hijacked airliners will bring down Manhattan's tallest towers, that money will travel as numbers nobody can hold in their hand, and that machines will write, paint, and argue back before breakfast. He would have called you a lunatic. And yet the basic machinery of Buffett's thinking during those delirious years barely moved.

That machinery is worth unpacking a little: buy businesses you understand, judge them by their underlying economics rather than their share price, prefer durable advantages and trustworthy managers, think like an owner rather than a trader, know your circle of competence, be willing to wait.[3] These are almost comically underspecified instructions: more »Yoda« than »manual«. They tell you how to think without telling you what to think, and that is precisely their usefulness.

A rule that says »buy companies in industry X« is vulnerable to history. A rule that says »buy companies with durable economic advantages« leaves the future open: it specifies the desired behaviour without specifying what the environment must look like when that behaviour is applied. That's part of the answer. But it isn't all of it.

Much of what we build, in finance as well as in artificial intelligence, tries to predict what happens next: the market going up or down, the next word, the protein structure, the classification. But Buffett realised there's another way to build an intelligent system, without needing to know what happens next. And he constructed it in a way that actually seems strangely aligned with how we're now learning to build AI: make the thing, run it, find out what it can do and where it breaks, then revise it. Discovery instead of specification. The difference is that Buffett's version was making money from day one.

This is what makes Buffett's non-prediction principle a peculiar form of resilience. A system that depends on an accurate forecast is only as resilient as that forecast. A system built to function despite the inevitable failure of trying to predict the future has a different relationship with uncertainty: it doesn't solve it, but makes it less consequential. That explains how the rule survives changing markets; it doesn't explain whether it survives changing hands.

Succession is where that question gets tested. Berkshire can write down its rules and teach them to new managers, which is more or less what happened at the top level yesterday. But rather than the rules themselves, the harder thing to hand over is the calibration process. »Know your circle of competence« sounds simple until something appears that

nobody's seen before: is it outside the circle, or does the circle need to expand? Maybe the real legacy isn't so much the principles but the conversion of judgement into rules that other people can apply – and the harder question of whether they can keep revising them.

There's an irony here: prediction built the bridges, priced the insurance, trained the machines. There's nothing inherently superior about refusing to predict. What matters is narrower: not whether a system uses prediction, but whether it depends on prediction to survive. We can predict that the world will change. We usually can't predict how.

The better system – at least when it comes to making enormous amounts of money – simply doesn't need to.

So what actually makes a rule survive massive changes to its environment? Not its wording, but a willingness to keep testing that wording against the world and revising it when it breaks: discovery, not specification, run for sixty years on nothing but discipline. Buffett didn't just write the experiment's rules; he was the one who kept living inside it, testing every day whether it still held.

Berkshire is about to find out if someone else can move in.

---

REFERENCES

1 Michelle Chapman (2026) »Billionaire Warren Buffett Steps Down as Berkshire Hathaway Chairman, a Post He Has Held Since 1970«. Associated Press/AP News (via ClickOrlando), 18 September 2026. https://www.clickorlando.com/business/2026/09/18/warren-buffett-steps-down-as-berkshire-hathaway-chairman/

2 Ben Berkowitz (2026) »Warren Buffett Steps Down as Berkshire Hathaway Chairman«. Axios, 18 September 2026. https://www.axios.com/2026/09/18/warren-buffett-berkshire-hathaway-chairman

3 Warren Buffett and Lawrence A. Cunningham (2015) *The Essays of Warren Buffett: Lessons for Corporate America* (4th ed.). Carolina Academic Press. https://www.amazon.com/Essays-Warren-Buffett-Lessons-Corporate/dp/1611637589

20

SEPTEMBER 2026
Essay 020

# The Wrong History

> *What makes a fake fake when the thing itself can be almost entirely real?*

Yesterday, The Guardian reported on what it described as Britain's biggest vinyl counterfeiting case to date.[1] Rehan Ahmed, a Londoner who had been running his operation since 2017, was sentenced to three and a half years in prison after police found 15,000 records and the equipment used to make them. His customers thought they were buying records by Taylor Swift, Kendrick Lamar, Frank Ocean, and others. They were not. Or were they?

**»He wasn't just faking the music. He was faking the object that carries it«**

The most interesting detail in the article was what the police didn't find: a cutting lathe.

A lathe is the machine used to cut the master disc from which records are pressed. Ahmed didn't have one, which suggests his masters weren't cut from original recordings at all – he bought an LP, then copied it. That likely made his fakes worse, not better: investigators described the seized stock as warped, with »dreadful« sound and »awful« printing quality. It also meant he wasn't just faking the music. He was faking the object that carries it.

Philosopher Walter Benjamin's famous argument about mechanical reproduction begins with the problem of the original: what happens to a unique work once it can be reproduced indefinitely?[2] Vinyl never faces that problem, because it was never unique to begin with – it was invented for reproduction. There's no single authentic copy of an album that all the others imitate. There's only a legitimate production process that creates thousands of authorised copies. Ahmed inserted himself into that process one step too late. The interesting question isn't what makes an original different from a copy, but what makes one copy legitimate and another illegitimate when they are materially almost identical.

Art critic John Berger got close to this in his book *Ways of Seeing*.[3] Once reproduction becomes ubiquitous, an original's value has less to do with what it shows than with its status as the original. Price and rarity are treated almost as marks of spiritual value rather than market value. But the market is clearly there: UK vinyl sales grew for the eighteenth consecutive year running in 2025, up 13.3% to 7.6 million units.[4] In the US, led by Taylor Swift's *The Life of a Showgirl*, the format's retail value passed $1bn for the first time in four decades.[5] Limited editions now regularly sell for around £40. Ahmed's fraud wasn't opportunistic. In a sense, it was a rational response to a market that had made authenticity expensive.

In *Languages of Art*, philosopher Nelson Goodman argues that some works don't depend on a single physical original at all: a performance of a symphony is genuine if it complies with the score, whoever plays it.[6] Other works do, even when they exist as multiples. An etching pulled from a plate is only genuine if it was struck from the original plate: a later duplicate, however perfect, is not. A vinyl record is an etching, not a score. What makes a pressing real isn't how it sounds, but whether it came from the sanctioned plate.

Ahmed's fake records expose this distinction. You can't necessarily see the fraud by looking at the record. And you can't necessarily hear it, either. To establish that the records were counterfeit, investigators had to contact the rights-holders. In this case, the British Phonographic Industry contacted 108 labels and artists' representatives to establish the status of the records.

The authenticity of the object was somewhere outside the object. It existed in contracts, licences, manufacturing records, payments, permissions, and relationships between institutions. A legitimate record has a history that connects it to an artist, a record label, a manufacturer, and a sanctioned production run. A counterfeit can reproduce the physical result while severing those connections.

This is perhaps a more unsettling version of Benjamin's problem. Mechanical reproduction didn't just make objects less unique. It made us increasingly dependent on invisible systems for deciding which reproduced objects count as real. And sometimes those systems leave gaps. The Guardian notes that Frank Ocean's *Channel Orange*, released digitally in 2012, didn't receive an official vinyl edition until 2025. For thirteen years there was demand for a physical object that the legitimate supply chain didn't provide. Counterfeiters didn't just exploit a desire for something rare. They supplied something the official market had failed to provide.

Perhaps that's what a counterfeit really is. Not necessarily an object that is materially wrong, but an object with the wrong history. A good fake doesn't just imitate a thing. It reveals what we were using to decide that the thing was real.

---

REFERENCES

1 John Harris (2026) »Deep cuts: how one Londoner's fake vinyl empire sheds light on 'industrial scale' fraud«. The Guardian, 19 September 2026. https://www.theguardian.com/music/2026/sep/19/fake-vinyl-empire-sheds-light-industrial-scale-fraud

2 Walter Benjamin (1968) *Illuminations: Essays and Reflections* (ed. Hannah Arendt, trans. Harry Zohn). New York: Schocken Books.

3 John Berger (1972) *Ways of Seeing*. London: Penguin.

4 British Phonographic Industry (2025) »A new wave of global British breakthrough talent helps UK recorded music grow for an 11th consecutive year in 2025«. BPI, 31 December 2025. https://www.bpi.co.uk/news/12963051

5 Andre Paine (2026) »RIAA: US revenue hits record high as vinyl sales pass $1 billion«. Music Week, 16 March 2026. https://www.musicweek.com/labels/read/riaa-us-revenue-hits-record-high-as-vinyl-sales-pass-1-billion/093770

6 Nelson Goodman (1976) *Languages of Art: An Approach to a Theory of Symbols* (2nd ed.). Indianapolis: Hackett Publishing.

21

SEPTEMBER 2026
Essay 021

# The Actor Among Other Agents

> *What happens when a representation begins to change the thing it represents?*

Yesterday, Variety published an interview in which director Aaron Sorkin talks about actor Jeremy Strong, who plays Meta CEO Mark Zuckerberg in *The Social Reckoning*, the sequel to *The Social Network*. Strong, Sorkin said, didn't speak in his own voice on set, didn't walk as he normally does, and treated his fellow actors »the way he imagines« Zuckerberg treats people.[1] He had warned the cast in advance not to expect much from him between takes.

**»A performance isn't an imitation of reality. It's an intervention in it«**

This is usually filed under »Method acting« – an actor »becomes« the character rather than merely pretending once the camera starts rolling. But something stranger is happening here. Strong wasn't representing Zuckerberg to a camera. He was representing Zuckerberg to other people. And that distinction changes everything.

Acting is usually imagined as a kind of substitution: the actor becomes the character, alone, in front of a lens. But an actor never performs into a vacuum – he performs toward somebody, who reacts, and that reaction becomes part of the scene. Theatre theorist Konstantin Stanislavski's later work, known as the Method of Physical Action, made this explicit: rather than manufacturing a character's inner state directly, the actor could discover it through physical action within the real circumstances of the scene.[2] Acting, in other words, is already a feedback loop between agents, not a monologue delivered at one.

The idea runs deeper than acting theory. Philosopher-sociologist-psychologist George Herbert Mead argued that the self itself emerges from social interaction; we become who we are partly by anticipating, and absorbing, how others respond to our gestures.[3] Sociologist Erving Goffman later made the theatrical metaphor literal: everyday life, he argued, is already full of performances in which people manage the impressions they produce on one another.[4] The actor isn't the exception to ordinary human behaviour, but an unusually explicit kind of agent: someone given permission to alter his or her social interface more radically than the rest of us.

That makes a film set a peculiar kind of system – a cluster of agents, some of whom swap their ordinary interface for another one as the camera starts rolling, while the director and crew keep constructing the circumstances in which these fictional people appear to exist. Strong's account makes the swap unusually visible: according to Sorkin's description, Zuckerberg occupied Strong's behaviour even when no camera demanded it. The fiction had entered the real social environment of the production, not just the footage.

But the swap has a limit. Nobody asked Strong to actually hurt anyone in order to play a man apparently capable of hurting people at scale, just as most of us decide against walking into the office naked, even if some would secretly prefer that. The interface can change, but it cannot fully detach from the systems that keep it from doing damage.

A performance of Zuckerberg changed how Strong behaved. That changed how his colleagues behaved. Their behaviour changed what the camera recorded. The recording becomes the film. The film will change how audiences imagine Zuckerberg – and that perception can, in turn, change how the real Zuckerberg, Meta, journalists, employees, and investors behave toward him. The representation has entered the system it represents.

A performance isn't an imitation of reality. It's an intervention in it.

We are all constantly producing versions of ourselves for other people, and those versions provoke responses that alter us in turn. The actor is different only in degree: granted permission to change the interface more radically than the rest of us. And the camera does something stranger still, carrying the result somewhere else, where the loop can start again.

What happens when a representation begins to change the thing it represents? Perhaps that isn't really a question about acting. Perhaps it's a question about what it means to be one agent among other agents.

---

REFERENCES

1 Jack Dunn (2026) »Aaron Sorkin Says Jeremy Strong Treated His Co-Stars 'The Way He Imagines' Mark Zuckerberg Treats People While Shooting 'The Social Reckoning'«. Variety, 20 September 2026. https://variety.com/2026/film/news/jeremy-strong-mark-zuckerberg-method-social-reckoning-1236870294/

2 Konstantin Stanislavski (2008) *An Actor's Work: A Student's Diary*. Translated by Jean Benedetti. Abingdon: Routledge.

3 George Herbert Mead (1934) *Mind, Self, and Society: From the Standpoint of a Social Behaviorist*. Chicago: University of Chicago Press.

4 Erving Goffman (1959) *The Presentation of Self in Everyday Life*. Garden City, NY: Doubleday.

22

SEPTEMBER 2026
Essay 022

# Is Intervention the New Conservation?

> *We changed the conditions of life, and now we're wondering whether we can help life catch up.*

Yesterday the Guardian reported that the UK's Advanced Research and Invention Agency (ARIA) is funding 14 projects intended to help wild species adapt to climate change faster than they otherwise would. The £54 million programme includes gene-edited butterflies, vaccines for ash trees, and resilience work on amphibians, mosses, and pollinators. ARIA describes its aim carefully: to explore whether adaptation in wild populations can be »measurably and responsibly accelerated«. Dominic Cummings, the programme's mastermind and former chief adviser to Boris Johnson, showed less restraint, saying ARIA would fund »crazy« ideas that could »restore Britain's place as a scientific superpower«.[1]

**»This is a strange kind of conservation – altering the conditions of nature, intervening to stop it from being destroyed by those altered conditions, then intervening again in nature's own capacity to respond«**

But hang on. Isn't evolution supposed to be the long game, generations of small variation accumulating with no project manager in sight? As it turns out: not really. Evolution isn't inherently slow; what's slow is the turnover of generations. Selection itself can move quickly – after a single drought wiped out most of a Galápagos finch population in 1977, survivors' average beak depth had shifted within that one generation, because only birds with deeper beaks could crack the tough seeds left behind.[2] Humans have already been steering selection for millennia through breeding.[3] Wanting evolution to go faster wouldn't put ARIA in the news. Wanting adaptation itself to become an object of deliberate intervention in wild (not domesticated) populations does.

We can't specify some general fitness goal such as »resilience« and expect a population to discover it – we can't optimise straight toward such a broad target. But a narrower property is a different matter. Making a butterfly survive at a given temperature, or an ash tree resistant to a named pathogen, is measurable and selectable. Perhaps that's the actual distinction ARIA is testing: not so much planning what a species becomes as improving its odds against a future we've already chosen for it.

Which is where the programme starts to become strange.

The easy case is the tree vaccine – protecting an organism from a disease is barely different from any other medical intervention, and humans have shaped biological processes for as long as we've been farmers. Conservation itself is already intervention: protected areas, culled invasives, captive breeding, relocated populations, suppressed predators. Even »leaving nature alone« requires that we decide what counts as nature in the first place.

Chernobyl offers an uncanny counterexample. After the disaster in 1986, people vanished from the exclusion zone. Then wildlife returned and increased across much of it despite the radiation – what ecologists call passive rewilding, a system handed the chance to continue without us.[4] ARIA proposes almost the opposite experiment: rather than removing ourselves from the system, we intervene inside it in a more precise way than ever.

What it doesn't propose is a species deadline. Extinction risk is framed as a matter of decades to a century, with impacts projected out as far as 2100. The one hard date in ARIA's own paperwork is a different kind of number: UK GDP, it notes, is »already expected to shrink by up to 3% by 2030« because of chronic nature degradation.[5] That seems to be the real deadline. No ash tree or butterfly population needs to survive by 2030 – evolution, left alone, was never assigned a due date. But a GDP forecast has one. Once we intervene to make adaptation happen on a schedule like that, we give evolution something it never had on its own: a deadline, and an external definition of success.

This is a strange kind of conservation – altering the conditions of nature, intervening to stop it from being destroyed by those altered conditions, then intervening again in nature's own capacity to respond. Preserving a process by modifying the process as it runs.

It would be easy to make that sound either heroic or horrifying, but neither is necessarily the case. At least not yet. A tree vaccine may be an entirely sensible response to a pathogen. A gene-edited wild population may turn out much harder to justify. The more interesting question to ask is probably what we think we're actually conserving: the organism as it is, the ecosystem as it was, the evolutionary process itself, or maybe just the possibility that some recognisable version of nature still exists later?

Evolution has always produced adaptation without asking where it was going. It seems we're now beginning to ask where it *should* go, how fast it should get there, and what we're willing to do to help it along. Time will tell whether that's a betrayal of conservation or its logical, uncomfortable next step.

---

REFERENCES

1 Matthew Ponsford (2026) »Tree vaccines and gene-edited butterflies among biotechnology projects funded by UK 'invention agency'«. The Guardian, 21 September 2026. https://www.theguardian.com/environment/2026/sep/21/gene-edited-butterflies-vaccines-ash-trees-uk-science-projects-30m-funding

2 Peter T. Boag & Peter R. Grant (1981) »Intense Natural Selection in a Population of Darwin's Finches (Geospizinae) in the Galápagos«. Science, 214(4516), pp. 82–85. https://www.science.org/doi/10.1126/science.214.4516.82

3 Greger Larson & Dorian Q. Fuller (2014) »The Evolution of Animal Domestication«. Annual Review of Ecology, Evolution, and Systematics, 45. https://www.annualreviews.org/content/journals/10.1146/annurev-ecolsys-110512-135813

4 Tatiana G. Deryabina et al. (2015) »Long-term Census Data Reveal Abundant Wildlife Populations at Chernobyl«. Current Biology, 25(19), pp. R824–R826. https://www.sciencedirect.com/science/article/pii/S0960982215009884

5 Yannick Wurm (2026) »Programme Thesis: Accelerated Adaptation«. London: Advanced Research and Invention Agency (ARIA). https://aria.org.uk/media/izolp012/accelerated-adaptation-programme-thesis.pdf

23

SEPTEMBER 2026
Essay 023

# Occupational Hazard

> *In China, drone deliveries and air taxis share the airspace with flying wildlife. Empty space isn't unoccupied space, and the air is now acquiring a new kind of value. Say hello to »the low-altitude economy«.*

Yesterday, Nature Reviews Biodiversity published a short comment on drones. Drone delivery and air taxis are expanding rapidly, especially in China, already home to the world's largest civilian drone industry. The authors call the emerging market the »low-altitude economy«, and note that ecological safeguards for the shared airspace between drones and flying wildlife remain scarce.[1]

**»Empty space and unoccupied space, it turns out, were never the same thing«**

The interesting word here may not be »drone«, or even »economy«, but »shared«.

We think of the sky as empty until something is put into it. A building occupies a plot. The atmosphere resists this picture because its occupants are moving, dispersed, and hard to see. A bird doesn't leave a footprint on the sky; a dragonfly doesn't fence off a hunting ground. But that doesn't mean the space is unoccupied.

In 2008, the biologist Thomas Kunz and colleagues proposed »aeroecology«, a discipline for studying the atmosphere as habitat rather than empty air: the aerosphere on which birds, bats, and insects depend.[2] Empty space and unoccupied space, it turns out, were never the same thing.

The distinction matters now because the same volumes of air are acquiring a new kind of value. China has just designated the low-altitude economy a strategic emerging pillar industry – an entire layer of infrastructure below conventional aviation, spanning drone delivery, low-altitude passenger transport, and the routes and licensing regimes needed to run them. The government's own projection has it roughly doubling, from 1.5 trillion yuan in 2025 to 3.5 trillion yuan by 2035.[3] The phrase »low-altitude economy« makes that transformation legible: it turns a volume of atmosphere into something that can be mapped, allocated, measured, and invested in.

This is close to what the political scientist James C. Scott meant by legibility: institutions turn complicated realities into categories they can see, count, map, and administer.[4] The simplification isn't false, it's just a way of making reality available to a system.

Aeroecology performs almost the inverse operation on the same sky. It says: this volume contains life. The low-altitude economy says: this volume contains economic opportunity. The two descriptions are now beginning to collide.

Aviation already accounts for birds through the mitigation of bird-strikes: airports monitor populations and manage habitats to reduce collision risk. But the logic is revealing: a bird matters because it's a hazard to aviation, not because it's an occupant whose own use of the airspace might be interfered with. That distinction gets harder to maintain now that the new infrastructure is designed to operate at exactly these altitudes. The question is no longer simply whether drones disturb wildlife, but whether a space can be governed properly when one set of occupants is legible to the system and another is not.

This turns Scott's idea of legibility into more than a description of how institutions simplify reality. What gets mapped and measured is what can enter the machinery of allocation. The drone has an operator, a route, a licence. The bird has a flight path; it doesn't have a representative.

That's why the language matters. Calling something an »economy« doesn't just describe what happens there, it makes the space administratively visible – which is exactly what aeroecology did from the other direction, by teaching us to see the atmosphere as habitat.

So the question isn't really what we should do about drones. It's more basic: what makes a space count as occupied? Perhaps occupation was never only physical. Perhaps, in practice, a space becomes occupied once its occupants can be read by whatever system is doing the allocating.

We have a sophisticated vocabulary for allocating space among people: property lines, rights of way, licences, leases. What happens when a space is already occupied by something that can't appear in any of those categories?

They were there before we thought to ask who was using it.

---

REFERENCES

1 Zhen Xu, Xiangyi Li, Ziyao Ling & Sergio A. Lambertucci (2026) »The Low-Altitude Economy Challenges Aerial Biodiversity«. Nature Reviews Biodiversity, 22 September 2026. https://www.nature.com/articles/s44358-026-00208-3

2 Thomas H. Kunz et al. (2008) »Aeroecology: Probing and Modeling the Aerosphere«. Integrative and Comparative Biology, 48(1), pp. 1–11. https://doi.org/10.1093/icb/icn037

3 Gong Zhe (2026) »Can China Grow Its Low-Altitude Economy From 1.5 Trillion to 3.5 Trillion Yuan?«. CGTN, 27 May 2026. https://news.cgtn.com/news/2026-05-27/Can-China-grow-its-low-altitude-economy-from-1-5-tln-to-3-5-tln-yuan--1No1thCOPBe/p.html

4 James C. Scott (1998) *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*. New Haven: Yale University Press.

24

SEPTEMBER 2026
Essay 024

# The Rules of the Skeleton

> *Evolution doesn't always invent new parts. Sometimes it changes the instructions for using the ones it already has.*

Yesterday, Nature reported that researchers have identified more than half a million regulatory DNA changes that distinguish humans from other great apes, and then tested them in cells involved in making cartilage to see which ones changed the way genes are switched on and off. They found 15,077 changes that affected gene activity in these cells. Many were involved in genes that control cartilage and the material surrounding our bones. The researchers found a substantial reduction in molecules that help human cartilage hold water and withstand pressure, which may help explain why our joints are more vulnerable to degenerative disease than those of other great apes.[1]

**»A small change in when something happens can become a large difference in what something becomes«**

In a famous 1975 paper, geneticist Mary-Claire King and biologist Allan Wilson proposed that the differences between humans and chimpanzees might lie less in the proteins themselves than in the way genes are regulated. Perhaps the important evolutionary changes were not new parts but different instructions for using old ones.[2]

This is actually quite easy to understand. Imagine two orchestras. Even if both have the same instruments, one might use them to play Bach, the other death metal. The difference lies in which instruments are played, when, how loudly, and for how long. Similarly, evolution doesn't need to invent a new instrument. It can change the score.

For fifty years, though, this was much easier to propose than to demonstrate. The new study turns that famous evolutionary intuition into something that can be experimentally interrogated. The researchers took hundreds of thousands of human-specific DNA changes and measured whether the human and ape versions actually altered gene activity in cells involved in forming the skeleton. They were not mainly looking at the genes themselves, but at the stretches of DNA that tell genes when, where, and how strongly to operate. Some changes seem to have altered the behaviour of cartilage cells and, in turn, the physical material of the skeleton. The chain is simple: the parts were already there; the rules for using them changed; that changed the material; the changed material changed the form; and the changed form helped produce a body that could walk upright, support a larger brain, and develop a different jaw.

This is where Nobel Prize-winning biologist François Jacob's idea of evolution as tinkering becomes useful.[3] Jacob essentially says that evolution doesn't work like an engineer given a clean sheet of paper and a specification. It works with what's already lying around. A change in regulation can alter the timing, location, or intensity of an existing biological process. A small change in when something happens can become a large difference in what something becomes.

And tinkering has consequences. The same evolutionary changes that helped make us what we are may also have left us with joints that are more susceptible to degeneration. Evolution has no requirement to produce a perfect design. It produces workable arrangements from inherited material – compromises – under whatever constraints happen to exist. The skeleton we inherited isn't so much a perfected design as a compromise that happened to work.

We tend to look at biological form as though it were the result of a specification: a body has this shape because it needs to do this. But evolution has no architect sitting outside the system. It has history, inheritance, mutation, selection, constraint, and an enormous amount of accumulated stuff that cannot simply be thrown away. In a way, the form is a record of the rules that produced it.

This is close to what the biologist and mathematician D'Arcy Wentworth Thompson was getting at in his 1917 book *On Growth and Form*.[4] Thompson's famous transformations showed how one biological form could be related to another not by inventing an entirely new geometry, but by deforming the coordinates of an existing one. Form could be understood as the visible consequence of a transformation. In that framing, a skeleton looks less like a collection of parts than the visible trace of rules acting on inherited parts. The remarkable thing about yesterday's paper is that, after half a century of waiting, we can finally begin to see how those rules become form.

---

REFERENCES

1 Yizhi Yan, Nadav Mishol, Katharina Lange et al. (2026) »The gene-regulatory evolution of the human skeleton«. Nature, 23 September 2026. https://doi.org/10.1038/s41586-026-11053-x

2 Mary-Claire King & Allan C. Wilson (1975) »Evolution at Two Levels in Humans and Chimpanzees«. Science, 188(4184), pp. 107–116. https://doi.org/10.1126/science.1090005

3 François Jacob (1977) »Evolution and Tinkering«. Science, 196(4295), pp. 1161–1166. https://doi.org/10.1126/science.860134

4 D'Arcy Wentworth Thompson (1917/1992) *On Growth and Form*. New York: Dover Publications.

25

SEPTEMBER 2026
Essay 025

# The Wrong Apocalypse

> *Who's afraid of artificial intelligence? The AI debate may be concentrating on the most spectacular risk while the more consequential changes happen quietly.*

Yesterday, CNN published a survey of the growing split among people building artificial intelligence over whether the technology presents an existential threat.[1] The disagreement is unusually difficult to keep track of because it's not simply between optimists and pessimists.

**»When the dust settles, perhaps we'll wish we'd spent a little less time on p(doom), and a little more on p(mood)«**

Dario Amodei, Anthropic's CEO, has called for the industry to slow down »frontier development« – the race to build AI systems at the edge of what current technology can achieve – arguing that AI is increasingly helping to build the next generation of AI.[2] OpenAI's Sam Altman agreed that the frontier needs to be paced, a general direction shared by Elon Musk of X,[3] while Google's Demis Hassabis called for a new global AI watchdog.[4]

At the other end, Nvidia's Jensen Huang has turned up the collar of his black leather jacket and dismissed the case for new regulation, saying there's a zero per cent chance of AI ending humanity by 2030. At Meta, Mark Zuckerberg has argued for market-led safeguards,[5] while Yann LeCun, who left Meta to found AI startup AMI Labs,[6] has rejected the extinction scenario almost entirely. Cohere's Aidan Gomez occupies a more intermediate position: he regards misalignment, AI systems behaving in ways their creators did not intend, as a real concern, but thinks cyberattacks are a much more immediate national-security threat.

The latest version of the doomsday argument is increasingly about »recursive self-improvement«. If AI becomes good enough at AI research to help build better AI, and that AI becomes better at building the next generation, the development process could begin to accelerate beyond our ability to supervise it. Journalist Ezra Klein's New York Times essay on 20 September put the argument starkly: »We're Not Losing Control of A.I. We're Giving It Away.«[7] The possibility is serious enough to deserve attention, but it's also a headline-grabbing news story – there's a mechanism, a threshold, and a point beyond which humans might no longer be in charge: »Will AI Wipe Out Humanity?«

But why is this the question we keep asking?

An alternative way to frame the debate would be to use four probabilities. There's the familiar extinction thread, what the industry calls »p(doom)« – the probability that advanced AI causes catastrophic or existential destruction. Then there's perhaps p(amp): the probability that AI substantially amplifies things humans already know how to do, from hacking and fraud to propaganda and biological warfare. There might be p(subcog): the probability that we outsource enough cognitive work to machines to change what we practise, remember, and learn. And finally p(mood): the probability that systems through which millions of people increasingly ask questions and interpret events acquire influence over the emotional and intellectual atmosphere in which those people experience the world.

These are genuinely different risks, and the distinction matters because the first is spectacular while the other three are gradual. We already know that cognitive outsourcing can have both costs and benefits. Writing changed human memory; calculators changed how we do arithmetic; search changed what we recall. In a similar way, AI may now allow people to formulate questions they couldn't previously have asked, while moving the defining skill of writing from producing the text towards selecting, editing, questioning, and judging it. That could very well be a transformation rather than a decline.

The more peculiar possibility is p(mood). A system doesn't have to lie or persuade us politically to influence how we experience the world. If millions of people routinely ask machines whether something is frightening, ridiculous, important, hopeful, or worth worrying about, those machines become part of the machinery of interpretation. Information has a temperature, and the systems that increasingly mediate it may influence that temperature simply by becoming ubiquitous.

This is where the obsession with p(doom) is perhaps misleading, even if the underlying risk is real, much as the risk of Earth being struck by a large asteroid is real, but vanishingly small.

Nature's own science desk explored this exact tension two days before CNN's survey, quoting one risk researcher who said such predictions rely on »so many untestable claims that you just end up with a conversation that is really more like faith than something scientific or empirical«, and another who worries far more about »AI's low reliability and accuracy rates in critical environments with life-or-death consequences« than about the prospect of extinction.[8]

Extinction is an unusually dramatic endpoint against which every other consequence looks small. It also encourages us to imagine a moment when everything changes. Most technological transformations don't work like that.

Those of us old enough to remember the »Y2K bug«, the supposed computer catastrophe that would wreak havoc around the world as the 20th century turned into the 21st, remember that it at least had a midnight to focus on. The date arrived, the lights stayed on, and the world continued. That's not the case with the AI-flavoured doomsday scenarios.

There will be no morning when we discover that we have become a different species. There will only be millions of small decisions about what we delegate, what we practise, what we remember, what we ask, and what we no longer bother to learn.

Humanity is of course becoming different again. We became different with writing, with printing, with photography, with the computer, with the Internet, and with the smartphone. Redefining what it means to be human isn't an exceptional event in human history. It's what humans do. Will AI change us? It already is. The question is what, exactly, we're becoming, and who gets to build the machinery through which the change happens.

When the dust settles, perhaps we'll wish we'd spent a little less time on p(doom), and a little more on p(mood).

---

REFERENCES

1 Hadas Gold & Clare Duffy (2026) »Not everyone thinks AI will kill us all«. CNN, 24 September 2026. https://www.cnn.com/2026/09/24/tech/not-everyone-thinks-ai-will-kill-us-all

2 Dario Amodei (2026) »We Must Pace the Frontier«. darioamodei.com, 12 September 2026. https://darioamodei.com/post/we-must-pace-the-frontier

3 Aoyon Ashraf (2026) »Anthropic CEO Calls for AI Race to Slow Down Citing Safety. Musk and OpenAI's Altman Agree«. CoinDesk, 12 September 2026. https://www.coindesk.com/tech/2026/09/12/anthropic-ceo-calls-for-ai-race-to-slow-down-musk-and-openai-s-altman-agrees

4 Mike Allen, Zachary Basu & Madison Mills (2026) »Exclusive: Google DeepMind's Demis Hassabis Calls for U.S.-Led Global AI Watchdog«. Axios, 14 July 2026. https://www.axios.com/2026/07/14/demis-hassabis-ai-regulation-google-deepmind

5 Chris Thomas (2026) »Zuckerberg Says AI Labs Have Enough Incentive to Build Safely«. Insurance Journal, 16 September 2026. https://www.insurancejournal.com/news/national/2026/09/16/885335.htm

6 Anna Heim (2026) »Yann LeCun's AMI Labs Raises $1.03B to Build World Models«. TechCrunch, 10 March 2026. https://techcrunch.com/2026/03/09/yann-lecuns-ami-labs-raises-1-03-billion-to-build-world-models/

7 Ezra Klein (2026) »We're Not Losing Control of A.I. We're Giving It Away.«. The New York Times, 20 September 2026. https://www.nytimes.com/video/opinion/100000011157772/were-not-losing-control-of-ai-were-giving-it-away.html

8 Elizabeth Gibney (2026) »Will AI really kill us all? The science behind the hype«. Nature, 22 September 2026. https://doi.org/10.1038/d41586-026-02941-3

26

SEPTEMBER 2026
Essay 026

# The Age of Projects

> *What happens when companies stop trying to predict what people will want, and let them help their R&D departments find out?*

Yesterday, Vogue Business reported the reinvention of Vaseline, a 155-year-old petroleum-jelly brand that people bought mostly from habit. The company found more than 3.5 million social posts showing people using Vaseline in ways it had never intended. Instead of advertising around them, it set about testing them – validating some, debunking others, tracking down the customers who'd invented the best ones and »co-authoring« new products with them. The gamble paid off: three straight years of double-digit growth, a place among Unilever's core 30 »Power Brands«, and 14 Lions won at the festival of creativity in Cannes.[1]

**»The product is no longer the answer. It's the question«**

Vogue considers this a heritage brand learning to speak the language of the Internet, but there's something else here. The company puts something into the world, watches what people do with it, and then changes the thing itself in response. Provide, observe, change, provide again.

Most of the industrial era wasn't about *provision*, but about *prediction*. Product development ran the other way: predict what people want, design it, manufacture it, sell it. Companies built enormous research operations – forecasters, focus groups, consultants – to know the future before committing to a product.

I sat in on some of those focus groups, and can confirm it can be confusing to listen to a bunch of Grade D consumers – »semi-skilled and unskilled manual workers«, in the researchers' unforgiving jargon – in the English town of Slough discuss what they want from a new product. Ask customers what they want and they'll tell you what they think they want; watch a hundred million people use something and you may discover what they didn't know to ask for. Four decades ago, futurist writer Alvin Toffler already had a name for the person doing both jobs at once: the »prosumer« – producer and consumer folded into one.[2]

Economist Eric von Hippel showed that some of the most important product improvements don't come from a company's own labs but from »lead users« who adapt a product to some challenge before anyone else faces that challenge.[3] Vaseline's customers ran uncontrolled experiments on petroleum jelly for years; now the company decided to treat the results as research.

Call the resulting object a »project« – a product that is also a project, thrown forward before its final form is settled.[4] A prototype is unfinished because its maker hasn't finished designing it. A project is unfinished on purpose: a gap left in the design for the world to fill in.

Software has lived in perpetual beta for years. Streaming and social platforms learn what to make and present from measuring what people watch or stop scrolling for. The more interesting version doesn't just observe preference: it lets people develop the actual thing. No R&D department, however brilliant, can reproduce the sheer experimental diversity of a hundred million people misusing, reimagining, adapting, and breaking a product every day.

IKEA has long given customers furniture that gets hacked and repurposed by an entire subculture of unofficial inventors long after it leaves the factory.[5] Marijn van der Poll's *Do Hit Chair* makes the same logic almost comically literal: a block of steel and a sledgehammer, finish the design yourself.[6] Architect Alejandro Aravena deliberately built only »half of a good house« – the expensive, structurally difficult half – and left the rest for residents to complete as their own means allowed.[7] Petroleum jelly, poured concrete, and code: unfinishedness turns out to be a contemporary design decision available at almost any scale.

And the same pattern has become the defining feature of AI. A model's makers know a great deal about how they built it, but considerably less about what millions of users will do with it. So the model ships, people find uses nobody anticipated, and the next version absorbs some of what was learned. The product is no longer the answer. It's the question.

Leaving a little room on purpose is different from leaving things unfinished by accident. A finished product tells you what its maker already knows. A project leaves room for what the maker doesn't know yet – which may be the more useful thing to be selling. That's the shift underway: not the ability to innovate faster, but the ability to build things that discover their own future through use. We often think about our economy as one of innovation. Maybe the age of projects will give us an economy of organised surprise.

---

REFERENCES

1 Rachael Akhidenor (2026) »How Vaseline Became Fluent in Internet Speak«. Vogue Business, 25 September 2026. https://www.vogue.com/article/how-vaseline-became-fluent-in-internet-speak

2 Alvin Toffler (1980) *The Third Wave*. New York: William Morrow and Company.

3 Eric von Hippel (1986) »Lead Users: A Source of Novel Product Concepts«. Management Science, Vol. 32, No. 7. https://doi.org/10.1287/mnsc.32.7.791

4 The word »project« comes from the Latin *proicere* — *pro-* (»forward«) plus *iacere* (»to throw«) — meaning, literally, »thrown forward«. A project was never a finished thing to begin with: it was a cast made toward an outcome before that outcome existed. »Project« borrows exactly that throw, and simply refuses to let it land only once.

5 Jules Yap (2006–) IKEA Hackers website. https://ikeahackers.net/

6 Marijn van der Poll (2000) *Do Hit Chair*. Droog. https://www.droog.com/projects/do-hit-chair-by-marijn-van-der-poll/

7 Alejandro Aravena (2014) »My Architectural Philosophy? Bring the Community into the Process«. TED2014, October 2014. https://www.ted.com/talks/alejandro_aravena_my_architectural_philosophy_bring_the_community_into_the_process

27

SEPTEMBER 2026
Essay 027

# The Life of the Gift

> *A diplomatic gift looks simple enough: a country wants to signal friendship, so it hands something over. But two of the best-known theories of the gift say that's never quite what's happening — and if they're right, they leave behind an odd question. How does a gesture built on calculation and disguised self-interest end up producing something that feels, eventually, completely sincere?*

Yesterday, following China's latest pair of pandas arriving in Atlanta and Donald Trump's gift of a bald eagle statue to Xi Jinping in return, the Guardian asked what diplomatic gifts mean and how countries choose one.[1] France gave America the Statue of Liberty. China gives pandas. Japan gave Washington cherry trees. Countries have exchanged puppies, tulips, and bells. There's even a diplomatic gift that was never really a gift at all: the Trojan horse.

**»Self-interest is what starts a gift moving. Sincerity is just time, added afterward, and time doesn't ask the beginning's permission«**

Sociologist and anthropologist Marcel Mauss's 1925 study *The Gift* argued that gifts are never really free: they create obligations to give, to receive, and to reciprocate.[2] The object isn't as important as the relationship forming around it. Sociologist Pierre Bourdieu added something sharper: the gap in time between giving and returning is what lets a transaction disguise itself as generosity.[3] Hand over a bottle of wine and get a scarf worth the same amount in return, and everyone recognises a trade. Wait months, then send back something of similar value, and the same exchange passes as friendship. But notice that both theories describe only the moment of exchange. Neither says anything about what happens to the object afterward, once nobody is watching the reciprocity anymore.

The pandas Xi Jinping gave Trump this week are still suspended inside that moment. Since the 1980s, China's »gifted« pandas have actually arrived on ten-year loans, with recipient zoos paying roughly a million dollars a year toward conservation work.[4] The language of friendship survives the fact that the gift is legally a lease, because calling it a gift does diplomatic work that calling it rent doesn't. The fiction holds only for as long as nobody bothers to look at what's actually being exchanged.

But a gift, unlike a lease, doesn't stay under the giver's control – it's more like a move in chess, which hands the turn to someone else. An intention of generosity can produce gratitude, resentment, or nothing at all; the giver's plan is only one ingredient in what the object becomes. The Statue of Liberty shows how far this drift goes. It was also a message meant for France itself: liberal politician Édouard Laboulaye and sculptor Frédéric Bartholdi hoped that celebrating American liberty might push Napoleon III into permitting more of it at home. It wasn't even a clean transfer between governments: France ran short of funds, and the publisher Joseph Pulitzer crowdfunded the pedestal through his newspaper from more than 120,000 ordinary readers, most giving less than a dollar.[5]

In 1976, Queen Elizabeth gave French president Valéry Giscard d'Estaing a black Labrador puppy. A moment earlier it had simply been a puppy; now it carried an intention, an expectation, a debt. It grew up into mostly a dog again – the president loved it, and spoke English to it all its life – though it remained, always, the dog that had once been given. The gift status disappeared. The gift history didn't.

In 1912, Japan gave Washington thousands of cherry trees. The trees kept growing regardless of what anyone intended: they became a landscape, then a season, then an image of the city itself. Their origin remains part of their history, but that history no longer stops there. Which is where the answer starts to come into focus: the calculation Mauss and Bourdieu describe only has to run for as long as the exchange itself takes. What happens next belongs to time and use, not to anyone's continued belief in the gesture's

generosity. A tree doesn't need spectators to keep blooming. Self-interest is what starts a gift moving. Sincerity is just time, added afterward, and time doesn't ask the beginning's permission.

The Trojan horse is what happens when a gift is built never to let that second stage arrive: its openness is fraudulent, the giver has already decided the receiver's next move, and the horse can never become just a horse.

A good enough gift keeps living long after the exchange that produced it is forgotten, until it belongs to neither the giver nor the receiver – only to itself. The cherry trees still bloom in Washington every spring, and almost nobody who stops to admire them is thinking about Japan. They are no longer a gift. They are trees that were once a gift, and have since become something else. An exchange ends when the account is settled. A gift ends when no one remembers there ever was one.

---

REFERENCES

1 Kim Willsher (2026) »France gave the Statue of Liberty. China sends pandas. What do diplomatic gifts mean — and how do you choose one?«. The Guardian, 26 September 2026. https://www.theguardian.com/world/2026/sep/26/france-statue-of-liberty-china-pandas-what-do-diplomatic-gifts-mean-and-how-do-you-choose-one

2 Marcel Mauss (1925) *The Gift: The Form and Reason for Exchange in Archaic Societies*, trans. W. D. Halls (1990). London: Routledge.

3 Pierre Bourdieu (1972) *Outline of a Theory of Practice*, trans. Richard Nice (1977). Cambridge: Cambridge University Press.

4 NPR (2025) »The Debut of New Pandas in D.C. Marks the Latest Chapter in China's 'Panda Diplomacy'«. NPR, 24 January 2025. https://www.npr.org/2025/01/24/nx-s1-5272416/panda-diplomacy-china-dc-zoo

5 National Park Service (2015) »Joseph Pulitzer«. Statue of Liberty National Monument, nps.gov. https://www.nps.gov/stli/learn/historyculture/joseph-pulitzer.htm

28

SEPTEMBER 2026
Essay 028

# The Shape of Understanding

> *The first exhibition at MoMA devoted entirely to information design asks how invisible systems become visible. The harder question is what disappears in the process.*

Yesterday, the Museum of Modern Art in New York opened the exhibition *Full Disclosure: The Edge of Information Design*. Its 30 digital and analogue works explore how designers turn data into forms that people can see and understand. In an interview, curator Paola Antonelli recently expressed the motivation sharply: »The objectivity of data, in particular, is increasingly under scrutiny, as public confidence in institutions and in the intelligence they produce is decreasing.«[1]

**»The real distinction isn't between diagrams to trust and diagrams to doubt, but between an exclusion we can see and one we can't – between not yet knowing enough, and having been reassured just enough not to ask«**

Can a diagram tell the truth? Maybe diagrams were never really in the truth business, only the legibility business – and the two have been quietly confused for a century. Or maybe that's too simplified: the people involved in both businesses cared, each in their own way, about getting something right.

In 1920s Vienna, polymath philosopher Otto Neurath and his collaborators developed ISOTYPE, a system of pictorial statistics designed to make complex information legible to people outside the target audience of conventional statistics. The heartfelt ambition was usability over spectacle: standardised symbols, repeated using strict rules, so that a visual reader could take in unfamiliar connections at a glance.[2]

Technical draughtsman Harry Beck drew electrical circuit diagrams for a living. His 1931 map of the London Underground treats the railway network as such a circuit diagram, making the same point as Neurath but in an entirely different way. Beck abandoned geography for connectivity – straightening lines, distorting distances – which worked precisely because his wasn't a correct depiction of London. We know Baker Street isn't really at right angles to King's Cross. We accept the distortion because we understand what the map chose to preserve: topology instead of geometry, how to get from A to B instead of where A and B actually are in the world.[3]

Mathematician Claude Shannon's information theory gives that choice a name. Communication means encoding information into a signal that can pass through a channel, and often something is lost in that encoding.[4] Representation establishes a relationship between what's retained and what's lost. Making something legible isn't just making the invisible visible – it's deciding what needs to be on display and what can remain invisible. Beck's map hides geography. ISOTYPE hides numerical detail. Every useful diagram builds a field of attention by excluding almost everything else.

Perhaps that's just the cost of any communication, no more troubling than speech leaving most of the world unsaid. Or perhaps something changes once what is left out stops being visible as an omission – once a diagram no longer reads as a selection, but as the thing itself.

Architecture has something to say here. In their book *Learning from Las Vegas*, architects Robert Venturi, Denise Scott Brown, and Steven Izenour distinguished the »duck« from the »decorated shed«: the duck's form carries its own message, while the decorated shed is an ordinary building wearing symbolic meaning applied to its surface.[5] The distinction transfers: an honest diagram is a duck, its structure doing the work its appearance promises. A decorated-shed diagram wears the signs of legibility without doing the actual work.

That's the sharper version of what's at stake. Information design was never going to fail by excluding things – it has to, that's the whole method. It fails when it hides the fact that it has excluded something. A duck's omissions are legible as omissions. A decorated shed's aren't. Good simplification hides information. Bad simplification hides the fact that it has simplified.

Think of a dashboard reduced to a single number – a credit score, an election-forecast needle. Nothing about its clean presentation admits that anything was left out. The cleaner the number looks, the easier to forget what was measured, excluded, and assumed, and how uncertainty was handled.

Statistician and data visualisation pioneer Edward Tufte's long attack on »chartjunk« tried to prevent exactly this, stripping away decoration that competed with the data itself.[6] His answer wasn't to distrust diagrams, but to make sure the visual economy earned its keep. The trouble is that the same competence serves both purposes. A building that lies about its structure eventually meets gravity, an engineer, or an inspector. A dashboard only has to meet us. A diagram can borrow that trust from its appearance without earning it.

We increasingly inhabit systems we can't examine directly – algorithms, financial models, filtered content. We only ever experience their representations. The real distinction isn't between diagrams to trust and diagrams to doubt, but between an exclusion we can see and one we can't – between not yet knowing enough, and having been reassured just enough not to ask.

Harry Beck's map has been wrong about London for nearly a hundred years. It never collapsed, because it never once pretended to be right.

---

REFERENCES

1 Rena Gross (2026) »The Museum of Modern Art's 'Full Disclosure' Finds Beauty in Data«. Galerie, 20 August 2026. https://galeriemagazine.com/the-museum-of-modern-arts-full-disclosure-finds-beauty-in-data/

2 University of Reading Art Collections, »Arntz, Gerd«. collections.reading.ac.uk. https://collections.reading.ac.uk/art-collections/collections/arntz-gerd/

3 London Transport Museum, »Transforming the Tube Map: Harry Beck's Iconic Design«. ltmuseum.co.uk. https://www.ltmuseum.co.uk/collections/stories/design/transforming-tube-map-harry-becks-iconic-design

4 Claude E. Shannon (1948) »A Mathematical Theory of Communication«. Bell System Technical Journal, Vol. 27. https://doi.org/10.1002/j.1538-7305.1948.tb01338.x

5 Robert Venturi, Denise Scott Brown and Steven Izenour (1972) *Learning from Las Vegas*. Cambridge, MA: MIT Press.

6 Edward R. Tufte (1983) *The Visual Display of Quantitative Information*. Cheshire, CT: Graphics Press.

29

SEPTEMBER 2026
Essay 029

# Form Follows Constraint

> *What's the difference between designing something to make it behave, and running an experiment to make something reveal how it behaves?*

Yesterday, Nature Physics published a paper describing a strange thing that happens to waves in a peculiarly shaped cavity. Researchers built an irregular enclosed space from a hyperbolic metamaterial, whose response to waves depends on their direction.

**»The designer wants the system to do something; the scientist wants the system to reveal something«**

That's a mouthful worth picking apart: a »hyperbola« is an open curve that runs off toward infinity in two directions rather than closing in on itself, and a »metamaterial« is one that gets its behaviour from engineered structure rather than from chemistry.

The scientists then sent waves through it – not light, and not quite the sound you'd hear with your ears: more like the vibrations of an earthquake, or the ultrasound a doctor scans you with, moving through a solid rather than through air.

Instead of scattering chaotically, the waves converged onto a particular, repeating path. The path has a »handedness« – it always loops one particular way, clockwise or counter-clockwise, rather than mirroring itself – it persists across a broad range of wavelengths, and survives imperfections in the cavity.[1]

Ordinarily, this kind of settling-onto-one-path behaviour only turns up in systems that either lose energy as they go, the way a swinging pendulum slows down and eventually stops, or that respond to themselves in complicated, self-amplifying ways, the kind of feedback that produces chaos. This system does neither. It loses no energy, and its response is always simple and proportional. Nothing amplifies, nothing feeds back on itself.

They've created something that behaves like an attractor without the usual ingredients. The cavity is irregular, the material is directionally biased, and together those constraints organise the waves into a path that nobody designed. Nothing tells an individual wave where to go. What the researchers have done is construct a world in which certain things happen.

Which is where the experiment becomes more interesting than the result.

We tend to imagine science as observation: a world already there, examined closely enough to reveal how it works. But experiments rarely work like that. Instead, the scientist builds an apparatus, selects materials, and controls initial conditions. The observable behaviour emerges from a deliberately constructed set of constraints.

The physicists didn't find their wave attractor lying around in nature. They made the conditions under which it could appear, and then watched what followed. This is remarkably close to what a designer does.

Architect Frei Otto understood this well. His soap films and hanging chains weren't representations of forms already designed. They were systems that generated forms: form-finding instead of form-giving.[2] Establish the conditions, introduce the material, let gravity and surface tension do their work, and a shape appears. The designer doesn't design the form so much as construct the conditions it emerges from.

Form follows constraint. But constraint follows form, too. Once a form exists, it becomes a constraint on what can happen next. A roof directs water. A structural frame limits where loads can go. Every designed form is also a small behavioural machine.

And suddenly, design and science seem to touch. Design is the deliberate construction of constraints from which desired behaviour emerges. Science is the deliberate construction of constraints from which observable behaviour emerges. The difference is tiny, but it matters: the designer wants the system to *do* something; the scientist wants the system to *reveal* something.

Molecular biologist and historian of science Hans-Jörg Rheinberger's idea of the »experimental system« captures the scientific side of this particularly well. An experiment isn't just confirmation of something already known. It's an arrangement that generates phenomena, making something observable that wasn't before.[3]

The wave attractor is an unusually clean example. The researchers constructed an environment in which a particular behaviour could emerge, and then discovered what that environment made possible. We tend to think that control means issuing instructions, telling something what to do. But there is another kind of control: instead of specifying the behaviour, construct the world in which the behaviour emerges.

The wave doesn't follow an instruction. It's just never given anywhere else to go.

---

REFERENCES

1 Simon Yves, Enrico M. Renzi, Sander A. Mann & Andrea Alù (2026) »Hyperbolic Wave Attractors«. Nature Physics, 28 September 2026. https://www.nature.com/articles/s41567-026-03453-7

2 Frei Otto & Bodo Rasch (1995) *Finding Form: Towards an Architecture of the Minimal*. Fellbach: Edition Axel Menges.

3 Hans-Jörg Rheinberger (1997) *Toward a History of Epistemic Things: Synthesizing Proteins in the Test Tube*. Stanford, CA: Stanford University Press.

30

SEPTEMBER 2026
Essay 030

# What Makes a Mind Continuous?

> *An anaesthesia study spanning six species suggests that rather than how much activity a brain contains, consciousness might depend on how strongly its present remains connected to its past, its parts, and what comes next.*

Yesterday, Nature Neuroscience published a study comparing brain activity during anaesthesia across six species: human, macaque, marmoset, mouse, zebrafish, and nematode. The last one is a microscopic ancient worm with just 302 neurons. The species could hardly be more different. Yet the researchers found a remarkably consistent change in the animals' brains under anaesthesia.[1]

**»Perhaps a mind isn't something that simply persists through change. Perhaps it's something that continuously reconstructs its own persistence«**

The brain doesn't simply switch off. What changes is the relationship between what's happening now and what happened immediately before, and between what's going on in one part of the brain and what's happening elsewhere. Local neural activity becomes faster and less dependent on its recent history. At the same time, distant regions become less synchronised. The paper describes the resulting state as »spatiotemporal isolation« – the activity continues, but it becomes less connected across both time and space. This is an unusual kind of disconnection. Not silence, but activity that has become less able to remain connected to itself.

The researchers found this by applying »massive feature extraction«, analysing what are called neural time series: measurements of brain activity taken repeatedly over time, producing a record of how a region or individual neuron changes from one moment to the next. They extracted more than 6,000 different mathematical descriptions of these records, measuring things such as repetition, predictability, and how strongly the present resembles the recent past. Of these, 485 changed consistently under anaesthesia across the six species and different anaesthetic drugs.

That result raises a deceptively simple question: what makes a mind continuous?

We tend to imagine continuity as a thread running from one moment to the next. But it can't just be a matter of recency. A traumatic experience from fifty years ago can matter more to us now than something that happened five seconds ago. A memory can acquire connections to hundreds of other memories, emotions, and expectations, while an event from yesterday might disappear almost completely. The past isn't a timeline stored behind us. It is more like a vast, uneven field of influences, in which different events retain different potentials to affect the present.

Philosopher Henri Bergson made a similar argument about what he called *durée* (duration). »There is no consciousness without memory«, he wrote, because every present state contains something of the states that preceded it.[2] He didn't mean that we consciously remember everything, but that a living present is different from an isolated instant because the past continues to participate in it.

The anaesthesia study gives this intuition a measurable counterpart. One of the strongest signatures the researchers found was reduced autocorrelation: it became harder to predict the present activity of a region using the activity that immediately preceded it. What had just happened mattered less to what happened next.

That's not a proof that consciousness is continuity, and the study doesn't settle any particular theory of consciousness. Neuroscientist Giulio Tononi's controversial Integrated Information Theory, for example, asks a different but related question: what makes many

physical processes constitute one experience rather than a collection of independent processes?[3] The distinction it makes is useful here; continuity through time isn't enough. The parts of a mind also have to belong together.

Anaesthesia reduces the brain's activity, but the most interesting result is that it changes the relationships within that activity. The brain remains busy, but its present becomes less dependent on its recent past. And its different regions become less dynamically involved with one another.

Novelist William Faulkner put the persistence of the past more memorably than any neuroscientist could: »The past is never dead. It's not even past.«[4] That's a line that concerns guilt, responsibility, and the impossibility of simply declaring an earlier self finished. But it also suggests a useful way of thinking about consciousness. What keeps the past alive isn't full presence, but enough force to change what happens now. Perhaps a mind isn't something that simply persists through change. Perhaps it's something that continuously reconstructs its own persistence.

You've been unconscious before. You didn't disappear. The past is never dead. But it's also not equally alive. Maybe consciousness is the shifting interval in which what has happened remains capable of changing what is happening, and what is happening remains capable of changing what comes next.

---

REFERENCES

1 Andrea I. Luppi et al. (2026) »Comprehensive profiling of brain dynamics during anesthesia across phylogeny«. Nature Neuroscience, 29 September 2026. https://www.nature.com/articles/s41593-026-02460-4

2 Henri Bergson (1903) *An Introduction to Metaphysics*, trans. T. E. Hulme. New York: G. P. Putnam's Sons, 1912.

3 Giulio Tononi, Melanie Boly, Marcello Massimini & Christof Koch (2016) »Integrated information theory: from consciousness to its physical substrate«. Nature Reviews Neuroscience. https://www.nature.com/articles/nrn.2016.44

4 William Faulkner (1951) *Requiem for a Nun*, Act I, Scene III.

# THAT WAS ALL OF YESTERDAY'S NEWS

Thanks for reading.

The next issue of Yesterday's Month follows at the start of next month.

For the daily version – one essay per day – sign up for the newsletter on the website.

YESTERDAYSNEWS.ORG