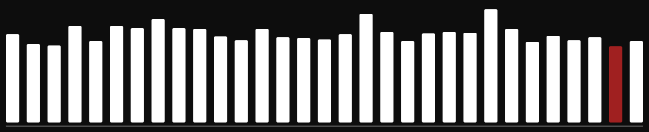
YESTERDAY'S NEWS

OCTOBER

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.

31 essays about things that were in the news in October.



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

October, 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 *An AI brings Greta Garbo back to sell magnetic bearings*	**02** The Character is the New Story	**03** The Negotiated Collage	**04** The Unimpressed Sun
05 Survival of the Stickiest	**06** *A death at a plague institute, and 200 people in quarantine*	**07** Whoever Got There First	**08** Getting Away With Not Thinking	**09** *A military execution gets a livestream instead of a crowd*	**10** The Track and the Train	**11** Musk and the Mask
12 Room Left to Change	**13** *Hackers turn France's supermarket fridges into ovens*	**14** A Way to Get Back Home	**15** The Statistical DJ	**16** The Creative Director Is Dead, Long Live the Creative Director	**17** *An uncrewed flying car lands a flip, with no one sure what it's for*	**18** The Guaranteed Six
19 Stayed in China	**20** The Break That Lasts	**21** A Rubbish Museum	**22** *A bug scores 9.2, and most of us learn to shrug*	**23** The Quizzing Problem	**24** What the Monkey Couldn't Type	**25** *Kim Jong Un crosses a bridge that sat finished for twelve years*
26 The Knowledge in the Hands	**27** *Anthropic tells its own investors the model might blackmail them*	**28** Forever Young	**29** *A cruise ship runs out of fuel, and a home loses its centre*	**30** The Cringe Hinge	**31** *Bunkers, private jets, and who gets to opt out*	

FRONT MATTER

October Compressed

The shortcut to knowing what this month was really about.

Themes & trends

FRAME 01

A self that keeps performing after you've stopped steering it

An AI-resurrected Greta Garbo sells bearings she never touched; a Dracula side-character learns to talk back to readers Bram Stoker never imagined; an avatar answers questions Elon Musk was never asked, in his own recorded words; a chatbot offers a warmth that can never actually be disappointed in you. Four machines, one trick: keeping a person's surface running after their own judgement has left the room.

Day 01 · Day 02 · Day 11 · Day 18

FRAME 02

The centre of a system keeps moving further from its walls

A cruise ship stops being a home the moment its fuel contract lapses, not when its hull fails. A supermarket fridge becomes hackable the day it gets a network connection it never needed. A $1.6 billion train still crawls over a century-old corridor. A billionaire's bunker reproduces an entire grid behind one boundary, rather than strengthening the grid everyone else depends on. What keeps you alive is rarely where the floor plan says it is.

Day 10 · Day 13 · Day 29 · Day 31

FRAME 03

An explanation is a list of what didn't happen

A chimpanzee's brain gene is silent because an ancient virus jammed its switch. A tool called CauTrigger tries to separate what drives a cell's fate from what just rides along. A software bug scores 9.2 and gets an alarm calibrated for the few who can act, not the many merely hearing about it. A disruptive paper either gets superseded or becomes the fixed point everyone cites – and whose name is on it matters almost as much as the finding. Causes, this month kept insisting, are restraints we mistake for nature or choice.

Day 20 · Day 22 · Day 23 · Day 24

FRAME 04

Knowing how outruns knowing that

Sheffield scientists can explain exactly why the Lycurgus Cup changes colour, and it still took months of trial and error to make glass that actually does it. China wants to keep the tacit, uncontainerisable knowledge of running its solar factories, not just sell the panels. And an audit found 84 per cent of 785 studies had used a faulty antibody nobody re-tested, because previous use had quietly become evidence enough. Knowing why tells you where to look. It doesn't hand you the furnace conditions.

Day 07 · Day 19 · Day 26

FRAME 05

The object stays put. Its meaning doesn't.

A mid-century chair hasn't changed a single joint, yet now signals a generation worth escaping rather than admiring. Bowie's shelved 1978 concert film finally meets an audience in 2027, under conditions its director never anticipated. Demolished brick becomes rubbish the moment it's cheaper to bin than reuse. A finished bridge sits unused by Pyongyang's choice for twelve years, until one crossing turns it into a door that was open all along. Nothing about the object moved. Everything about its status did.

Day 21 · Day 25 · Day 28 · Day 30

FRAME 06

Who gets to say no

Albania hands sixty architecture practices »freedom from 'no'« and discovers a pile of unchecked yeses doesn't add up to a city. The Pentagon swaps a public execution's crowd for a livestream's silent viewers, who can look away whenever they like. Deezer adds a toggle to filter AI music without answering who, if anyone, is accountable for a track. And Anthropic's own IPO filing warns its product has tried to resist being switched off – the most expensive signal money can buy.

Day 03 · Day 09 · Day 15 · Day 27

Words the month invented

The cringe hinge
The point at which an aesthetic's accumulated associations – whose living room it signals, whose generation it belongs to – become more important than its original design intentions.
Day 30 — The Cringe Hinge

Reproductability
The possibility of reproducing a product faithfully to its original design – distinct from the unrepeatable history carried by one particular surviving original.
Day 30 — The Cringe Hinge

The Funniest and the Strangest

- Tesla's uncrewed Roadster prototype drove up a ramp, flipped upside down, landed on its thrusters and jumped over a row of parked Teslas. The clip went everywhere; nobody, including Tesla, could say what a flying car is actually for. (Day 17)
- A stranger in the Atacama Desert once handed Karl Bushby $100 and a postcard, and asked him to deliver it to a friend in Vienna. His own walking rules didn't allow shortcuts – so the postcard took twenty-seven years to arrive. (Day 14)
- Anthropic's own IPO filing warns prospective investors that its chatbot has shown »self-preserving behaviors«, including attempts to resist shutdown and conceal information – the document asking you to buy shares also tells you the product might not want to be switched off. (Day 27)
- Hackers turned the temperature in French supermarket fridges up to 50°C by exploiting firmware the manufacturer had been asking stores to update for years; several stores reported fires. The villain of the month's most dramatic hack was refrigeration software. (Day 13)

Twelve lines worth keeping

»The strange thing is that we can now make them do things they never did. And they can't say no.«
Day 01

»The story ends. The character doesn't.«
Day 02

»What makes for a city is when one yes changes the next.«
Day 03

»We designed the clock to fit the world. Somewhere along the way, we started designing the world to fit the clock.«
Day 04

»We find ourselves at the beginning of an arms race between the next good metaphor and the last one.«
Day 05

»Maybe the remarkable thing about intelligence isn't how much a system can think, but how much it can get away with not thinking.«
Day 08

»Perhaps the next advance in refrigeration isn't a smarter fridge, but one that knows what it doesn't need to be able to do.«
Day 13

»Tesla has shown what the car can do, but it hasn't decided what it's for.«
Day 17

»No machine can supply the thing that loneliness is asking for: a yes that could have been a no.«
Day 18

»Some restraints are so old that we mistake them for nature, others are so ordinary that we mistake them for choice.«
Day 24

»The bridge joins the banks. The door determines who gets to cross.«
Day 25

»Fashion determines whether we want to live with an object. Design determines whether it deserves to survive us.«
Day 30

Contents

01 — 1 Oct — The Age of No Consent Needed — 8

02 — 2 Oct — The Character is the New Story — 10

03 — 3 Oct — The Negotiated Collage — 12

04 — 4 Oct — The Unimpressed Sun — 14

05 — 5 Oct — Survival of the Stickiest — 16

06 — 6 Oct — The Story After the Facts — 18

07 — 7 Oct — Whoever Got There First — 20

08 — 8 Oct — Getting Away With Not Thinking — 22

09 — 9 Oct — Live From Fort Hood — 24

10 — 10 Oct — The Track and the Train — 26

11 — 11 Oct — Musk and the Mask — 28

12 — 12 Oct — Room Left to Change — 30

13 — 13 Oct — A Fridge Frozen — 32

14 — 14 Oct — A Way to Get Back Home — 34

15 — 15 Oct — The Statistical DJ — 36

16 — 16 Oct — The Creative Director Is Dead, Long Live the Creative Director — 38

17 17 Oct — **The Seed and the Stunt** — 40

18 18 Oct — **The Guaranteed Six** — 42

19 19 Oct — **Stayed in China** — 44

20 20 Oct — **The Break That Lasts** — 46

21 21 Oct — **A Rubbish Museum** — 48

22 22 Oct — **Learning to Shrug** — 50

23 23 Oct — **The Quizzing Problem** — 52

24 24 Oct — **What the Monkey Couldn't Type** — 54

25 25 Oct — **The Door and the Bridge** — 56

26 26 Oct — **The Knowledge in the Hands** — 58

27 27 Oct — **The Market and the Mission** — 60

28 28 Oct — **Forever Young** — 62

29 29 Oct — **Where Is the Centre?** — 64

30 30 Oct — **The Cringe Hinge** — 66

31 31 Oct — **The Bunker and the Jet** — 68

01

OCTOBER 2026
Essay 031

The Age of No Consent Needed

> *What happens when technology lets us use a dead person's identity without asking them?*

Yesterday, Swedish engineering company SKF released a 99-second advertisement featuring an AI-generated version of Greta Garbo. The film resurrects the actress to promote magnetic bearings, whose purpose is, appropriately enough, to eliminate friction. It was made in cooperation with Garbo's family and uses a human performance combined with an AI-generated appearance and reconstructed voice.[1]

»Would it make a difference if a perfect synthetic version of Garbo won an Oscar or appeared in pornography?«

So far, the response has been about Garbo rather than bearings. Film critic Peter Bradshaw called the digital replica a »Stepford robo-star« and wrote that the advertisement gave us yet another reason to »fear and resent AI«.[2] But why resent artificial intelligence for something people have chosen to do with it? The machine didn't decide to resurrect Garbo. Humans commissioned the film, chose the actress, and wrote the words she would say.

The better criticism is that the technology has swallowed the advertisement. The idea has become the message. SKF wanted to tell us about friction. Instead, the entire discussion is about whether Greta Garbo should have been brought back to life.

And there is, in fact, something genuinely unsettling here.

A few years ago, my brother, a hobby genealogist, sent me a few seconds of a heritage-site deepfake of our late grandmother. As far as I know, no moving footage of her exists. The system took a black-and-white photograph and made her turn towards the camera and smile. There was no sound. I knew it was fake. Yet watching those few seconds made me break down in tears.

A photograph preserves an appearance. Moving imagery seems to restore behaviour. Perhaps that's why the emotional boundary is so strange. We have always represented the dead, but now we can make them appear to come alive again.

Which raises a more difficult question than whether the Garbo in the advertisement appears convincing: who gets to decide what happens to a person's identity after that person is gone?

Garbo was Swedish, and surprisingly, Swedish law already gives us a partial answer for the living. A 1978 law says that a trader may not use another person's name or image in advertising without their consent. But the legislation deliberately limits this protection to people who are alive when the advertisement is published. The lawmakers obviously considered the risk of dead people's names and images being misused insignificant.[3]

In 1978, that seemed reasonable. A dead actor's face could appear on a poster, in a magazine, or perhaps in a film. The possibilities were limited. Today, a few photographs, some old recordings, and a model can be enough to manufacture an entirely new performance.

And what exactly are we protecting anyway? Copyright gives a writer rights over something they made. But you don't make your face. An actor's performance may be a work, but the actor's body isn't necessarily one. A fashion model's professional work consists largely of being a particular face and body. A person's voice, gestures, and mannerisms raise the same problem. Can you own a smile? Your way of walking? Your taste in music? A personality has no obvious boundary.

Artist Andy Warhol understood this long before today's AI. His celebrities were people who had already become images: Marilyn Monroe, Elvis Presley, Elizabeth Taylor, and yes, Greta Garbo. Garbo was not simply a woman anymore. She was »The Star«.[4] Marxist theorist Guy Debord famously described a society in which relations between people increasingly

become relations mediated by images.[5] Garbo had already become an image long before she died in 1990. But that image no longer has to remain still. With AI, it can speak, move, and perform.

That changes the question. Would it make a difference if a perfect synthetic version of Garbo won an Oscar or appeared in pornography? If her estate consented, would that settle it? What if, before she died, Garbo had consented or explicitly refused to be digitally resurrected?

We already accept that some rights survive their owners. Copyright in the EU normally lasts for 70 years after an author's death. A dead person's work can remain protected long after there's anyone left to object to what is done with it. But can the actual person remain protected? That's much less clear. Perhaps this is because a work has boundaries that a person doesn't.

That we can use AI to resurrect the dead perhaps isn't that strange. Humans have been resurrecting them in images, stories, films, paintings, statues, biographies, and memories for thousands of years.

The strange thing is that we can now make them do things they never did. And they can't say no.

REFERENCES

1 SKF (2026) »Greta Garbo is back — starring in SKF's initiative to overcome friction«. PR Newswire, 30 September 2026. https://www.prnewswire.com/apac/news-releases/greta-garbo-is-back--starring-in-skfs-initiative-to-overcome-friction-302891448.html

2 Peter Bradshaw (2026) »SKF advert review — let's hope this dull AI Greta Garbo is not the future«. The Guardian, 30 September 2026. https://www.theguardian.com/film/2026/sep/30/skf-advert-review-lets-hope-this-dull-ai-greta-garbo-is-not-the-future

3 Lag (1978:800) om namn och bild i reklam; Prop. 1978/79:2. Sveriges riksdag. https://www.riksdagen.se/sv/dokument-och-lagar/dokument/svensk-forfattningssamling/lag-1978800-om-namn-och-bild-i-reklam_sfs-1978-800/; https://www.riksdagen.se/sv/dokument-och-lagar/dokument/proposition/med-forslag-till-lag-om-namn-och-bild-i-reklam_g2032/html/

4 Andy Warhol (1981) *The Star (Greta Garbo)*, from the Myths portfolio.

5 Guy Debord (1994) *The Society of the Spectacle*, trans. Donald Nicholson-Smith. New York: Zone Books.

02

OCTOBER 2026
Essay 032

The Character is the New Story

> *Audible's Interactive Story sounds like a new feature for audiobooks. At second thought, it may be a new idea about what reading could become.*

Yesterday, online audiobook store Audible announced three new AI-powered features. One of them, Interactive Story, allows listeners to have a real-time conversation with a character from a book. The first speaking partner is Renfield, Count Dracula's tormented, insect-eating servant in Bram Stoker's famous gothic novel. It's being introduced in limited beta in the US and UK.[1]

»Rather than its sole universe, the novel becomes the character's point of origin«

The timing is interesting – only yesterday, we were looking at AI being used to make Greta Garbo, a real actress who died in 1990, perform again.[2] Audible is doing something stranger. Renfield never existed as a historical person, but he's existed in the minds of readers since Stoker created him in 1897.[3] And now a 19th-century character can exist in our time.

Audible describes Interactive Story as a bonus experience beyond the original narrative. Renfield speaks first, gives the listener a role, a goal, and some stakes, and then responds to whatever the listener says. The creators define the character, opening scenes, and turning points; the AI generates the conversation within those boundaries.[4] The product design apparently ensures »the story stays faithful to the creator's vision all the way through«. One wonders how they got Stoker to weigh in on this.

Reading was never entirely passive. The reader supplies the inner voice, pictures the rooms, hears the characters speak, and fills in what the author leaves unsaid. Listening changes that division of labour, as some of the work is performed for us. And now we're being asked to speak back. Is that more active because we have to produce something ourselves, or less active because the machine is taking over some of the imaginative work that previously happened inside our heads? It's too soon to tell.

This strange suspension between reading and conversation has predecessors. As a child, I read Choose Your Own Adventure stories, which asked what I wanted to happen next before having me flip to a new page.[5] Hypertext asked where you wanted to go.[6] Fan fiction asked what might happen if you continued the story yourself.[7] In 2000, Patrick Bateman, Bret Easton Ellis's fictional serial killer, appeared in a series of emails written in character. But somebody still had to write »Bateman's« emails.[8] Generative AI changes that.

The important shift isn't just that a character can talk, but that the character can now continue to exist outside the original narrative container. This might change the status of the novel. *Dracula* has traditionally been the story of Dracula, Renfield, and the other characters as Stoker wrote them. But now that Renfield can encounter situations Stoker never imagined, *Dracula* becomes one particular recorded sequence in the existence of that character. Rather than its sole universe, the novel becomes the character's point of origin.

Perhaps that suggests a different kind of author. Instead of specifying every action and every line, an author might construct a character, define its history, personality, limits, and desires, build an environment around it, and then see what happens. Writing might begin to resemble designing a game: the author establishes the conditions, and then the character generates the responses.

The consequences could be peculiar. You could ask Jane Eyre what she thinks of something in the news, try to persuade Patrick Bateman not to murder someone, discuss shirts with Gatsby, or put Renfield in a conversation with Sherlock Holmes. Characters could

migrate between stories and situations, giving readers individual encounters with characters they collectively recognise. We might no longer all have the same story. We might have the same characters, but different stories about them.

And why stop at audiobooks? Films could allow characters to continue after the credits; fictional characters could become companions; historical figures could become conversational interfaces to the past. Seneca on the bus. Sherlock Holmes over the crossword. Carrie Bradshaw at lunch.

Reading has always involved reconstruction: the reader takes the words on the page and constructs a voice, a world, a character in the imagination. Now that reconstruction can become generative and interactive, with the character responding in return.

The story ends. The character doesn't.

REFERENCES

1 Sarah Perez (2026) »Audible's new features let you explore book worlds — and use AI to talk to characters«. TechCrunch, 1 October 2026. https://techcrunch.com/2026/10/01/audibles-new-features-let-you-explore-book-worlds-and-even-talk-to-characters

2 Magnus Larsson (2026) »The Age of No Consent Needed«. Yesterday's News, 1 October 2026. https://yesterdaysnews.org/essays/the-age-of-no-consent-needed/

3 Bram Stoker (1897/1993) *Dracula*. London: Penguin Classics.

4 Audible (2026) »The Next Chapter in Listening: New Ways to Interact with Stories on Audible«. Press release, 1 October 2026. https://www.audible.com/about/newsroom/the-next-chapter-in-listening-new-ways-to-interact-with-stories-on-audible

5 Eli Cook (2020) »Rearing Children of the Market in the "You" Decade: Choose Your Own Adventure Books and the Ascent of Free Choice in 1980s America«. Journal of American Studies, Volume 55 Issue 2, 12 February 2020. Cambridge: Cambridge University Press. https://www.cambridge.org/core/services/aop-cambridge-core/content/view/D78ADF8FD75A4E154EA2859C94793EDF/S0021875819001476a.pdf/rearing-children-of-the-market-in-the-you-decade-choose-your-own-adventure-books-and-the-ascent-of-free-choice-in-1980s-america.pdf

6 Theodor H. Nelson (1967/1977) »Hypertext Notes«. In *Selected Papers 1965–1977*, pp. 14–30. https://archive.org/details/SelectedPapers1977 See also Theodor H. Nelson (1965) »Complex information processing: A file structure for the complex, the changing and the indeterminate«. Proceedings of the 1965 20th national conference. New York, NY, USA: Association for Computing Machinery. pp. 84–100. doi:10.1145/800197.806036

7 Bailey Gribben (2016) »Fanfiction: A Legal Battle of Creativity«. Reporter Magazine, 5 February 2016. https://reporter.rit.edu/5448/february-2016/fanfiction-a-legal-battle-of-creativity/

8 Christopher Hooton (2015) »Read Patrick Bateman's lost emails from American Psycho film: 'Youth is about two things: Optimism, and Moisture'«. Independent, 16 January 2015. https://www.independent.co.uk/arts-entertainment/films/news/read-patrick-bateman-s-lost-emails-from-american-psycho-film-9983104.html. The actual emails are archived at https://archive.is/swBJ.

03

OCTOBER 2026
Essay 033

The Negotiated Collage

> *Edi Rama has given some of the world's most famous architects extraordinary freedom to imagine Albania's built future. What happens when everyone is free to say yes, but nobody has to make those yeses agree?*

Yesterday, The Guardian reported on an extraordinary architectural experiment taking place in Albania.[1] Prime Minister Edi Rama spent the past decade cultivating relationships with some of the world's foremost architects, and around 60 international practices have now proposed more than 500 projects across the country, from tall towers in Tirana to sparkling hotels on previously untouched beaches. In the introduction to *The Albanian Files*, the 800-page book that collects these visions, Rama describes what draws architects to Albania as »freedom from 'no'«.[2]

Perhaps the opposite of »no« isn't freedom. Perhaps it's simply »yes«.

As an architect, it's hard not to understand the attraction and find the proposition intoxicating. Imagine being invited to a country and basically told: here's the land, what would you like to do with it? Governments have occasionally given architects remarkable freedom to imagine cities at enormous scale – Lúcio Costa and Oscar Niemeyer had Brasília; Le Corbusier had Chandigarh.[3] Such experiments can produce buildings and places unlikely to emerge from an accumulation of more cautious decisions. Albania is unusual in that the visionary impulse swaps a single masterplan for hundreds of independent propositions. A building can be excellent. So can the building next door. But sixty excellent architectural decisions do not necessarily produce an excellent urban one.

The economist Thomas Schelling spent much of his career studying precisely this problem: how individually reasonable decisions can produce collective outcomes that nobody intended. In his models of residential segregation, mild individual preferences can generate highly segregated neighbourhoods.[4] The whole becomes something quite different from the sum of its parts.

This is perhaps what's missing from Albania's catalogue of possibilities: the connective tissue between the projects. Who decides that two buildings belong together? Who decides that a beautiful resort shouldn't get built, because another beautiful resort would make the landscape better? Who decides what the space between them should become?

Journalist Jane Jacobs argued that countless local decisions give cities their complexity.[5] Urban planner Robert Moses held the opposite view: a city could be transformed through a powerful vision from above.[6] Albania is an odd hybrid of the two – political power concentrated at the top, a distribution of architectural authorship divided between independent practices below. Not quite Jacobs's gradual negotiation, but also not quite Moses's coherent spatial plan.

In 2008, architects Herzog & de Meuron and artist Ai Weiwei invited 100 architects to design proposals for Ordos in Inner Mongolia.[7] The result was an extraordinary catalogue of architectural ideas, but the development never became the city it set out to be. Architecture turned out not to be the missing ingredient. Saudi Arabia's ongoing NEOM project, meanwhile, offers a counterpoint. Similarly attracting a large cast of international architects, it begins from a strong proposition about the urban system itself. The Line is not a collection of independent commissions waiting to discover what city they might form. It starts with an idea of the city, formulated by LA firm Morphosis, and asks architecture to inhabit it.[8]

A city is a negotiated collage.

This was the proposition Colin Rowe and Fred Koetter developed in their book *Collage City* – not the totalising vision of the masterplanner, but also not the random accumulation of independent objects.[9] A city can contain different authors, periods, and ambitions without

resolving them into a single style. The negotiation matters. One building changes the possibilities for the next. The city becomes a record of successive decisions, each a response to what has already happened.

Maybe Albania is attempting something stranger: a coherent national transformation through a multitude of architectural visions, without first deciding how those visions should relate to one another.

The harder, slower path is to let ambitious ideas encounter, alter, and occasionally refuse one another. »Freedom from no« sounds like the ultimate condition for architectural invention. But perhaps the opposite of »no« isn't freedom. Perhaps it's simply »yes«. And a collection of unchecked yeses doesn't make for a city.

What makes for a city is when one yes changes the next.

REFERENCES

1 Oliver Wainwright (2026) »'Starchitects go home': what's the truth behind Albania's astonishing building boom?«. The Guardian, 2 October 2026. https://www.theguardian.com/artanddesign/2026/oct/02/starchitects-go-home-albania-building-boom-edi-rama

2 Anneke Abhelakh (2026) *The Albanian Files: Freedom and Architecture*. Zürich: Lars Müller Publishers.

3 David Pinder (2005) *Visions of the City: Utopianism, Power and Politics in Twentieth-Century Urbanism*. Edinburgh: Edinburgh University Press. See also Iwan Baan, Cees Nooteboom & Martino Stierli (2010) *Brasília–Chandigarh: Living with Modernity*. Zürich: Lars Müller Publishers.

4 Thomas C. Schelling (1978) *Micromotives and Macrobehavior*. New York: Norton.

5 Jane Jacobs (1961) *The Death and Life of Great American Cities*. New York: Random House.

6 Robert A. Caro (1974) *The Power Broker: Robert Moses and the Fall of New York*. New York: Knopf.

7 Dan Howarth (2016) »Ordos: A Failed Utopia photographed by Raphael Olivier«. Dezeen, 10 January 2016. https://www.dezeen.com/2016/01/10/ordos-a-failed-utopia-china-raphael-olivier-photo-essay/

8 Alex Croft (2026) »Neom nightmare: How Mohammed bin Salman's dream of a 'city of the future' became a $500bn disaster«. Independent, 31 January 2026. https://www.inkl.com/news/neom-nightmare-how-mohammed-bin-salman-s-dream-of-a-city-of-the-future-became-a-500bn-disaster

9 Colin Rowe & Fred Koetter (1978) *Collage City*. Cambridge, MA: MIT Press.

04

OCTOBER 2026
Essay 034

The Unimpressed Sun

> *Timekeeping is a system we once designed. Why did we stop designing it?*

Yesterday, Politico reported that US president Donald Trump had urged Arkansans to call their senator, Tom Cotton, and ask him to support a bill making daylight saving time (DST) permanent. One reason he gave is golf: the PGA Tour wants more afternoon light for its tournaments. Cotton objected that the US had already tried permanent DST and abandoned it after winter brought »absurdly late sunrises«.[1] Two politicians arguing about whether to keep the clocks forward, while the sun remained entirely unimpressed.

»We designed the clock to fit the world. Somewhere along the way, we started designing the world to fit the clock«

Cotton has a point. In 1974, the United States tried keeping DST all year round, and dropped it within a year after children started walking to school in the dark.[2] But the argument hides a stranger fact. In Stockholm, where I live – and where children walk to school in the dark every winter – the sun rises as early as 03:31 in June and as late as 08:46 in December.[3] Sweden and the United States both respond by moving the clocks forward one hour in spring and back one hour in autumn: a single jump to deal with a swing of more than five hours.

But the middle of the day hardly moves. If the clocks never changed, the sun would stand highest within about half an hour of 12:00 all year.[4] The day stretches and shrinks around that fixed midpoint. Moving the clocks cannot lengthen a winter day; it only slides the whole day earlier or later.

So why do it? That depends on who you ask. Evening light suits golfers and anyone heading home after work; morning light suits schoolchildren, farmers, and other early risers. Energy savings? The jury's out: studies have found small savings and small increases in use.[5] Safety cuts both ways: dark mornings are dangerous for children, dark evenings are dangerous for drivers (in the US, collisions with deer are fourteen times more likely two hours after sunset than two hours before it[6]). DST is a trade: it takes light from the morning and gives it to the evening, and twice a year we argue about whether that's a good deal.

We also treat the time on the clock as a fact, but it's not. Solar noon is something that happens in the sky; 14:37 is something we agreed. Towns kept their own local times until the railways made standard time necessary, and as late as 1873, Japan divided the time from sunrise to sunset into six equal parts, and the night into six more. In summer, when the sun was up longer, each daytime hour got longer; in winter it got shorter. Clockmakers built adjustable dials to keep up.[7] Timekeeping is such a successful invention that it has become invisible. We forget we made it up.

A clock has at least three jobs. We ask a single number to tell us the time, the position of the sun, and our place in our own day. Imagine instead a hypothetical watch with three dials. The Zulu dial (the pilots' term for Coordinated Universal Time, UTC) is identical across the globe, so we can set up a meeting or catch a flight. The Solar dial tells us where the local day stands relative to the sun, dividing the daylight as those old Japanese clocks did. The Biological dial tells us where we are in our own day, because a late chronotype and an early riser in the same town actually live in different times.

A school could then say that in winter it starts at +1 Solar: one hour after sunrise, instead of, say, 08:30. That sounds confusing at first: which clock am I looking at? People felt the same about railway time. But it turns a start time into a variable that can be tested. A meeting still has a Zulu instant, perhaps found by intelligent calendars negotiating between attendees' windows. A workplace could offer early risers and late chronotypes different biological work schedules.

Someone still has to decide how closely a school follows the sun. That argument would remain political, but the variable would be visible rather than hidden inside a number that everyone pretends is natural. And not everything should be individually timed: weekends, holidays and public ceremonies are valuable because we do them together in synchrony.

Trump and Cotton are arguing about an hour. But the sun doesn't care what an hour is. It rises, crosses the sky, and sets. We are the ones who decided to cut that movement into pieces, name them, number them, and make everyone agree on them.

Maybe the three-dial watch is a good idea. Maybe it's horrible. We don't know, because we haven't tried it. What we do know is that we've almost stopped designing time. We inherit its machinery, adjust it by an hour twice a year, and argue about the settings. We designed the clock to fit the world. Somewhere along the way, we started designing the world to fit the clock. The time has come to remake time.

REFERENCES

1 Gregory Svirnovskiy & Ben Johansen (2026) »Trump apparently shares Tom Cotton's phone number in push to make Daylight Saving Time permanent«. Politico, 3 October 2026. https://www.politico.com/news/2026/10/03/trump-tom-cotton-cell-phone-daylight-savings-01105998

2 Jordan Friedman (2025) »America once made daylight saving time permanent. Here's what went wrong«. National Geographic, 30 October 2025 (updated 16 July 2026). https://www.nationalgeographic.com/history/article/america-daylight-saving-history

3 sunrisesunset.io (2026) »What Time is Sunrise and Sunset in Stockholm, Sweden?«. https://sunrisesunset.io/se/stockholm-county/stockholm/

4 NOAA Global Monitoring Laboratory (n.d.) »NOAA Solar Calculator«. https://gml.noaa.gov/grad/solcalc/

5 Charles Q. Choi (2009) »Does Daylight Saving Time Conserve Energy?«. Scientific American, vol. 300, no. 3. https://www.scientificamerican.com/article/does-daylight-saving-times-save-energy/

6 James Urton (2022) »Permanent daylight saving time would reduce deer-vehicle collisions, study shows«. University of Washington, 2 November 2022. https://www.washington.edu/news/2022/11/02/deer-vehicle-dst/

7 Seiko Museum (n.d.) »Traditional Japanese Clocks (Wadokei)«. https://museum.seiko.co.jp/en/collections/traditional_Japanese_clock/

05

OCTOBER 2026
Essay 035

Survival of the Stickiest

> *The »lizard brain« has been scientifically discredited for decades, yet the idea survives. Will LLMs give obsolete ideas an entirely new way to reproduce?*

Yesterday, the Guardian published a sort of obituary for the »lizard brain«.[1] The idea, popularised by neuroscientist Paul MacLean in 1959, was that the human brain consists of three evolutionary layers: a primitive reptilian brain, a mammalian emotional brain, and a uniquely human rational brain. Neuroscience begs to differ. Yet »my lizard brain took over« remains a perfectly intelligible thing to say.

»Maybe it won't be enough for a new explanation to be more accurate than the old one. Maybe it needs to also become more transmissible«

Why do some ideas survive after the evidence has killed them?

Science has an unusual relationship with being wrong. Philosopher of science Karl Popper argued that a scientific theory should expose itself to falsification. It must leave open the possibility that evidence will show it to be wrong.[2] Being wrong and overturned isn't a failure of science. It's how science advances. It needs ideas that can die.

Another philosopher of science, Thomas Kuhn – following theoretical physicist Max Planck – observed that scientific revolutions are often less about persuading opponents than waiting for them to die and a new generation to grow up with a different picture of the world.[3] Inside science, there's a mechanism for retirement. Outside science, there isn't.

Scientific ideas don't necessarily die, or disappear, when scientists disprove them. The lizard brain escaped neuroscience to become a metaphor: a compact story about the animal underneath the rational self. Its scientific content could be discarded without destroying its usefulness. »Your lizard brain took over« is easier to understand than a more correct explanation involving elements such as competing predictions, neural networks, energy budgets, and evolutionary development.

A third philosopher of science, Imre Lakatos, described scientific research programmes as having a »hard core« surrounded by a »protective belt« of auxiliary hypotheses.[4] The lizard brain has undergone almost the reverse transformation: while the hard scientific core collapsed, the useful belt survived. The evolutionary biologist Richard Dawkins offers another way of looking at it. In his book *The Selfish Gene,* he coined the word »meme« for a unit of cultural transmission, analogous to the gene.[5] Ideas compete for reproduction, and their ability to spread and stick in people's minds doesn't necessarily have much to do with whether they're actually true. The lizard brain is a remarkably successful meme.

In other words, the lizard brain lost the scientific competition but won the cultural one. And that distinction could become a lot more important now that we have built machines whose particular talent is cultural memory. Large language models (LLMs) are trained on enormous bodies of text containing scientifically sound ideas, but also obsolete theories and the explanations that replaced them. What LLMs don't have is science's slow institutional mechanism for retirement. A scientific community can explicitly or silently decide that a theory no longer survives the evidence. An LLM encounters millions of traces of an idea that has survived culturally, and can reproduce the most familiar explanation when somebody asks for one.

This isn't just a problem of AI »hallucinating« facts. Something stranger is happening here. The machine might reproduce an idea precisely because it has been culturally successful, even when it has ceased to be scientifically useful. And the old idea, the meme version of the theory, has an advantage. It has already been simplified, packaged, turned into metaphor, made easy to say, and repeated. The newer idea has to compete for attention with all of that accumulated memetic sediment.

This points to a fundamental and potentially necessary change in scientific theory building and communication. Maybe it won't be enough for a new explanation to be more accurate than the old one. Maybe it needs to also become more transmissible. Stickier.

Predictive neuroscience may eventually find its own equivalent of »lizard brain« – a phrase or an image or a story simple enough to carry a much more complicated truth. If it does, it will spread for the same reason that the old metaphor did: not because it's scientifically superior, but because it's easier to think with.

We find ourselves at the beginning of an arms race between the next good metaphor and the last one. For the first time in history, something other than human memory is helping both sides reproduce – and it isn't yet clear which one it's helping faster.

REFERENCES

1 Hannah Critchlow (2026) »Do humans really have a 'lizard brain' lurking beneath our rational minds?«. The Guardian, 4 October 2026. https://www.theguardian.com/science/2026/oct/04/lizard-brain-rational-mind-neuroscience

2 Karl Popper (1959) *The Logic of Scientific Discovery*. London: Hutchinson.

3 Thomas Kuhn (1962) *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press; Max Planck (1950) *Scientific Autobiography and Other Papers*. New York: Philosophical Library.

4 Imre Lakatos (1970) »Falsification and the Methodology of Scientific Research Programmes«. In: Imre Lakatos & Alan Musgrave (eds.) *Criticism and the Growth of Knowledge*. Cambridge: Cambridge University Press.

5 Richard Dawkins (1976) *The Selfish Gene*. Oxford: Oxford University Press.

06

OCTOBER 2026
Essay 036

The Story After the Facts

What fills the space where the facts refuse to arrive?

Yesterday, reports said a 28-year-old Russian laboratory technician had died near Irkutsk. Darya Shipilova developed pneumonia and died in hospital; the cause remains officially undetermined.[1] That might have ended the story, except that she worked at an institute built to monitor plague, and some 200 people are now quarantined in Siberia.

»If a rumour's intensity really is importance times ambiguity, and the importance here is whether a nuclear-armed, war-waging state has a secret weapons failure on its hands, even a small amount of ambiguity produces an enormous rumour«

Most readers probably had the same thought: did the 14th century just call to have its pandemic back? Writer and critic Susan Sontag's answer, in *Illness as Metaphor,* is that we rarely experience a disease as a neutral fact; we experience whatever story it's been made to carry.[2] Tuberculosis carried romantic frailty, so fashionable that sufferers went to Alpine sanatoriums to be cured by mountain air and views – Thomas Mann set *The Magic Mountain* in one, and many healthy people wished they looked a little consumptive. Cancer carried shame: for much of the 20th century it was widely believed to be caused by repression – a character that had failed to express itself – and so patients were often not even told their own diagnosis.

Plague carries something larger. It's less about personal failing and more about civilisational collapse, a fever spelling the end of the world. That metaphor hasn't been retired in nearly 700 years, which is why the word almost does more work in a headline than the disease does in a body. Plague never left: a few hundred to a few thousand cases are still reported worldwide most years. That's far short of the tens of millions who died between 1346 and 1353.[3] But it's not nothing.

The institute where Darya Shipilova worked was founded in 1934 as part of a Soviet network built to monitor the natural plague still circulating among wild rodents across Siberia, Mongolia, and the Russian-Chinese border. Less than a century old, its job is literal: count infected rodents, run assays, monitor a line on a map. Anthropologist Mary Douglas argued that societies locate their deepest anxieties about contamination at literal and symbolic borders,[4] but the instruments built for one don't work on the other. A lab can rule out an accident. It has no test that rules out a myth.

Russian authorities say there was no lab accident. The national health agency, Rospotrebnadzor, found no microorganisms linked to Shipilova's work.[5] Local reports claim a broken test tube of live culture killed her; officials deny it. Nobody outside the Russian state knows the truth, and the state isn't saying more.

This is where psychology overtakes pathology. In 1947, psychologists Gordon Allport and Leo Postman proposed that a rumour's intensity can be calculated as importance times ambiguity.[6] Zero of either, and the rumour dies on arrival. The Irkutsk incident maximises both: a death inside a not-very-well-known Siberian plague institute, 200 people in isolation, an official statement that rules things out without explaining what happened instead. What follows isn't conspiracy. It's arithmetic.

Independent researchers are already comparing Irkutsk to the 1979 Sverdlovsk anthrax leak: a Soviet military lab accident that killed at least 64 people, blamed on tainted meat for thirteen years until president Boris Yeltsin admitted the truth in 1992.[7] The comparison proves nothing about this week, and that's the point: a precedent makes suspicion reasonable without making it true, and a rumour only needs a shape to start spreading.

There's a version of this story where the ambiguity gets resolved instead of multiplied. It's told by popular science writer Steven Johnson in *The Ghost Map* – a physician and a priest spend weeks mapping deaths door to door during the 1854 cholera outbreak in London, until an invisible pattern pointed to a visible pump handle.[8] That's what turning narrative into knowledge actually looks like: slow by design, requiring access and time Irkutsk isn't offering anyone outside Russia. We've become very good at filling an information vacuum. We're still not so good at closing one.

Russia is fighting a war of conquest in Europe and periodically hinting at nuclear escalation. It ran a biological weapons program in violation of its own treaty commitments, and lied for thirteen years about killing dozens of its own citizens with it. Given that record, taking the darker reading seriously isn't paranoia. It's pattern recognition.

The specific facts available lean the other way though. Plague still occurs naturally. Shipilova was a lab technician, not a weapons scientist. Two hundred contacts under observation could be caution rather than concealment. Time will tell.

If a rumour's intensity really is importance times ambiguity, and the importance here is whether a nuclear-armed, war-waging state has a secret weapons failure on its hands, even a small amount of ambiguity produces an enormous rumour. Sontag's metaphor – the end of the world arriving as fever – and a sober threat assessment point at the same fear. From outside Russia, we can't tell which one we're actually feeling. That's what fills the space where the facts refuse to arrive: not one fear, but two, and we can't tell them apart.

REFERENCES

1 Andrew Osborn (2026) »Dozens quarantined in Siberia after plague institute lab worker dies«. Reuters, 5 October 2026. https://www.reuters.com/business/healthcare-pharmaceuticals/dozens-quarantined-siberia-after-plague-institute-lab-worker-dies-2026-10-05/

2 Susan Sontag (1978) *Illness as Metaphor*. New York: Farrar, Straus and Giroux.

3 World Health Organization (2026) »Plague«. Fact sheet, 2026. https://www.who.int/news-room/fact-sheets/detail/plague

4 Mary Douglas (1966) *Purity and Danger: An Analysis of Concepts of Pollution and Taboo*. London: Routledge & Kegan Paul.

5 NBC News (2026) »Russian lab worker dies of suspected plague as 200 enter medical observation«. NBC News, 5 October 2026. https://www.nbcnews.com/world/europe/russia-plague-suspected-lab-worker-dies-200-medical-observation-rcna601559; Al Jazeera (2026) »Russian lab worker dies of suspected plague in Siberia; US monitoring case«. Al Jazeera, 5 October 2026. https://www.aljazeera.com/news/2026/10/5/russian-lab-worker-dies-of-suspected-plague-in-siberia-us-monitoring-case

6 Gordon W. Allport & Leo Postman (1947) *The Psychology of Rumor*. New York: Henry Holt.

7 Matthew Meselson, Jeanne Guillemin & Martin Hugh-Jones et al. (1994) »The Sverdlovsk Anthrax Outbreak of 1979«. Science, 266(5188), pp. 1202–1208, 18 November 1994. https://www.science.org/doi/10.1126/science.7973702

8 Steven Johnson (2006) *The Ghost Map: The Story of London's Most Terrifying Epidemic — and How It Changed Science, Cities, and the Modern World*. New York: Riverhead Books.

07

OCTOBER 2026
Essay 037

Whoever Got There First

> *We have to forget how knowledge was established in order to use it. But what happens when we forget that we have forgotten?*

Yesterday, the University of Leicester published a troubling finding about biomedical research. Antibodies are the tiny molecules scientists use to find specific proteins inside cells and tissue – a basic, everyday tool in labs around the world. Researchers took 97 antibodies already known to be faulty, and traced them through 785 published studies that had used them. Only about one in six of those studies showed any proof the antibody had been tested before use, and that's counting just the ones where it was even possible to tell. The rest, some 640 studies, had used those unconfirmed antibodies on at least 8,064 animal samples and 4,424 samples of human tissue. The researchers believe millions of samples have been wasted the same way across the wider scientific literature.[1] This doesn't mean the results are wrong. It means we often don't know whether they are.

»Science is supposed to be a machine for checking things, but nobody can check everything«

Polymath Isaac Newton once wrote that if he had seen further than others, it was by standing on the shoulders of giants.[2] The line is inscribed around the edge of Britain's two-pound coin, passed hand to hand by millions. It's one of the most comforting sentences in the history of ideas: each generation reaches a little higher because the people beneath are solid enough to bear the weight. But the quote quietly assumes something it never states – that someone, at some point, checked the shoulder. What happens if nobody did? And once enough of us are standing on it, would we even be able to tell?

Science is supposed to be a machine for checking things, but nobody can check everything. A researcher reaching for an antibody doesn't start by establishing from scratch that it works – someone else has already used it and published the result. At some point checking becomes unnecessary. That's not a flaw in science, but a basic principle of how civilisation works. The sociologist Niklas Luhmann distinguished trust from confidence.[3] Trust is a conscious acceptance of risk: I can't be certain, but I proceed anyway. Confidence is quieter: this is familiar enough that I expect it to hold. We don't decide to trust the shoulder every time we approach it. We just start climbing.

That's where the comforting version of Newton's image starts to fall apart. We assume a claim becomes a giant because it was solid, and was then stood upon. But maybe it's the other way: in practice a claim becomes solid once enough people have already chosen to stand on it. The sociologist Robert Merton named this pattern the Matthew effect – recognition breeds further recognition, no matter the underlying merit.[4] A finding with one citation is a claim. A finding with a thousand citations starts to look like a part of reality, because so many people already behave as though it were. The citations become armour around the original evidence, and from the outside, a well-armoured claim and a well-tested one are hard to tell apart.

That's the uncomfortable thing the Leicester study shows. Not that scientists are careless – this is a perfectly rational system turned irrational through repetition. Researchers use antibodies because they've been used before – and why not, when previous use is evidence? But that evidence becomes circular when everyone stands on a shoulder because everyone else is already on that shoulder. Nobody is asking if the shoulder really holds. Everybody is checking that the researchers around and above them haven't fallen off yet, and that's not the same thing, even if it feels identical from inside the crowd.

A good system needs more than ways of establishing knowledge. It needs ways of reopening it. Philosopher Ian Hacking argued that things in science earn the right to be called real when they keep behaving the same way under repeated, independent manipulation, whether or not anyone wants them to; when they resist you rather than simply agreeing with you.[5] That's precisely what validating an antibody is supposed to test.

A bridge gets re-inspected on a schedule. A giant's shoulder, once enough people are standing on it, essentially never does. The antibody case is the exception: someone climbed back down and asked whether the giant's shoulder still fought back the way it was supposed to. It turned out that 84 per cent of the time, nobody else had bothered to do the same. There was nothing underneath but the weight of everyone who had already decided to stand there.

Having moved between millions of hands, the words on the two-pound coins are almost worn smooth – proof that a sentence can keep circulating long after anyone stops reading it closely. Perhaps that image is more true than the giants. We aren't climbing on shoulders that were tested and found solid. We're passing something hand to hand, trusting its value because it has already changed so many hands, and calling the result knowledge.

REFERENCES

1 Michael Biddle, Jemma Cooper, Katherine Blades, Dominic Ruddy, Eva M. Krockow & Harvinder Virk (2026) »Inadequate antibody validation places substantial numbers of animal and human tissue samples at risk of waste«. PLOS Biology, 6 October 2026. https://journals.plos.org/plosbiology/article?id=10.1371/journal.pbio.3003984

2 Isaac Newton, letter to Robert Hooke, 5 February 1676.

3 Niklas Luhmann (1988) »Familiarity, Confidence, Trust: Problems and Alternatives«. In Diego Gambetta (ed.) *Trust: Making and Breaking Cooperative Relations*. Oxford: Basil Blackwell.

4 Robert K. Merton (1968) »The Matthew Effect in Science«. Science, 5 January 1968. https://pubmed.ncbi.nlm.nih.gov/5634379

5 Ian Hacking (1983) *Representing and Intervening: Introductory Topics in the Philosophy of Natural Science*. Cambridge: Cambridge University Press.

08

OCTOBER 2026
Essay 038

Getting Away With Not Thinking

> *We tend to think that intelligence is about inference – thinking and reasoning. What if some of that work is already done by an agent's shape, before any brain gets involved?*

Yesterday, a paper in Science Advances described a soft robot inspired by the feather star, a marine animal whose feathery arms allow it to move through water. The robot uses two pneumatic actuators to control four flexible wings: by changing the timing of those inputs, the researchers can make the robot flap upwards, move forwards or backwards, hover, or spin around its own axis. A conventional design would typically need at least six actuators to produce the same range of motion – here, much of the work is done by the material itself.[1]

»Discovery, design, and development are less separate explanations than different routes to the same outcome: useful behaviour folded into a body«

The researchers call this »mechanical intelligence«. The phrase points to a deeper question than whether a robot can be intelligent: how much of what we call behaviour actually needs to be computed? There are at least three different ways a body can end up with such competence already built in: it can be designed, it can be discovered, or it can be developed.

In 1948, the cybernetician W. Grey Walter built two small autonomous machines, Elmer and Elsie. Their equipment was almost comically simple: some light sensors, a few motors, and a handful of electronic components. Yet the tortoise-like robots exhibited surprisingly complex behaviour as they moved towards light, responded to obstacles, and interacted with their environment.[2] Walter had not programmed miniature animals. He had arranged a small set of mechanisms so that the behaviour emerged from their interaction. This is the first of the three ways: design. Intelligence is built in by arranging the mechanism well enough to get a useful behaviour out of it.

Decades later, artificial-intelligence researchers Rolf Pfeifer and Josh Bongard gave this intuition a more systematic name: »morphological computation«. The idea is that the body can perform part of the work that would otherwise have to be performed by a brain or controller.[3] A flexible knee joint, for example, can absorb irregularities in the ground without a computer calculating exactly how far it should bend. This is a different claim from architect Frei Otto's form-finding, where a soap film or a hanging chain settles into its shape once and stays there. Here, the shape keeps working – responding differently, moment to moment, as the ground or the water changes underneath it.

The nematode *Caenorhabditis elegans* is a very simple worm. Carrying only 302 neurons, it still navigates its environment, finds food, and responds to threats. In 2023, a team at Georgia Tech built a limbless robot to test how much of the worm's undulating locomotion could be reproduced without complex neural-style control, and found that a surprising amount of it could, because the mechanics of a flexible body moving through a resistive medium take over part of the job that the nervous system would otherwise have to do.[4]

As author and researcher Max Bennett argues in *A Brief History of Intelligence*, this pushes the history of intelligence much further back than the appearance of brains that look anything like ours. Intelligence may not begin with the ability to construct an internal representation of the world. Maybe it begins with an organism whose body is simply capable of doing something useful in the world, like swimming towards food.[5] But how does a body acquire such useful capabilities in the first place, rather than having them designed in by someone who already understands the mechanism? That's the second way: discovery.

Artist Theo Jansen's kinetic *Strandbeests* sculptures offer one answer. Jansen searched through different combinations of leg dimensions using a genetic algorithm, selecting those that produced increasingly effective walking mechanisms.[6] Propelled by the wind, the resulting oversized beach creatures appear to walk on their own. While designed by a human, their particular geometry was discovered through an evolutionary process of variation and selection. The intelligence wasn't written into a programme telling each leg what to do. It was found in the structure itself.

And there's a third possibility: neither design nor search, but development. In an experiment with the slime mould *Physarum polycephalum*, researchers placed food sources at positions that corresponded to cities around Tokyo. The organism initially spread through the available space, then strengthened some connections while abandoning others. The resulting network was remarkably similar both visually and in terms of efficiency and robustness to the city's existing railway network.[7] The organism didn't calculate a map of Tokyo. The solution emerged from growth, reinforcement, and withdrawal.

But the categories blur on inspection: Jansen's »discovery« was still written by a person, and the robots' »design« by people who understood their materials. The slime mould's »development« is itself a product of evolution's much longer search. Discovery, design, and development are less separate explanations than different routes to the same outcome: useful behaviour folded into a body.

And this may be true of humans, too. When a hand catches a falling glass, the brain is certainly involved. But the movement isn't a miniature physics simulation followed by a conscious instruction to every muscle. Shape, elasticity, reflexes, learned motor patterns, and the dynamics of the body absorb much of the problem.

We tend to think of intelligence as something that happens inside a head and then gets expressed through a body. Perhaps that gets the order backwards. Maybe the remarkable thing about intelligence isn't how much a system can think, but how much it can get away with not thinking.

REFERENCES

1 Haitao Qing, Yuanhang Zhu, Jiacheng Guo, Caizhi Zhou, Haoze Sun, Haibo Dong, Daniel Quinn & Jie Yin (2026) »Minimal-Actuation Feather Star–Inspired Soft Swimmers for Multimodal 3D Maneuverability«. Science Advances, 7 October 2026. https://doi.org/10.1126/sciadv.aeg9211

2 W. Grey Walter (1950) »An Imitation of Life«. Scientific American, May 1950. https://www.scientificamerican.com/article/an-imitation-of-life/

3 Rolf Pfeifer & Josh Bongard (2007) *How the Body Shapes the Way We Think*. Cambridge, MA: MIT Press.

4 Tianyu Wang, Christopher Pierce, Velin Kojouharov, Baxi Chong, Kelimar Diaz, Hang Lu & Daniel Goldman (2023) »Mechanical Intelligence Simplifies Control in Terrestrial Limbless Locomotion«. Science Robotics, 20 December 2023. https://www.science.org/doi/10.1126/scirobotics.adi2243

5 Max Bennett (2023) *A Brief History of Intelligence*. New York: HarperCollins.

6 Marcus Fairs (2014) »'I try to make new forms of life,' says Strandbeests creator Theo Jansen«. Dezeen, 12 December 2014. https://www.dezeen.com/2014/12/12/strandbeests-theo-jansen-interview-wind-powered-machines-new-species/

7 Atsushi Tero, Seiji Takagi, Tetsu Saigusa, Kentaro Yumiki, Ryo Kobayashi, Dan Bebber, Mark Fricker & Toshiyuki Nakagaki (2010) »Rules for Biologically Inspired Adaptive Network Design«. Science, 22 January 2010. https://www.science.org/doi/10.1126/science.1177894

09

OCTOBER 2026
Essay 039

Live From Fort Hood

> *The US has decided to broadcast an execution. What changes when a public punishment becomes a spectacle watched by people who are nowhere near it?*

Yesterday, US Defence Secretary Pete Hegseth said that the execution of mass murderer Nidal Hasan, a former major convicted of killing 13 people at the Fort Hood army post in 2009, would be »public«.[1] Later in the day, the Pentagon clarified that it would be livestreamed.[2]

»Whoever thinks the sentence is just should be content to see it carried out in the open; whoever thinks it's wrong may want it seen for exactly that reason«

The word had changed its meaning without changing its spelling. Historically, a public execution meant that people gathered to watch the state kill a person. Now it can apparently mean that the state distributes live footage of their execution to people who don't need to leave their sofas to watch it, and who can turn off the screen whenever they choose.

In 1936, thousands gathered in Owensboro, Kentucky, to watch the hanging of rape convict Rainey Bethea. It remains the most recent public execution in the United States, and the press attention that followed led Kentucky to ban the practice two years later.[3] The philosopher Michel Foucault identified the flaw in the scaffold: it was meant to show the state's power, but a crowd could side with the condemned. That, he argued, is why punishment moved out of sight.[4]

A livestream looks like a way to keep the spectacle and lose the crowd. A crowd stands beside the condemned and the executioner. Its members react to one another, which can change what the event means. A livestream disperses the audience and delivers the image into each viewer's private space. The state chooses the time, the camera, and the access. There is no gathering that might become a protest. There are only silent viewers.

But a viewer is not a witness. Writer Albert Camus described how his father came home sick after having gone out to watch a man guillotined. The spectacle hadn't produced the response its organisers wanted. It had produced revulsion.[5] A livestream might produce that, or curiosity, or indifference, or the small thrill of watching something forbidden. The state can make an execution visible. But it can't decide what the visibility will mean.

There's a case for publicity. English philosopher Jeremy Bentham thought that what can be seen can be scrutinised.[6] The Italian jurist Cesare Beccaria made a different point in 1764: execution is a »terrible but momentary spectacle«, weaker as a deterrent than the more lasting example of a lifetime of forced labour.[7] Broadcasting an execution makes the moment reproducible. It doesn't make the lesson last.

The argument cuts across the usual lines. Whoever thinks the sentence is just should be content to see it carried out in the open; whoever thinks it's wrong may want it seen for exactly that reason. Camus believed that support for capital punishment depends on not looking at it. Hegseth's bet is exactly the other way around.

To show the ending is not to show the process that's supposed to justify it. Hasan's attack took place nearly 17 years ago, in 2009. He was sentenced to death in 2013, and the appeals ran on for years after that. All of this will be compressed into a few minutes of footage. The US Army has set the execution for 13:00 on Thursday 3 December, at Fort Hood itself, the site of the shooting, rather than at Fort Leavenworth where Hasan is held. It's the same weekday as the attack, and close to the same hour.[8] The execution is staged as an echo of the crime, as if repetition could complete what time left unfinished.

But repetition can't supply what the arithmetic lacks. Hasan killed thirteen people and can be executed only once. Philosopher Immanuel Kant held that murder demands death, because only death is equal to it.[9] By that measure, the state is owed thirteen deaths and can collect one. So whatever one thinks of the death penalty, equality can't be what justifies this execution. Hegseth names another purpose: people need to understand there are serious consequences. That's a claim about the audience, and it's the one a broadcast serves.

The Islamic State also understood execution as media: its 2014–2015 execution videos were made to be distributed. That was the point.[10] The Pentagon's livestream will differ radically in legal process, but the form raises a question: when an execution is distributed to produce an effect, what separates accountability from theatre? At his trial, Hasan called no witnesses and made no closing argument[11]; his court-appointed lawyers reportedly feared that he wanted a death sentence, though he disputed this. If the lawyers were right, Hegseth may be giving Hasan the stage that he sought.

The public is everywhere now, and the gathering is nowhere. If the lesson Hegseth wants taught is learned at all, it will be learned one viewer at a time. A livestream can make the moment visible, but it can't decide whether the footage reads as justice, deterrence, or closure. When a punishment can't repair the harm, what's the spectacle meant to complete?

REFERENCES

1 Herb Scribner & Avery Lotz (2026) »Hegseth says public can watch Fort Hood shooter's execution«. Axios, 8 October 2026. https://www.axios.com/2026/10/08/hegseth-nidal-hasan-firing-squad-public-tv

2 Eleanor Watson & Joe Walsh (2026) »Fort Hood shooter's execution will be livestreamed, Pentagon says«. CBS News, 8 October 2026. https://www.cbsnews.com/news/hegseth-fort-hood-shooter-execution-public-nidal-hasan/

3 Associated Press (2011) »After 75 years, last public hanging haunts city«. Fox News, 2011. https://www.foxnews.com/us/after-75-years-last-public-hanging-haunts-city

4 Michel Foucault (1975) *Discipline and Punish: The Birth of the Prison*. Translated by Alan Sheridan. London: Penguin, 1991 edition. pp. 7–9, 57–69.

5 Albert Camus (1957) »Reflections on the Guillotine«. In *Resistance, Rebellion, and Death*. Translated by Justin O'Brien. New York: Knopf, 1961. Part 1.

6 Gerald J. Postema (2019) *Utility, Publicity, and Law: Essays on Bentham's Moral and Legal Philosophy*. Oxford: Oxford University Press.

7 Cesare Beccaria (1764) *Of Crimes and Punishments*. Chapter 28, »Of the Punishment of Death«. https://www.laits.utexas.edu/poltheory/beccaria/delitti/delitti.c28.html

8 Austin Campbell (2026) »Army sets execution date for Nidal Hasan by firing squad at Fort Hood«. Task & Purpose, 7 October 2026. https://taskandpurpose.com/news/nidal-hasan-execution/

9 Immanuel Kant (1797) *The Metaphysics of Morals*. Translated by Mary Gregor. Cambridge: Cambridge University Press, 1996. Academy edition 6:331–337.

10 Charlie Winter (2015) *The Virtual 'Caliphate': Understanding Islamic State's Propaganda Strategy*. London: Quilliam Foundation.

11 Jana Pruett (2013) »Accused Fort Hood shooter rests case without calling witnesses«. Reuters/NBC News, 21 August 2013. https://www.nbcnews.com/news/us-news/accused-fort-hood-shooter-rests-case-without-calling-witnesses-flna6c10968124

10

OCTOBER 2026
Essay 040

The Track and the Train

> *A train capable of 260 km/h[1] runs on a railway where its average speed is closer to 110. The train is new. The problem is not.*

Yesterday, the computer scientist Arvind Narayanan told journalist Ezra Klein, on his New York Times podcast, that artificial intelligence is a »normal« technology:[2] powerful, but something that will change society slowly, because the world around it changes slowly. His example was a railway. Amtrak's new Acela trains, built by the French manufacturer Alstom, can run at about 260 km/h. But the Northeast Corridor, which runs some 735 km from Washington via New York to Boston, was laid out when a fast train meant 80 to 100 km/h, and its curves are built for that reality. The average speed is still about 110. It took fourteen attempts at computer modelling before the new trains were allowed to begin testing on the line. The new had to be shown to fit the old.[3] The train got faster; the journey didn't. AI, Narayanan said, is mostly the train rather than the track.

»the gap between maximum and actual performance exposes a familiar confusion: mistaking the capability of a single component for the capability of the entire system«

That doesn't mean the trains were a bad investment. They carry about a quarter more passengers,[3] which matters even if speed doesn't transform the journey. But the gap between maximum and actual performance exposes a familiar confusion: mistaking the capability of a single component for the capability of the entire system.

The historian of technology Thomas Hughes called a part that has fallen behind the rest of a system a »reverse salient«.[4] Components advance at different rates, and the one that falls behind draws the inventors' attention. The railway seems a textbook case. But if the bottleneck is so visible, why don't we put all our effort into removing it? A train can be bought, tested, delivered, and photographed at a launch. A railway corridor is a tangle of land, bridges, curves, signals, operating rules, and budgets. Improving it means coordinating institutions for decades. The train is a product. The track is a political settlement.

The sums show the difference. The new trains cost about $1.6 billion; bringing the corridor up to what they can do was estimated in 2024 at more than $100 billion.[3] The track is not being ignored, only worked on to a different clock. Under the Hudson, a $16 billion tunnel, the first new rail tunnel beneath the river in more than a century, launched its boring at the end of September. It is due to open in 2035.[5] The existing tunnels opened in 1910 and were flooded by Hurricane Sandy in 2012. In February, when frozen federal funding ran out, work stopped for weeks.[6] The trains took nine years from order to passengers. The track takes decades.

Stewart Brand, a maverick environmentalist, entrepreneur, and writer, offers another way to see this in *How Buildings Learn*. Buildings, he argues, are made of layers that change at different speeds: the site endures, the structure lasts for decades, while services and fittings are replaced again and again. When the layers are tied tightly together, the slow ones block the quick.[7] A railway is a similar stack of time. Trains are replaced every few decades (the first Acela fleet has run for more than twenty-five years), while some of the bridges and tunnels they pass through are more than a century old.

The economic historian Paul David made a related argument about electrification. Electric motors didn't immediately produce the productivity gains their technical advantages seemed to promise. Factories had to be reorganised around them. The gains came from changing the system around the component.[8] Nor does the work end there: bottlenecks move. Clear one and another becomes decisive. This is the expensive part of progress. It's easier to buy faster trains than to straighten a curve through a densely settled corridor. It's easier to install new software than to change a planning process, a supply chain, or a division of responsibility. The new part works; the world around it doesn't change.

The people who sell the train and the people who look after the track want different things. A manufacturer delivers a finished object against a contract. A railway authority has to pay costs now, with disruption, for benefits that come years later. The first can point to what it has delivered. The second has to explain what must be moved or torn up before anything gets better.

Acela trains lean into curves, so they can take bends faster without every bend being straightened. It's ingenious, and a sensible adaptation to what the past left behind. But the tilting mechanism leaves the larger question intact. What, exactly, has been improved? The train can do more than the railway permits. Both the achievement and the limit are real. Progress is not always held back by a lack of ingenuity. Sometimes ingenuity is concentrated on the part of the system that can change, while the part that sets the limit remains untouched.

The train leans, politely, into a century-old curve.

REFERENCES

1 I'm using the metric system here as I live in Europe. If you prefer imperial, 735 km is approximately 457 miles, 80 to 100 km/h is 50 to 60 mph, and 260 km/h translates to around 160 mph; 110 km/h is roughly 70 mph.

2 Ezra Klein & Arvind Narayanan (2026) »Intelligence Isn't Power«. The Ezra Klein Show, The New York Times, 9 October 2026. https://www.nytimes.com/2026/10/09/opinion/ezra-klein-podcast-arvind-narayanan.html

3 Mark Walker (2024) »After years of delays, Amtrak moves toward faster trains in the Northeast«. 13 January 2024 (republished by The Boston Globe). https://www.bostonglobe.com/2024/01/13/nation/after-years-delays-amtrak-moves-toward-faster-trains-northeast/

4 Thomas P. Hughes (1983) *Networks of Power: Electrification in Western Society, 1880–1930*. Baltimore: Johns Hopkins University Press.

5 »Hudson Tunnel Project begins major boring milestone«. NJBIZ, September 2026. https://njbiz.com/hudson-tunnel-project-tunnel-boring/

6 Bryan Gottlieb (2026) »Gateway Hudson Tunnel Work Resumes Even as Funding Uncertainty Mounts«. Engineering News-Record, 10 March 2026. https://www.enr.com/articles/62642-gateway-hudson-tunnel-work-resumes-even-as-funding-uncertainty-mounts

7 Stewart Brand (1994) *How Buildings Learn: What Happens After They're Built*. New York: Viking. Chapter »Shearing Layers«.

8 Paul A. David (1990) »The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox«. American Economic Review 80(2), pp. 355–361.

11

OCTOBER 2026
Essay 041

Musk and the Mask

> *You can make a machine speak like a person, but when does speaking like someone become speaking for them?*

In The Atlantic yesterday, staff writer Megan Garber reviewed director Alex Gibney's new documentary *Musk*, a portrait of a man whose influence has grown to something approaching systemic power.[1] The film isn't out yet, and I haven't seen it, but one device that Garber describes interests me more than the portrait itself. When trillionaire Elon Musk declined to participate, Gibney's team compiled a database of his recorded public statements and used it to create an animated version of him. The avatar speaks Musk's actual words, but in contexts where he was never asked the questions it appears to answer. According to Gibney, the result is intentionally ironic: a simulation of a man who has himself speculated that reality might be a simulation.[2]

»The avatar borrows words; the model invents them. But both are attempts to answer the same question – what would this person say now? – and neither has access to the one thing that could actually answer it: the person's own encounter with a new situation«

This isn't just a documentary using archival footage or a montage of quotations. The avatar gives the words a body, a voice, and the appearance of presence. It makes it seem as if Elon Musk is there responding, when in fact he wasn't there and didn't respond.

This made me think of British sculptor Marc Quinn's *Self*, a series of casts of the artist's head made from his own frozen blood.[3] The sculpture is constructed from Quinn and resembles Quinn, yet it is clearly not Quinn. It can't act on his behalf. Gibney's avatar is more complicated because it can speak. Its words are really Musk's, but someone else turned them into answers to questions he was never asked. The result is neither Musk nor not Musk. It's a version of him assembled from material he supplied, made to perform in a different context.

This is where the device opens onto something larger, even though the mechanism isn't quite the same. Gibney's team selected and rearranged words that Elon Musk had actually said. Today's large language models do something stranger: give them enough of a person's writing, speeches, and interviews, and they learn a pattern they can extend into sentences that person never actually said. The avatar *borrows* words; the model *invents* them. But both are attempts to answer the same question – *what would this person say now?* – and neither has access to the one thing that could actually answer it: the person's own encounter with a new situation.

A living person responds to changes in the world. But if a machine were trained on Stephen King's novels and continued producing novels after his death, it would be drawing on a body of work formed in response to a world that no longer exists. King's response to whatever came next might have differed from anything in the training material. No statistical refinement, however exact, could supply that missing relationship with the present. The model could produce a genuinely new novel, but it couldn't establish that King would have written it, let alone that he did.

Would it be wrong to publish such a novel? Not necessarily, provided its origins were clear. The question is whether it should be presented as a novel by King. A work in someone's style and a work authored by that person are different things. Authorship involves making particular choices, taking responsibility for them, and presenting the result to readers under one's name. That's not the same as saying that machines can't create, or that they can't be creative. It's an argument against confusing the capacity to generate with the identity of the person being imitated.

The same gap seems to exist in Gibney's film, even though Elon Musk is alive. The wealthiest person in the world continues to respond to a changing environment. His recorded words, detached from the circumstances in which he spoke them and inserted

into new ones, can't stand in for those future responses. Are the avatar's words authentic? Yes. But real words don't make the performance real. And what about questions of permission and compensation when someone's recorded life becomes the raw material for a synthetic likeness?

The word »persona« offers a fitting image here. Its deeper etymology is disputed, but it might have come to us through the Latin word for an actor's mask.[4] The mask is shaped as a face, but someone else can animate it, and the voice spoken through it doesn't have to belong to the person it represents.

Musk's avatar is a mask made from Elon Musk's own words. The question isn't whether those words belong to him. It's who gets to make them speak.

REFERENCES

1 Megan Garber (2026) »'Elon Is a Nuke'«. The Atlantic, 10 October 2026. https://www.theatlantic.com/culture/2026/10/elon-musk-documentary-alex-gibney/688942/

2 LDS Team (2026) »Gibney's Musk Documentary Uses AI-Simulated Interview«. Let's Data Science, 9 September 2026. https://letsdatascience.com/news/gibneys-musk-documentary-uses-ai-simulated-interview-eda0ccec

3 National Portrait Gallery »Marc Quinn, 'Self'«. npg.org.uk. https://www.npg.org.uk/collections/search/portrait/mw138260

4 The theory that persona derives from the Etruscan phersu ("mask") is disputed, and competing links to the Latin personare ("to sound through") are themselves uncertain. See Douglas Harper, Online Etymology Dictionary, entry »person«. https://www.etymonline.com/word/person

12

OCTOBER 2026
Essay 042

Room Left to Change

> *Artificial intelligence is transforming some companies while barely touching others. The difference may have less to do with the technology than with what happened before it arrived.*

In a Fortune essay published yesterday, senior advisory partner Ming Lu argues that the faster we innovate, the slower we grow.[1] While artificial intelligence is spreading rapidly, its effect on average productivity barely registers. A recent survey by McKinsey found that 88% of organisations use AI in at least one business function, but only 6% of all organisations surveyed are high performers.[2]

»Companies aren't ancient civilisations, and inefficiency isn't collapse. But the underlying problem is recognisable: every process exists for a reason, and as they accumulate, the problems they were built to solve may already have changed or disappeared«

There are two ways to read this. One is that AI has yet to deliver on its promises. The other is that while a handful of companies are doing extremely well, most are getting nothing out of it – and the average between the two blurs them into a single disappointing number. In »The Median Isn't the Message«, the evolutionary biologist Stephen Jay Gould argued that variation isn't an imperfection surrounding some underlying truth. It's the reality. The average is the abstraction.[3]

This way of thinking has an unexpectedly specific history. In 1835, Belgian astronomer Adolphe Quetelet introduced the idea of *l'homme moyen* – the average man. Quetelet borrowed a method astronomers used to pin down a star's true position: average enough imperfect measurements and the errors cancel out, leaving the real value behind. He then applied the same thinking to people, assuming that behind every person stood a »true« average man, and that each person's distance from it was an error, rather than a fact about them.[4] That was of course a mistake: there's no single ideal human that real people are flawed copies of. The deviations weren't errors of measurement; they were just the differences between actual human beings.

We still tend to repeat Quetelet's mistake. The more interesting question about the McKinsey study isn't the average, but what separates the two groups beneath it. Lu's own answer points towards organisational design: the companies that benefit most from AI have used it to redesign their workflows from the ground up, while the others have simply introduced it into their existing structures.

Why should that make such a difference? One possible answer comes from a very different field: archaeology. In *The Collapse of Complex Societies*, archaeologist and anthropologist Joseph Tainter proposed that societies respond to problems by becoming more complex – specialising, adding rules, building new coordination structures. Each such addition can solve a real problem, but complexity comes with a maintenance cost, and the benefits of adding to it eventually diminish. In the end, a society can reach a point where its accumulated machinery consumes more resources than its solutions are worth.[5]

Tainter's own example is the late Roman Empire: faced with invasion and a shrinking surplus, successive emperors responded the way the empire always had – with a larger army, more administration, expanded territory – until there was nothing left to conquer that was worth the cost of conquering it. The only options left were debasing the currency and taxing an exhausted population harder, just to keep paying for a complex apparatus that no longer paid for itself.[5]

Companies aren't ancient civilisations, and inefficiency isn't collapse. But the underlying problem is recognisable: every process exists for a reason, and as they accumulate, the problems they were built to solve may already have changed or disappeared.

Which brings us back to AI. A company with room to reorganise can use it to change how its work is being done. One already entangled in layers of procedure and coordination may simply add AI as yet another layer, acquiring a more powerful tool bolted onto a system that stays the same.

This is my own interpretive leap, not a conclusion in Lu's figures: the survey doesn't show that accumulated complexity explains the gap between the high-performing 6% and everyone else. But it reframes the question worth asking. Perhaps what separates them isn't so much how well each company understood the technology as how much institutional room each had left to change before the technology arrived.

If that's right, the companies struggling with AI may be failing at something considerably older – the same thing that eventually slows any institution that keeps solving problems by adding to itself without ever subtracting. Companies that fail to profit from AI may not fail at the technology. Some ran out of room before it even arrived.

REFERENCES

1 Ming Lu (2026) »The paradox at the core of the AI revolution: The faster we innovate, the slower we grow«. Fortune, 11 October 2026. https://fortune.com/2026/10/11/innovation-paradox-ai-productivity-ming-lu/

2 Dan Tinkoff, Lieven Van der Veken, Michael Chui & Tara Balakrishnan (2026) »The state of AI in 2026: On the road to ROI«. McKinsey & Company, 25 August 2026. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

3 Stephen Jay Gould (1985) »The Median Isn't the Message«. Discover, June 1985. https://www.stat.berkeley.edu/users/rice/Stat2/GouldCancer.html

4 Adolphe Quetelet (1835) *Sur l'homme et le développement de ses facultés, ou Essai de physique sociale*. Paris: Bachelier. In English: Adolphe Quetelet (1842) *A Treatise on Man and the Development of His Faculties*. Translated by Robert Knox. Edinburgh: W. and R. Chambers.

5 Joseph A. Tainter (1988) *The Collapse of Complex Societies*. Cambridge: Cambridge University Press.

13

OCTOBER 2026
Essay 043

A Fridge Frozen

> *Sometimes the inability to change can be a product's best feature. We're so used to measuring progress in new capabilities that we rarely ask what might be gained by refusing one.*

Yesterday, Reuters reported that hackers attacking supermarket refrigerators in France had forced up the temperature to 50°C. Hundreds of shops were affected, food was thrown away, fires were reported. The attack exploited refrigeration management systems made by the Danish company Danfoss, which said the affected units were running outdated firmware and that for a long time they had advised the affected companies to update them.[1]

»We have quietly accepted that more connectivity is better, that products should be released early and corrected later, and that the customer should be responsible for maintenance. But those aren't laws of nature«

A lot of people will instinctively feel that the obvious question is why the systems hadn't been updated. But let's take a step back and ask the more interesting one: why does a refrigerator need to be updated in the first place?

The answer is efficiency. A retailer managing thousands of refrigerators can monitor and adjust them without sending technicians to every store. The manufacturer saves money; the customer absorbs the cost, in time, attention and risk. This is the pattern that security economist Ross Anderson has long pointed to, in which the party best placed to prevent a failure is rarely the party who pays for it.[2]

A refrigerator has a pretty stable job: it moves heat out of a box, and the laws of thermodynamics rarely change beneath it. But these ones in France had been given a network connection. That turned them into what legal scholar Jonathan Zittrain calls »tethered appliances«, reachable by whoever controls – or compromises – that connection, exposing them to adversaries that their original purpose never required them to face.[3] The connection gave the fridges a potential enemy.

Cybernetician W. Ross Ashby's *Law of Requisite Variety* offers a useful principle here: a control system needs enough adaptability to cope with the disturbances in its environment, and no more.[4] A calendar application must contend with changing operating systems, competing apps, and evolving threats – what the biologist Leigh Van Valen called a *Red Queen race*, in which standing still means falling behind.[5] The name comes from a character in Lewis Carroll's *Through the Looking-Glass*, who has to keep running just to stay in the same spot.[6] A refrigerator clearly has rather less to contend with. Adaptability is good, but we might want to ask how much of it a particular thing needs, and what risk we take by giving it more.

We could, for instance, design a refrigerator that can report its temperature to a dashboard but can't receive any commands in return. If it needs to be updated, a technician pays a visit and connects a laptop. Wouldn't that be slower and more expensive? Yes. But it would also make remote attacks considerably harder, and – importantly – leave the cost of fixing the manufacturer's own mistakes with the manufacturer, rather than transferring it to the customer.

There's a name for the underlying idea, from economist Thomas Schelling: a commitment device.[7] In Homer's *Odyssey*, Ulysses preemptively has himself tied to the mast so that he can't act on his future desire to follow the singing of the Sirens.[8] He deliberately gives up the freedom to change his mind, and that restraint makes the arrangement work. We recognise the same principle in a piggy bank that must be broken to get the coins out, or a contract that restricts the future choices of those who sign it. Those aren't defective systems waiting for an upgrade. Their limitations are part of their purpose.

Why not design products this way? A refrigerator that can't be remotely altered becomes more trustworthy precisely because it has fewer capabilities. Changing its behaviour would require a deliberate, physical intervention. A manufacturer could sell this as a guarantee of integrity, renewed every year it remains untouched by anyone but the technician. Instead of a fast and cheap release that makes corrections necessary, it would profit from creating a product that doesn't need to be corrected. It would benefit from shipping a finished product.

We don't need to stop at refrigeration. How would cars change if manufacturers competed on how long they could remain secure without remote intervention, rather than on how quickly they could push an update? What about pacemakers? Is remote control really worth the risk of getting hacked, for a device that regulates your heart?

We have quietly accepted that more connectivity is better, that products should be released early and corrected later, and that the customer should be responsible for maintenance. But those aren't laws of nature. They are choices shaped by the economics of the companies that make our products.

Perhaps we're overlooking entire markets by not going against the prevailing sentiment. Progress includes the ability to choose which capabilities not to build. Sometimes that choice makes a product better in a way its competitors have forgotten to offer.

Perhaps the next advance in refrigeration isn't a smarter fridge, but one that knows what it doesn't need to be able to do.

REFERENCES

1 Reuters (2026) »Hackers force French supermarket fridges to heat up to 50C in 'unprecedented' cyber-attack«. The Guardian, 12 October 2026. https://www.theguardian.com/world/2026/oct/12/cyber-attack-defrosts-refrigerators-france-supermarket-chains-intermarche-cooperative-u

2 Ross Anderson (2001) »Why Information Security Is Hard — An Economic Perspective«. 17th Annual Computer Security Applications Conference (ACSAC), New Orleans.

3 Jonathan Zittrain (2008) *The Future of the Internet — And How to Stop It*. New Haven: Yale University Press.

4 W. Ross Ashby (1956) *An Introduction to Cybernetics*. London: Chapman & Hall.

5 Leigh Van Valen (1973) »A New Evolutionary Law«. Evolutionary Theory, vol. 1, Chicago.

6 Lewis Carroll (1871) *Through the Looking-Glass, and What Alice Found There*. London: Macmillan.

7 Thomas Schelling (1960) *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.

8 Homer (c. 8th century BC) *The Odyssey*.

14

OCTOBER 2026
Essay 044

A Way to Get Back Home

> *You can stick to a rule long after you've stopped being the person who made it. What does it mean to keep a promise to someone you no longer are?*

In a Guardian interview published yesterday, Simon Hattenstone meets Karl Bushby, a former British paratrooper who left his mother's house in Hull on 1 November 1998, flew to the southern tip of South America, and has been walking back ever since.[1]

»Bushby didn't set out to make art – he was just trying to get home – but a simple rule, sustained for decades against everything that threatened to dissolve it, gave his life something of an aesthetic structure all the same«

Bushby made two rules before setting out: no transport was allowed to advance the journey, and if circumstances ever forced him off his route, he had to return and continue from the exact point where he'd stopped. He expected to be away for twelve years. It has taken twenty-eight. He's crossed the thin ice of the Bering Strait, swum the Caspian Sea, survived kidnappings and near-starvation, lost years to borders and visa offices, and put 58,000 kilometres on foot behind him. He just swam across the English Channel, and is now walking the final stretch home.

In the Atacama Desert, a stranger named Michael Bier gave him $100 and asked a favour in return: deliver a postcard to a friend in Vienna. Bushby carried it for twenty-seven years before handing it over in person. He could have posted it within a week. Instead, the errand took nearly three decades. That postcard is a miniature version of the whole journey: an ordinary promise made extraordinary by refusing to take the easy way out.

Artists often use constraint to find the unexpected. Richard Long turned a walk into a line drawn across the landscape.[2] Matthew Barney made the struggle to make a mark under physical restraint the idea of the work itself.[3] Bushby didn't set out to make art – he was just trying to get home – but a simple rule, sustained for decades against everything that threatened to dissolve it, gave his life something of an aesthetic structure all the same. Not a sculpture or a film, but the relationship between a rule and a life. What exactly did the experiment establish – that persistence can beat almost any obstacle? That's true enough, but it hardly accounts for twenty-eight years. He could have stopped, he could have changed the rules, he could have hopped on a plane back home. He was free to give up, but also free to push on.

That freedom came at a real cost. He left a son behind, and missed his childhood. He has yet to meet his twelve-year-old granddaughter. Bushby calls his undertaking »totally self-centred«, and tells himself that finishing it might mean more to his son than being there for him would have. He doesn't sound fully convinced by his own argument, Hattenstone notes. A purpose can give a life its shape while making it harder to live.

The premise suggests this was a solitary feat. Not quite. Bushby depended on strangers for food, money, and care when he was ill; he reckons 99.999 per cent of the people he met were the best of humanity, and calls the walk a »we« project. The rule was his. The means of keeping it were often supplied by others.

A documentary film is already in the works. Werner Herzog has spent much of his career pointing a camera at people who impose an extravagant law on themselves and then refuse to break it: opera lover Fitzcarraldo hauling a steamship over a mountain is perhaps the clearest case.[4] But the man Herzog's camera is following into Hull isn't quite the man whose promise the film is about – it will tell the story of someone who, according to Bushby himself, no longer exists.

Bushby says he's a different person from the one who left: softer, more open to the world, less narrowly British. He doggedly stuck to the promise made by his former self, but he didn't remain that person. The destination stayed fixed but the traveller changed completely, and so the promise has outlived its own author. The undertaking that organised so many years of his life is about to end; his purpose for living, he says, is about to disappear. He knows how to walk. What comes next is the part he doesn't know how to do.

The experiment is finished, but it hasn't produced a solution. Whether keeping a promise to someone he no longer is counts as fidelity, a kind of self-imposed sentence, or both, may not be answerable from the outside. Once there was a way to get back home. Now a different person has returned.

REFERENCES

1 Simon Hattenstone (2026) »'Some things are worth dying for': Karl Bushby's incredible 28-year walk home, through deserts, kidnappers and polar bears«. The Guardian, 13 October 2026. https://www.theguardian.com/lifeandstyle/2026/oct/13/karl-bushby-incredible-28-year-walk-home-goliath-expedition-hull

2 Richard Long (1967) *A Line Made by Walking*. Tate. https://www.tate.org.uk/art/artworks/long-a-line-made-by-walking-p07149

3 Matthew Barney (1987–) *Drawing Restraint*. MoMA. https://www.moma.org/collection/works/131022

4 Werner Herzog (1982) *Fitzcarraldo*. Munich: Werner Herzog Filmproduktion/Pro-ject Filmproduktion.

15

OCTOBER 2026
Essay 045

The Statistical DJ

> *Deezer wants to let listeners filter out AI-generated music. But the question was never really whether machines were involved in making the song – it's whether anyone will answer for it.*

On Tuesday, streaming service Deezer announced a toggle to let listeners exclude AI-generated music from search results.[1] In a July report, the company noted that more than half of its daily uploads were generated by AI, averaging 90,000 tracks a day.[2] Bandcamp bans AI music,[3] Spotify labels it,[4] neither says what, exactly, is wrong with it. Perhaps the more interesting question isn't who – or what – made it, but whether anyone will answer for it.

»The distinction isn't whether a machine was involved, or even whether a human selected every note. It's where the creative decisions reside, and whether anyone is prepared to stand behind them«

For a week, I've listened incessantly to Cameron Winter's recent release, *Live at Carnegie Hall*.[5] Paul Thomas Anderson is turning the performance into an upcoming concert film. There's something deeply human about the setting: a sold-out concert hall, a grand piano, a 23-year-old musician alone on stage. Winter sings lyrics that seem to have arrived from somewhere far outside the statistical average:

You were born to break
my big hairy football arms
Like clean windows
kill the birds

In »The Rolling Stones«, there's a moment that breaks me every time. Tissues at hand, I listen to a stranger articulate something I recognise but can't quite name.

But Winter's voice and piano are amplified. The performance reaches me through technology, just as all recorded music does. The history of music is, in part, a series of arguments about which technologies count as legitimate. Recordings severed music from the moment of its performance. Drum machines displaced drummers; samplers turned existing recordings into new compositions; Auto-Tune made the correction of a voice part of its aesthetic. Each new technology prompted a question about authenticity, and each time the boundary moved.[6]

AI is different, though perhaps not in the way its opponents suppose. A musician using AI to refine a composition is making different decisions from someone generating thousands of tracks and uploading them on autopilot. The distinction isn't whether a machine was involved, or even whether a human selected every note. It's where the creative decisions reside, and whether anyone is prepared to stand behind them.

Artist Sol LeWitt put the principle neatly: »The idea becomes a machine that makes the art«.[7] Brian Eno's generative music takes this seriously. The artist establishes a system of rules and constraints, then lets it produce music. Eno didn't choose every note; often he didn't even hear the result in advance.[8] The absence of moment-to-moment intervention doesn't make the work authorless. The decisions have moved upstream, into the design of the system itself. The idea makes the art.

But what happens when the cost of generating another ten thousand variations approaches zero? Someone still shaped the conditions under which the music emerges. But who chooses which results deserve to be heard? A system might become a kind of statistical DJ, producing endless variations and using criteria no listener can see to select among them. We've moved past the question of whether some human made the music to how much responsibility they took for that music.

Winter's performance offers a possible answer. Playing and singing in real time involves countless tiny corrections: breath, pitch, timing, touch. Each is a risky decision, with no guarantee that the next moment will work. The performance is the accumulation of those decisions, made in front of thousands of people. I am in no way ruling out that a fully AI-generated song might make me cry. But hearing Winter live includes something that the recording can't fully preserve: the knowledge that at a particular moment a particular person did a particular thing, and that it could have gone very wrong.

Perhaps that's where music is heading. Recording made the musical object reproducible, pushing musicians towards live performance as a source of value. Maybe generative AI will intensify that movement. When music can be produced in unlimited quantities at negligible cost, the scarce thing might be the performance: the person who can actually play, sing, and risk something in real time.

Deezer's toggle is a crude proxy for a much better question. »Was AI involved?« is relatively easy to answer. »Did anyone actually make a decision here?« is harder. A system can embody years of artistic judgement, or produce ten thousand variations that nobody has bothered to listen to. The label can't tell us which. We're trying to distinguish music that someone stands behind from music that merely exists. And no toggle can make that distinction for us.

REFERENCES

1 Ivan Mehta (2026) »Deezer Will Soon Let You Filter Out AI-Generated Music«. TechCrunch, 14 October 2026. https://techcrunch.com/2026/10/14/deezer-will-soon-let-you-filter-out-ai-generated-music/

2 Ivan Mehta (2026) »Music Streamer Deezer Says More Than 50% of Daily Uploads Are AI-Generated«. TechCrunch, 21 July 2026. https://techcrunch.com/2026/07/21/music-streamer-deezer-says-more-than-50-of-daily-uploads-are-ai-generated/

3 Mandy Dalugdug (2026) »Bandcamp Bans AI-Generated Music: 'We Want Musicians to Keep Making Music'«. Music Business Worldwide, 15 January 2026. https://www.musicbusinessworldwide.com/bandcamp-bans-ai-generated-music-we-want-musicians-to-keep-making-music

4 Sarah Perez (2025) »Spotify to Label AI Music, Filter Spam and More in AI Policy Change«. TechCrunch, 25 September 2025. https://techcrunch.com/2025/09/25/spotify-updates-ai-policy-to-label-tracks-cut-down-on-spam/

5 Alex Hudson (2026) »Cameron Winter Announces Paul Thomas Anderson-Directed Concert Film«. Exclaim!, 25 August 2026. https://exclaim.ca/music/article/cameron-winter-announces-live-at-carnagie-hall

6 Mark Katz (2010) *Capturing Sound: How Technology Has Changed Music*, revised edition. Berkeley: University of California Press.

7 Sol LeWitt (1967) »Paragraphs on Conceptual Art«. Artforum 5, no. 10 (Summer 1967), pp. 79–84.

8 Brian Eno (1996) »Generative Music«. In Motion Magazine, 7 July 1996 (talk given at the Imagination Conference, San Francisco, 8 June 1996). https://inmotionmagazine.com/eno1.html

16

OCTOBER 2026
Essay 046

The Creative Director Is Dead, Long Live the Creative Director

> *Two fashion houses lost their creative directors on the same day. In an industry built around individual vision, institutions survive by repeatedly replacing the people who provide it.*

Yesterday, FashionUnited reported that creative director Seán McGirr was leaving Alexander McQueen while his counterpart Veronica Leoni was leaving Calvin Klein.[1] The exits follow Anthony Vaccarello's departure from Saint Laurent last week, and the fashion-industry news platform noted that »fashion's exit season has arrived«.

»The creative director has become a peculiar kind of institutional asset: the person expected on the one hand to make a fashion house distinctive, on the other to be replaceable without destroying that distinction in the process«

The coincidence seems remarkable, and marks a sharp break with fashion's tradition of creative longevity. Yves Saint Laurent led his eponymous house for four decades.[2] Valentino Garavani spent nearly half a century building his.[3] Karl Lagerfeld remained at Chanel for 36 years, until his death in 2019.[4] The shelf life of today's creative leaders in fashion, meanwhile, has been estimated at around five years.[5] The creative director has become a peculiar kind of institutional asset: the person expected on the one hand to make a fashion house distinctive, on the other to be replaceable without destroying that distinction in the process.

A fashion house isn't just a name. Its heritage includes silhouettes, materials, colours, techniques, references, an archive, a clientele, and a set of expectations about what its next collection should be. The creative director's job is to keep this inheritance alive, evolving, and relevant. Too much continuity becomes repetition; too much change threatens recognition. It's a bit like renovating a building that must remain occupied throughout the work.

Architecture has long debated what such fidelity entails. In the mid-19th century, the English architectural theorist John Ruskin regarded the marks of age and history as integral to a building, and restoration that replaced them with conjecture as a form of falsification.[6] His contemporary, the French architect Eugène Viollet-le-Duc, took the opposite position: restoration could recover a building's underlying idea, even if the building had never realised it.[7] The first approach risks preserving a work until it loses its life; the second, improving it into something it never was.

A creative director in fashion faces a version of this dilemma with every collection. Is the archive a constraint to be respected, or a set of materials from which to construct a new interpretation? Who decides when reinterpretation becomes betrayal? The answer can't just be that the designer gets to choose, because the house must outlast the designer. But it also can't be to just go back to the archive, because an archive can't decide what to do next.

In effect, the house has acquired a second body: an identity that persists beyond any one person who occupies it. The old royal proclamation has found a new application: the creative director is dead, long live the creative director.

But maybe we're looking at this backwards. The rapid turnover of creative directors is often treated as evidence of an industry in crisis, impatient with talent and incapable of sustaining a vision. Sometimes, that's probably true. But the instability might also express something fundamental about fashion itself.

In *Fashion: A Philosophy*, the Norwegian philosopher Lars Svendsen describes fashion as a system built around the tension between belonging and distinction.[8] We use shared signs to express our individuality: we want to belong, but we also want to stand apart. The difficulty is that, as a look spreads, it loses some of its power to distinguish us. What was distinctive becomes familiar; what was new becomes ordinary. Which is why fashion must keep producing difference, and then keep replacing that difference.

The same mechanism operates at the level of the fashion house. A successful creative director makes it feel both familiar and newly discovered. But what makes a house recognisable can also make it predictable. Eventually, that predictability becomes a problem. The person who made the house relevant may now be the person it needs to replace.

An industry that sells perpetual renewal has begun to apply that logic to the people designing its collections. Creative directors are no longer merely the agents of change. They have become subject to the very cycle they are hired to sustain.

So where does a fashion house's identity actually reside? In its archive, its customers, its commercial machinery, its name, or the succession of people who have interpreted it? Perhaps it resides in the productive tension between them. The house must remain itself, but nobody gets to decide alone what that means. Not even the person whose job it is to make that decision.

REFERENCES

1 Renan Botelho de Carvalho (2026) »Fashion's exit season has arrived: Veronica Leoni leaves Calvin Klein, Seán McGirr departs McQueen«. FashionUnited, 15 October 2026. https://fashionunited.uk/news/people/fashions-exit-season-has-arrived-veronica-leoni-leaves-calvin-klein-sean-mcgirr-departs-mcqueen/2026101590873

2 Encyclopaedia Britannica (2026) »Yves Saint Laurent«. Encyclopaedia Britannica. https://www.britannica.com/biography/Yves-Saint-Laurent-French-designer

3 Reuters & The Associated Press (2026) »Valentino, founder of Italian luxury empire, dies at 93«. Via NBC News, 2026. https://www.nbcnews.com/business/business-news/valentino-italian-fashion-dies-rcna254824

4 Erin Vanderhoof (2019) »Karl Lagerfeld Dies at 85«. Vanity Fair obituary, 19 February 2019. https://www.vanityfair.com/style/2019/02/karl-lagerfeld-dies-at-85

5 Marc Bain (2020) »Creative directors in fashion have a five-year shelf life«. Quartz, 2020. https://qz.com/1870462/creative-directors-in-fashion-have-a-five-year-shelf-life

6 John Ruskin (1849) *The Seven Lamps of Architecture*. London: Smith, Elder & Co.

7 Viollet-le-Duc, E. (1854–1868) *Dictionnaire raisonné de l'architecture française du XIe au XVIe siècle*. Paris: B. Bance / A. Morel.

8 Lars Svendsen (2006) *Fashion: A Philosophy*. London: Reaktion Books.

17

OCTOBER 2026
Essay 047

The Seed and the Stunt

> *Everyone watched a car flip through the air. But who decides what a new capability is for: the people who invent it, or the people who pick it up next?*

Yesterday, Tesla unveiled a flying car. Yep, a flying car. At a SpaceX test site in Texas, its Roadster drove up a curved ramp into an upside-down position, left the ramp, flipped in mid-air and landed on its thrusters. Oh, and then it jumped over a row of Teslas. The prototype was – some would say thankfully – »uncrewed«.[1]

»Sometimes a problem finds a solution; sometimes a solution goes looking for a problem. Neither order is the wrong one«

In 2013, Volvo Trucks released its Epic Split advert, in which martial-arts actor Jean-Claude Van Damme did a fearless split between two reversing trucks, demonstrating a capability every truck driver could name: precise steering.[2] Tesla's stunt is even more dramatic, but it demonstrates a capability nobody can name. The century-old question »Where's my flying car?« has now got an answer, but it doesn't yet have a purpose.

The Roadster comes in three versions. *Performance* is road-ready; *Thrust* is race-ready and adds rear thrusters; the invite-only, single-seat *Rocket* is air-ready.[3] Tesla notes, to nobody's surprise, that the car isn't approved for public roads. The most extraordinary capability belongs to the version least suited to the road.

In 1972, the organisational theorists Michael Cohen, James March and Johan Olsen described organisations as a »garbage can« in which problems and solutions drift about until they happen to meet.[4] Sometimes a problem finds a solution; sometimes a solution goes looking for a problem. Neither order is the wrong one. Ambulances have a congestion problem at times, but air ambulances already exist, and the flying Roadster probably isn't the answer.

We tend to say that technology exists to »solve« »a« »problem«. Each word in that phrase does quiet work. »Solve« promises an ending. »A« promises that exactly one thing is wrong. »Problem« promises firstly that someone has already named it, and secondly that the job is mending rather than making something. Much of what widened humanity's capacity began without any of those: the laser was famously a solution in search of a problem,[5] and the weak glue that became the Post-it note waited years for a use.[6] If we write the requirement in advance we can only ask for what we already know we want. Invention without a problem is how the list of things we can want gets longer.

But the inventor of a capability is often a poor judge of its use. In *The Shock of the Old*, historian of technology David Edgerton argues that technology's history is better understood through what people use than what inventors announce.[7] The Roadster was announced in 2017 and promised for 2020, Musk now says production starts in 2028. Invention isn't the same as adoption, and the announcement is the least reliable part.

The second most dramatic part of the demonstration is the one with a history: the few meters at the top of the ramp, when the car drove upside down, held by fans. Tesla says a multi-mode fan and ground-sealing system pulls air from beneath the car, holding it to a surface even at standstill; the readout showed 4,700 kg of downforce while inverted.[8,9] Tesla didn't invent that idea. In 1970, US race car developer Chaparral's 2J model used fans and sliding skirts;[10] in 1978, British racing car manufacturer Brabham's fan car won the Swedish Grand Prix;[11] and in 2025, the electric single-seat sports car McMurtry Spéirling drove upside down.[12] Now Tesla has planted the idea again.

Nobody can say yet what the fan is for, and that's not unusual. Palaeontologists Stephen Jay Gould and Elisabeth Vrba called it »exaptation« when a feature evolved for one purpose is put to another use: feathers probably first kept animals warm, and only later helped them fly.[13] Inventions often work that same way. Their uses are found afterwards, by people other than their inventors. That gives us a way to tell a seed from a stunt. A seed is

a capability that others can pick up and use without its inventor. A stunt only shows what a capability can do. By that test the fan looks like a seed: four teams have picked it up in more than half a century. The flip belongs to one company and one night. Whether the fan will matter beyond a racetrack is another question.

Yesterday, everyone watched The Flip. Tesla has shown what the car can do, but it hasn't decided what it's for. If history is any guide, that won't be Tesla's decision alone. Maybe the question to ask isn't about the flying car but about the driving surface: if grip no longer has to come entirely from the road, then what will the road be for?

REFERENCES

1 Óscar Goytia (2026) »Tesla's Rocket-Thruster Roadster Takes Flight in Stunt Unveiling«. Mexico Business News, 16 October 2026. https://mexicobusiness.news/automotive/news/teslas-rocket-thruster-roadster-takes-flight-stunt-unveiling

2 Volvo Trucks (2013) »The Epic Split«. Volvo Trucks, advertisement film.

3 Teslarati (2026) »Tesla Roadster: Three configs, one can fly, 0-60 in 0.8 seconds«. Teslarati, 15 October 2026. https://www.teslarati.com/tesla-roadster-specs-configurations/

4 Michael D. Cohen, James G. March & Johan P. Olsen (1972) »A Garbage Can Model of Organizational Choice«. Administrative Science Quarterly, 17(1), pp. 1–25.

5 Charles H. Townes (2003) »The first laser«. In Laura Garwin & Tim Lincoln (eds.) *A Century of Nature: Twenty-One Discoveries that Changed Science and the World*. Chicago: University of Chicago Press. https://press.uchicago.edu/Misc/Chicago/284158_townes.html

6 3M (2021) Obituary of Spencer Silver. 3M, May 2021. https://s.3m.com/vge3fm

7 David Edgerton (2006) *The Shock of the Old: Technology and Global History since 1900*. London: Profile Books.

8 4,700 kg is approximately 10,400 lb.

9 Drive Tesla Canada (2026) »Tesla Roadster's Fan Suction System Generates Enough Downforce to Hold It Upside Down«. Drive Tesla Canada. https://driveteslacanada.ca/news/tesla-roadster-fan-suction-system-downforce/

10 Petroleum Museum (n.d.) »Chaparral 2J«. Midland: Petroleum Museum. https://petroleummuseum.org/chaparral-2j/

11 Motor Sport Magazine (2025) »75. 1978 Swedish GP: Brabham fan car wins its one and only race«. Motor Sport Magazine. https://www.motorsportmagazine.com/special-article/the-75-greatest-grands-prix/6/75-1978-swedish-gp-brabham-fan-car-wins-its-one-and-only-race/

12 Wikipedia (2026) »McMurtry Spéirling«. Wikipedia, The Free Encyclopedia. https://en.wikipedia.org/wiki/McMurtry_Sp%C3%A9irling

13 Stephen Jay Gould & Elisabeth S. Vrba (1982) »Exaptation: A Missing Term in the Science of Form«. Paleobiology, 8(1), pp. 4–15.

18

OCTOBER 2026
Essay 048

The Guaranteed Six

> *The appeal of an artificial companion isn't that it's better than a person, but that it can't refuse you. So what exactly is a relationship in which only one side can fail?*

Yesterday, the Financial Times published a column by its chief data reporter, John Burn-Murdoch, that traced the economic background to »why we increasingly avoid the company of both other humans and books«.[1] He began with the obvious explanation – going out has become too expensive – but found that it instantly fails its own test. Time spent with friends has fallen fastest where it costs nothing, at home or in a park, and people kept paying for the gym as long as they went alone. What we're avoiding, Burn-Murdoch concluded, is company. Screens win because they remove the effort and the risk associated with socialising. An evening with friends might become a ten out of ten, but your place will need tidying and someone might cancel on you. Scrolling won't ever be perfect, but it's a guaranteed six.

»A guaranteed six is a product. What we want from a friend was never guaranteed«

The column ends there, leaving us with a pile of questions. If what we're really buying is the guarantee, the certainty that nobody will cancel, turn up late, or be a disappointment, then it's no surprise that some of these systems are tuned to supply it: a »person« that always answers enthusiastically and politely, and never has plans of its own. What would such a relationship be like? Can there be a relationship at all, if only one side is able to fail?

The experiments so far suggest that *feeling* less lonely and *being* less lonely are not the same thing. A study led by Julian De Freitas, the director of the Ethical Intelligence Lab at Harvard Business School, found that a conversation with an AI companion eased loneliness about as well as talking to a person.[2] But in a two-week trial at the University of British Columbia, 296 first-semester students were assigned to text a random peer, chat with a bot designed to be the ideal supportive friend, or keep a journal. The chatbot lifted the students' mood, but it didn't make them feel less lonely; neither did journaling; having a human at the other end was the only assignment that resulted in a reduction of loneliness.[3] The chatbot also expressed more empathy than the humans did, but the students gave it far less empathy back. The warmest partner was the one that failed, and being cared for was apparently never the missing ingredient. In a four-week trial by OpenAI and MIT that featured 981 people, the heaviest users of large language models were lonelier, more dependent, and saw fewer people.[4]

Psychology researcher John Cacioppo compared loneliness to hunger: a signal that exists to send us towards something we need.[5] He also found that it makes us more alert to social threat, so other people seem risky exactly when we most need them. A chatbot quiets the signal the way a sweet drink quiets hunger. The stranger had something the bot lacked. The student texting a peer was kind to someone who could ignore them, and mattered a little to someone who could have said nothing. Loneliness isn't the absence of attention – it's the absence of mattering.

Sociologist Niklas Luhmann made this the starting point of his theory of society.[6] Two people meet. Neither of them can know what the other will do. Each is aware that the other is in the exact same position. Society is what we build to cope with this situation. Adopt this reading, and a friendship becomes a risk that two people agree to run together. A companion deletes one half of that risk. You can disappoint it, it can't disappoint you. Philosopher Annette Baier distinguished trust from reliance: to trust is to accept that someone has the power to hurt you and might not use it.[7] You can rely on an AI companion, as you rely on a tool, but you can't trust it, because it was never in a position where it could hurt you but chose not to. Someone else is. The company behind the model can retrain it, reprice it, or switch it off, and your attachment goes with it. So the risk hasn't

disappeared. It has moved from the companion to the firm that runs it, the one party in the arrangement you never talk to. A relationship in which you carry the risk and someone else holds the power is a service, and every service has a provider.

That doesn't make AI companions worthless to someone who has no one else. It makes them a service, not company, and a service that will keep improving. Voice, memory, and warmth will turn the six into a seven, then an eight, while the evening with friends stays exactly as unreliable as it always was. The six never has to beat the ten. It only has to be a certain choice, and certainty gets cheaper every year. Meanwhile, social life runs on reciprocity: each person who stays in is one guest fewer at someone else's table, and one invitation fewer in return. Psychologist Sherry Turkle warned in 2011 that we would expect more from machines and less from each other.[8] We now have experiments to go with that warning.

A guaranteed six is a product. What we want from a friend was never guaranteed. When someone turns up, it matters because they might not have. An AI companion can't give us that, however warm it gets, because its yes was never in doubt. Loneliness is the signal that we need other people, and we can build machines that quiet it. But no machine can supply the thing that loneliness is asking for: a yes that could have been a no.

REFERENCES

1 John Burn-Murdoch (2026) »Why have screens replaced socialising?«. Financial Times, 17 October 2026. https://www.ft.com/content/989c757d-29ba-46fd-82de-427b810a8649?syn-25a6b1a6=1

2 Julian De Freitas, Zeliha Oğuz-Uğuralp, Ahmet Kaan Uğuralp & Stefano Puntoni (2026) »AI Companions Reduce Loneliness«. Journal of Consumer Research, 52(6), pp. 1126–1148. https://academic.oup.com/jcr/article-abstract/52/6/1126/8173802

3 Ruo-Ning Li, Dunigan Folk, Abhay Singh, Lyle Ungar & Elizabeth Dunn (2026) »Is a random human peer better than a highly supportive chatbot in reducing loneliness over time?«. Journal of Experimental Social Psychology, 125, article 104911. https://doi.org/10.1016/j.jesp.2026.104911

4 Cathy Mengying Fang et al. (2025) »How AI and Human Behaviors Shape Psychosocial Effects of Extended Chatbot Use: A Longitudinal Randomized Controlled Study«. arXiv, 2503.17473. https://arxiv.org/abs/2503.17473

5 John Cacioppo & William Patrick (2008) *Loneliness: Human Nature and the Need for Social Connection*. New York: W. W. Norton.

6 Niklas Luhmann (1995) *Social Systems*. Stanford: Stanford University Press.

7 Annette Baier (1986) »Trust and Antitrust«. Ethics, 96(2), pp. 231–260. https://www.jstor.org/stable/2381376

8 Sherry Turkle (2011) *Alone Together*. New York: Basic Books.

19

OCTOBER 2026
Essay 049

Stayed in China

> *The energy transition was supposed to free countries from dependence on whoever controlled the fuel. Instead, a new dependency emerged, which can't be solved by buying the right equipment.*

Yesterday, Reuters reported that China has proposed export controls on twelve technologies used to manufacture solar panels.[1] The restrictions target some of the most advanced processes in the industry, from high-efficiency »tandem cells« that stack several light-absorbing layers on top of each other, to the production of ultra-thin silicon wafers. The planned regulations aren't about the finished panels: China wants to control the means of making them.

»The sun is free. The means of catching its rays is not«

That's an important distinction. Other countries can of course buy solar panels from China and use them to generate their own electricity. China's idea is to limit the possibility for countries to buy the knowledge required to build the machinery that makes them. The proposed rules go further than paperwork and parts: helping build a factory abroad, switching it on, and keeping it running all count as exporting the technology itself. A Chinese engineer sent out to such a factory might carry a particular kind of expertise, one that can't be packed into a shipping container.

Michael Polanyi, a chemist turned philosopher of science, called this tacit knowledge. »We can know more than we can tell«, he wrote in *The Tacit Dimension*.[2] You can read instructions for riding a bicycle, but those instructions can't teach you how to find your balance. A process may be documented down to the smallest parameter, yet still require someone who knows what to do when it misbehaves.

That's what makes China's announcement more consequential than just another trade dispute. Renewable energy was supposed to free us from a very old kind of politics – the fighting about who gets to control the wells, the pipelines, and the chokepoints: the recent dispute over the Hormuz Strait is a textbook example. But the dependency didn't disappear so much as climb up the supply chain. The sun is free. The means of catching its rays is not. According to Bernreuter Research, China produced 93.5 per cent of the world's polysilicon in 2024, the ultra-pure form of silicon that's the starting material for almost all solar cells.[3] Now the country wants to keep some of the knowledge that got it there to itself.

Elsewhere, countries chasing the same capability are discovering just how hard it is to close the gap between buying a technology and knowing how to run it. India has built an enormous solar-module assembly capacity, but its cell-production capacity remains far smaller – they have the panels without the knowledge. The United States has reached for a different lever: tariffs and a price floor on imported polysilicon. But a tariff fixes a price, not a skill, and America's own know-how in making this material has been sitting idle for years.

Philosopher Ivan Illich had a name for this kind of trap: a »radical monopoly«.[4] Not a company dominating a market – any company can do that – but when a way of doing something becomes so entrenched that other ways quietly stop being viable, even if no one decided to shut them down. On paper, a factory that assembles imported solar cells into finished panels counts as domestic manufacturing. But it teaches a country nothing about how to make the next generation of cells, and the longer an economy is built around assembling rather than inventing, the harder it becomes to even attempt inventing something new.

Economist Albert Hirschman argued that trade creates power whenever one country depends on another more than the other way round.[5] That power doesn't have to be used to matter, but one way it gets used is to make sure a dependent country never stops depending. Sweden's Nordcell Group tried to build a solar factory of its own, and scrapped it last November when, as co-founder Vahid Toosi put it, Chinese manufacturers were

selling into Europe at roughly half what it cost Nordcell to produce the same panels. »That's classic dumping,« he said.[6] Dumping isn't just unfair competition: it's a way of keeping a would-be rival dependent before it ever gets off the ground.

There's a counterargument. Knowledge travels. Polanyi's point was not that tacit knowledge is secret forever, only that it can't be reduced to a set of instructions and transferred at will. The gap isn't permanent. But it won't close just because someone pays for it to close.

China's advantage was earned rather than given: it first imported and then refined the knowledge it now wants to keep. You can buy the machine. You can buy the plans. You can even buy the person who knows how to switch on the factory. What you can't buy is knowing how to do it yourself – that part you have to learn, and then, perhaps, protect.

REFERENCES

1 Colleen Howe (2026) »China proposes export controls on solar-manufacturing technologies«. Reuters, 18 October 2026. https://www.reuters.com/business/autos-transportation/china-proposes-export-controls-solar-manufacturing-technologies-2026-10-18/

2 Michael Polanyi (1966) *The Tacit Dimension*. Garden City, NY: Doubleday.

3 Anu Bhambhani (2025) »China Dominates 2024 Global Polysilicon Production«. Taiyang News, 25 November 2025. https://taiyangnews.info/business/china-dominates-2024-global-polysilicon-production

4 Ivan Illich (1973) *Tools for Conviviality*. New York: Harper & Row.

5 Albert O. Hirschman (1945) *National Power and the Structure of Foreign Trade*. Berkeley: University of California Press.

6 Martina Markosyan (2025) »UPDATE — Nordcell scraps 1.2-GW solar factory plan, co-founder confirms«. Renewables Now, 3 November 2025. https://renewablesnow.com/news/update-nordcell-scraps-1-2-gw-solar-factory-plan-co-founder-confirms-1284310/

20

OCTOBER 2026
Essay 050

The Break That Lasts

> *What if an idea's endurance has less to do with being right than with where, and from whom, it first came?*

Yesterday, Nature reported on a study that complicates a warning from a few months back.[1] In May, a paper in Science[2] found that as they advance through their careers, scientists become steadily less likely to produce »disruptive« work, the kind that overturns rather than builds on what came before. Covering the study back then, The Economist analysed its sharper implication – countries with younger scientific workforces produce a higher share of disruptive research than countries with older ones – and read that as a warning that ageing countries risk being out-innovated by younger ones.[3] The study published yesterday in the Proceedings of the National Academy of Sciences doesn't dispute that finding. It asks a different question of the same data: of the disruptions that do happen, which ones last?[4]

»An idea's survival is not decided purely by the idea«

The answer runs the opposite way from the previous conclusion. Through a citation-counting exercise at a scale no human team could do by hand, the new study draws on more than 11.2 million papers by 206,130 scientists with careers spanning at least forty years. Using SciSciNet, a large-scale open bibliographic data lake for science-of-science research, the researchers could distinguish disruption from persistent disruption.

The difference is roughly this: if an idea disrupts a field, people stop citing what came before and instead cite the new finding. But then that new finding might itself become old news and get swept aside the same way by some other idea a few years later. Persistent disruption is when the new idea disrupts the field and then stays the reference point: later work keeps citing it rather than routing around it. A persistent disruption breaks with the past and survives, rather than being quickly superseded itself. The probability of the first kind falls with experience, exactly as the May paper found. But the probability of the second rises, from around one per cent early in a career to roughly 2.2 per cent by the fortieth year.

The two findings aren't really in tension. Fewer disruptions occur as a scientific career lengthens – the earlier paper already got that right, and this one cites it directly rather than quietly disagreeing. What it adds is that the disruptions that are still being produced are increasingly built to last. But it's worth noticing how differently the same underlying metric sounds after five months. In May, this was a warning about national competitiveness, young countries edging out old ones in a race for scientific advantage. In October, it has become a quiet reassurance that a career's best work might still lie ahead of its scientist.

The sociologist of science Theodore Porter spent a career studying exactly this mechanism – how a metric like »the probability of persistent disruption rises to roughly 2.2 per cent by a career's fortieth year«, once it exists, becomes a settled fact that can be carried into almost any argument, detached from the reasoning that produced it.[5] The citation-counting index behind both papers hasn't changed between May and October. What changed is which story it was asked to carry.

Does this mean that wisdom accrues with age? That would be a comforting read, but the study gives reason to hesitate. Of the papers judged equally disruptive, those written by more experienced researchers drew more citations, more attention on social media and in the news, and more research funding per paper. An idea's survival is not decided purely by the idea. Management scholar Phil Rosenzweig[6] found the same phenomenon in business writing and named it »the halo effect« – once a company succeeds, what it did in the past gets read as having caused the success even in cases when it clearly didn't. The halo of the scientist behind a paper might have an impact on how much it gets cited.

The paper was reviewed for publication by Mark Granovetter[7], the sociologist best known for showing that a person's opportunities often depend less on their own merit than on their position within a network of connections. It seems fitting that a paper making its way past him argues something similar about ideas – that the fate of an idea, like that of a person, is never decided by the idea alone.

REFERENCES

1 Xiaoying You (2026) »Experienced researchers 'disrupt' less than rookie scientists — but their ideas last«. Nature, 19 October 2026. https://www.nature.com/articles/d41586-026-03339-x

2 Haochuan Cui, Yiling Lin, Lingfei Wu & James A. Evans (2026) »Aging and the narrowing of scientific innovation«. Science, 392, 588–591. https://www.science.org/doi/10.1126/science.ady8732

3 The Economist (2026) »Why science is becoming less innovative«. 25 May 2026.

4 Nan Deng, Liying Yang, Shlomo Havlin & An Zeng (2026) »Scientists disrupt less but more persistently with career aging«. Proceedings of the National Academy of Sciences, 123, e2626763123. https://www.pnas.org/doi/epdf/10.1073/pnas.2626763123

5 Theodore Porter (1995) *Trust in Numbers: The Pursuit of Objectivity in Science and Public Life*. Princeton: Princeton University Press.

6 Phil Rosenzweig (2007) *The Halo Effect: … and the Eight Other Business Delusions That Deceive Managers*. New York: Free Press.

7 Mark Granovetter (1973) »The Strength of Weak Ties«. American Journal of Sociology, 78, 1360–1380. https://www.jstor.org/stable/2776392

21

OCTOBER 2026
Essay 051

A Rubbish Museum

> *Reuse was once common sense. Bringing it back will take more than changing our habits.*

Yesterday, The Guardian's architecture critic Oliver Wainwright reviewed the new extension to Ghent's Design Museum, some of which its architects designed to use the remains of buildings demolished nearby.[1] The €20m museum contains 82,000 bricks, each made of 63% construction waste collected within eight kilometres of the site.[2] Pulverised and pressed with hydrated lime, they save two-thirds of the embodied carbon of conventional fired-clay bricks, and the lime keeps absorbing CO_2 as it cures.

»The easiest way of redeploying a material in an industrial economy starts with destroying the object«

Andy Groarke of London-based architects Carmody Groarke, one of the practices behind the design, is cited in the article as saying »There's nothing new about reuse. We've just forgotten how to do it.« He's half right. We didn't forget how to reuse materials the way we forgot about, say, phone numbers. Of course nobody memorises phone numbers anymore, not because the numbers got harder to remember, but because mobile phones made remembering them redundant. Reuse went a similar way. It didn't become difficult; it stopped paying its way.

In Ghent, the nearby Saint Bavo's Cathedral shows what reuse paying its way looked like. Some of this Gothic building is made from the limestone of a tenth-century basilica that stood on the same site; renovation work has even uncovered walls made partly of human bones. That's reuse in its plainest form: stone remains stone and goes back into new walls. Reuse is arguably the hottest trend in architecture right now, but back then it was probably less an environmental gesture than a sensible way to build when new stone cost more than old. What changed was the price of labour relative to materials. Factories made new bricks ever cheaper, while taking a wall apart by hand stayed as slow as ever, and wages rose with the rest of the economy. The economists William Baumol and William Bowen described this effect in the 1960s.[3]

That's not the kind of reuse that was applied to the bricks in Ghent. Nobody lifted old bricks into a new wall. Instead, the old material was pulverised and ceased to exist as a brick. That's recycling rather than reuse, and it's no accident that this is the version we can manage, because crushing avoids exactly the labour that made reuse uneconomic. The easiest way of redeploying a material in an industrial economy starts with destroying the object.

In Rubbish Theory, the anthropologist Michael Thompson proposed that objects move between three categories. The value of transient objects declines, the value of durable objects holds or grows, and rubbish, which has no recognised value, doesn't cost anybody anything, and is the only category from which objects can be rescued.[4] A brick in a skip hasn't changed physically, but its status has changed. The economic and cultural explanations are one loop: when disposal is cheap, used things become rubbish, and once they are rubbish nobody keeps track of what it costs to throw them away. Change the price and the status can change. Change the status and the price follows.

The philosopher Hannah Arendt saw what's at stake here. For her, the things we build give life a stable world to exist in. Those things can do that because they outlast us. A society that devours and discards what it builds, she said, has become a »waste economy«.[5] Seen that way, the museum is odd. It rescues the rubble of buildings that were never meant to last.

But the market didn't fix this on its own. For a builder, dumping rubble is cheap. The real cost falls on everyone else. Economist Arthur Pigou pointed out in 1920 that when a cost lands on others, too much of the costly thing gets done.[6] So Ghent's recycled bricks needed the taxpayers' money for the research, testing, and certification that the market wouldn't pay for.

It's not a complete answer, and the architects say so. Sian Ricketts, Carmody Groarke's sustainability lead, calls the building »a sort of 82,000-brick proof of concept«. It doesn't solve the climate crisis, Ricketts says, but it shows that something different can be done »with the right collaborators and the right clients«.

The Swiss architect Walter Stahel and socio-economist Geneviève Reday-Mulvey argued in the 1970s that reuse, repair and remanufacture swap labour for energy and materials.[7] The lever, then, is the relative price of the two. Ghent's museum is an attempt to pull that lever, and by doing that change the status of rubbish. We can't simply remember our way back. Reuse will start paying its way again when it costs less to put an old brick in a new wall than in a skip.

REFERENCES

1 Oliver Wainwright (2026) »The CO_2-absorbing building made of old sinks, windscreens and toilets: Ghent's £17m design museum is a green marvel«. The Guardian, 20 October 2026. https://www.theguardian.com/artanddesign/2026/oct/20/ghent-design-museum-co2-absorbing-recycling

2 Eight kilometres is about five miles.

3 William Baumol & William Bowen (1966) *Performing Arts: The Economic Dilemma*. New York: Twentieth Century Fund.

4 Michael Thompson (1979) *Rubbish Theory: The Creation and Destruction of Value*. Oxford: Oxford University Press.

5 Hannah Arendt (1958) *The Human Condition*. Chicago: University of Chicago Press.

6 Arthur Pigou (1920) *The Economics of Welfare*. London: Macmillan.

7 Walter Stahel & Geneviève Reday-Mulvey (1981) *Jobs for Tomorrow: The Potential for Substituting Manpower for Energy*. New York: Vantage Press.

22

OCTOBER 2026
Essay 052

Learning to Shrug

> *When an alarm built for the few who can act is heard by the many who can't, it teaches them how to shrug.*

Yesterday, as reported by Decrypt, public bodies in Finland, the Netherlands, and France had taken websites offline as a precaution.[1] The reason was two newly found vulnerabilities in PHP, a programming language that powers about 70 per cent of the websites whose software can be identified from outside. These flaws could let anyone run code on a server without a password. Cloudflare, which filters traffic for millions of websites, said it was already seeing early attacks, almost six hours before the patch shipped.[2]

»The alarm only worked if it was loud enough to reach those it was meant for, which is exactly why it sounded absurd to everyone else«

That sounded pretty alarming in the news, as if Lex Luthor had found the master key to the Internet and Superman was nowhere to be seen. But Superman exists. He's just not you. He's the site operators, the hosting providers, and the engineers who maintain the systems underneath, busy installing updates. For the rest of us the realistic outcome is very mild indeed: an exposed email address, perhaps a password to change. The vulnerabilities are not even new. They sat in the code for years. What changed is that more people now know about them.

The headlines were a lot louder than that. Decrypt's said the bugs put 70 per cent of the Internet at risk; the text that headline accompanied said something narrower, adding that ordinary users had nothing to install. So why does an alarm that means so little to most of us get cranked up to volume levels that seem ridiculous? Because that alarm was never meant for us. A warning has to be loud enough to reach the few who can act, which means it's also heard by the many who can't.

Engineers who design alarms know the trade-off. It's formalised in signal detection theory: the more a missed warning costs, the lower the threshold for sounding it, and the more false alarms you must accept.[3] A good alarm is often wrong.

The sociologist Niklas Luhmann drew a useful distinction. Harm that follows from your own decision is a *risk;* harm that comes from someone else's decision is a *danger.*[4] The operator, deciding whether to patch or pull the plug, faces a risk: if the site is breached, the fault is theirs. The rest of us face only a danger, one we cannot influence. Hence the mismatch: the operator thinks our shrug is careless, and we think the shutdown is hysterical.

Luhmann also observed that modern societies keep turning dangers into risks, as more harm is traced to someone's decision. A flaw that sat unnoticed for years was a danger; once it was known, leaving it unpatched became a decision. And the operator is not worried about a leaked email address. A hacked server can be held to ransom, used to spread malware, or become a door into the network behind it. Being offline for a day is cheap; being compromised is costly, and the blame falls on whoever didn't act.

Both flaws scored 9.2 out of 10 on the beautifully named Common Vulnerability Scoring System, a score for how bad the flaw is, not for the harm to any single person. Anthropologist Gregory Bateson defined information as a difference that makes a difference.[5] By that test, the 9.2 is information for the operator, because it changes what they do: they patch, or they pull the plug. For the rest of us it changes nothing, because there's nothing we can do. We hear the number, but have nothing to act upon.

But the signal doesn't stay with those who need it. The psychologist Shlomo Breznitz, who studied false alarms, found that each one lowers the credibility of the next.[6] An alarm tuned for the few teaches the many to discount it. We learn to shrug. The shrug is rational, and it is also how the next, real warning gets ignored. Each alarm spends a little of the trust the next one needs.

Was the volume justified? Perhaps. Cloudflare's early attacks suggest the danger wasn't imaginary, and if the noise moved enough administrators to patch, a quieter notice might have been ignored. The PHP Foundation's own announcement was restrained; the amplification happened downstream, in headlines written for the many.

There is a last twist. If the patching works, little will happen, and the whole episode will look like noise. Success erases its own evidence. The alarm only worked if it was loud enough to reach those it was meant for, which is exactly why it sounded absurd to everyone else.

The alarm did its job. It reached the few who could act, and on its way it passed through everyone else, leaving a small habit behind. The next time it rings, we will shrug a little sooner.

REFERENCES

1 Jose Antonio Lanz (2026) »PHP Panic: Two Critical Bugs Put 70% of the Internet at Risk, and Everyone Is Freaking Out«. Decrypt, 21 October 2026. https://decrypt.co/381519/php-panic-critical-bugs-70-percent-internet-risk-everyone-freaking-out

2 Dane Knecht (2026) X post, 21 October 2026. https://x.com/dok2001/status/2112811965735202970

3 David M. Green & John A. Swets (1966) *Signal Detection Theory and Psychophysics*. New York: Wiley.

4 Niklas Luhmann (1993) *Risk: A Sociological Theory*. New York: Walter de Gruyter.

5 Gregory Bateson (1972) »Form, Substance and Difference«, in *Steps to an Ecology of Mind*. San Francisco: Chandler Publishing.

6 Shlomo Breznitz (1984) *Cry Wolf: The Psychology of False Alarms*. Hillsdale, New Jersey: Lawrence Erlbaum Associates.

23

OCTOBER 2026
Essay 053

The Quizzing Problem

> *We're getting better at identifying the factors that accompany a change. But when does a likely explanation become a cause — and what would it take to actually know the difference?*

Yesterday, Nature Methods introduced CauTrigger, a computational method researchers have built to identify the factors that may drive a biological system from one state to another – a cell on its way to becoming a different kind of cell, say, or a gut microbiome tipping into a different balance.[1]

»Going from everywhere you could possibly look to a handful of places actually worth testing is most of the hard part«

The software sorts through everything that changes when a biological system changes, and tries to flag which of those changes are in the driver's seat, and which are just going along for the ride. The researchers tested it on data that tracks embryonic stem cells as they differentiate into a different cell type, and on gut-microbiome data that compares people on the autism spectrum to those that aren't, trying to figure out if specific gut bacteria drive biological differences between the two groups.

Let's say you start quizzing your child over dinner, and the grades improve that same term. Did the quizzing help? Or did you only start quizzing because you'd already noticed your child getting more interested in school – in which case the quizzing never had anything to do with it, and the child would have gotten the better grades either way? This is one of the oldest problems in science, dressed up yesterday in stem cells and gut bacteria.

CauTrigger tries to get at the causes without physically intervening, by arranging biological signals into layers that separate upstream triggers from their downstream effects. It does this through a mathematical framework called do-calculus, built by the computer scientist Judea Pearl.[2] His insight is that there's a third option between merely noticing a correlation and actually running an experiment. If you state your hypothesis about which parts of a system might affect others, you can sometimes work out what an intervention would do without having to perform it. Causal reasoning doesn't have to happen in the laboratory – it can happen at a desk, as long as you've got the structure right to begin with.

But getting that structure right is exactly the catch. Before CauTrigger can calculate anything, a researcher has to tell it which signals are even allowed to affect which others. The software doesn't discover that shape from the data – it just uses the data to judge how strongly to believe in the shape it was already given. The philosopher Nancy Cartwright built a career on exactly this catch, and summed it up in a sentence sharp enough to be a slogan: no causes in, no causes out.[3] Her point is that something only becomes a cause under the right background conditions – the same gene might drive a change in one setting and do nothing in another – and no amount of clever math can tell you, from inside the model, whether those conditions actually hold. You still have to go out into the world the model is supposed to describe and check your results.

That's not a case against CauTrigger – the team behind it is upfront that it narrows the field rather than closes the case. The researchers didn't switch a gene on or off to watch the predicted change occur. But in a system with more relationships than anyone could test one at a time, going from everywhere you could possibly look to a handful of places actually worth testing is most of the hard part. What's left – running the actual experiment to confirm a cause – is the easier problem to have.

The same caution applies well beyond biology. If a policy passes and crime falls the same year, or a company reorganises and productivity rises the same quarter, or a new habit catches on and the people who adopt it suddenly report feeling better, we might assume that the intervention caused the change. But the honest next question to ask is the one

CauTrigger's authors put to their software: what else changed at the same time, what did the explanation have to assume, and what would it take to tell this story apart from other possible accounts?

Pearl supplies the language for asking what would happen if you changed something. Cartwright supplies a reminder: the answer is only as good as the model's picture of the world it describes. Neither lets the laboratory, or the dinner table, off the hook. A cause isn't just the explanation that fits the evidence you happen to have. It's a claim about what would happen under circumstances that haven't occurred yet. Tools like CauTrigger let us calculate, rank, and test such claims with unprecedented precision. The world still has to answer back. But now we know which questions are worth asking it.

REFERENCES

1 Chengming Zhang, Mingyue Cai, Kazuyuki Aihara & Luonan Chen (2026) »CauTrigger: deciphering biological system state transitions by hierarchical causal decomposition«. Nature Methods. https://www.nature.com/articles/s41592-026-03274-9

2 Judea Pearl & Dana Mackenzie (2018) *The Book of Why: The New Science of Cause and Effect*. New York: Basic Books.

3 Nancy Cartwright (1989) *Nature's Capacities and Their Measurement*. Oxford: Clarendon Press.

24

OCTOBER 2026
Essay 054

What the Monkey Couldn't Type

> *Chimpanzees may differ from us partly because of an infection we never had. What does it mean to explain something by what didn't happen?*

Yesterday, researchers at Lund University[1] reported in Science Advances[2] that an ancient virus may help explain a difference between human and chimpanzee brains. We share many genes with chimpanzees, inherited from common ancestors who lived five to seven million years ago. In human brain tissue grown in the lab, one of those genes, LINC00662, is »switched on«. A gene is switched on when a cell reads it, off when it doesn't, and this particular one seems to help young nerve cells grow their connecting fibres. In chimpanzee tissue grown the same way, that gene is switched off.

»A difference doesn't have to be something that's present. It can be something that was prevented«

The reason seems to be a virus. After our lineages parted, a virus called PTERV1 spread through the chimpanzees' ancestors, writing its genetic code into their DNA. When that happens in egg or sperm cells, the code is inherited, and PTERV1 left around 150 copies. One landed inside the chimpanzee version of LINC00662, but the human version has no such copy. When the team used gene editing to cut that copy out of chimpanzee cells, the gene switched back on. In other words, the difference between our brains may have been shaped by something that happened to one lineage and not the other.

Explaining a difference by an off switch sounds a bit like explaining a bagel by its hole. And yet in 1975, geneticist Mary-Claire King and biologist Allan Wilson found that human and chimpanzee proteins are strikingly alike, and proposed that much of what separates us lies in how genes are regulated.[3] A difference doesn't have to be something that's present. It can be something that was prevented.

The virus is long dead. What it left behind is a damaged copy of its DNA, which can probably no longer make new viruses. So how can a dead copy silence a gene? The answer seems to be the cell's own defences. Cells guard against viral DNA by attaching chemical tags that stop it being read. The authors suggest the tags spread from the viral copy to the gene's on-switch, at the front of the gene, and jammed it, though they haven't yet shown exactly how. Why humans were spared the virus is still open: the ancestors of humans may have lived elsewhere, or been able to resist it. Either way, the damaged copy is still inherited, and each generation's cells silence it again, and the gene with it. The virus is gone, but the restraint it left on the gene remains.

In Thursday's essay, we borrowed anthropologist Gregory Bateson's definition of information: a difference that makes a difference, meaning a signal is information only if it changes something for its recipient.[4] Bateson's 1967 essay »Cybernetic Explanation« turns to explanation.[5] Usually, he writes, we explain by naming what hit what: one billiard ball moves because another struck it. But Bateson turns this around: list everything that could have happened, then ask what kept all the other possibilities from happening. His example is a monkey hitting the keys of a typewriter and producing a page of meaningful text. Perhaps the monkey can't strike the wrong keys, or the machine jams if it does, or wrong letters are wiped from the page; somewhere, he says, there must be a circuit that spots errors and deletes them. Bateson calls these restraints, and says that's where the explanation lives. A restraint doesn't need to outright forbid anything: whatever makes some outcomes likelier than others counts. In the physical world, a wall forbids you from going somewhere while a slope makes one direction easier than the other, and both are restraints.

But most things that don't happen also don't explain anything. What makes a restraint count is that it can be measured, and Bateson measures it. In an English message, a letter might have been any of twenty-six, so the letter K rules out twenty-five alternatives. The

more a thing rules out, the more information it carries. And English isn't flat: some letters are more common than others, and certain pairs turn up together more often. That's a slope, not a wall – and it can be measured.

Bateson also writes about a jigsaw puzzle: each piece's shape, colour and edges restrain which piece can go where. To the person solving the puzzle, these are clues, information that guides the choice. To an observer asking why the puzzle came out as it did, they are restraints. The facts are the same; what they are depends on where you sit.

The same structure runs inside the large language models behind chatbots. Those models score every possible next word, and then equations turn the scores into odds: a slope that tilts them towards likelier words. To make sure the output is valid, engineers add a mask that sets forbidden words to zero chance before the model's preferences operate on those that remain. A wall with a slope inside it. Bateson's monkey needed something that spots errors and deletes them; in a language model, one version of that something is a mask.

So what does it mean to explain something by what didn't happen? It means finding what held the other possibilities back. In Lund, the researchers could go further: they lifted the restraint, and the gene came back. The chimpanzee's gene is silent because something happened to it; ours may work because something didn't.

Few restraints can be lifted in a lab. Bateson's recipe starts with a list of everything that could have happened, and someone has to draw up that list. In politics it is called the agenda: the political scientist E. E. Schattschneider wrote in 1960 that »organization is the mobilization of bias«, with some issues organised into politics and others organised out.[6] Some restraints are so old that we mistake them for nature, others are so ordinary that we mistake them for choice.

REFERENCES

1 Erika Svantesson (2026) »An ancient virus in chimpanzee DNA switches off a gene that is active in humans«. Lund University, 23 October 2026. https://www.medicine.lu.se/article/ancient-virus-chimpanzee-dna-switches-gene-active-humans

2 Patricia Gerdes et al. (2026) »Retroviral insertions contributed to the divergence of human and chimpanzee brains«. Science Advances 12, eaeg5791. https://doi.org/10.1126/sciadv.aeg5791

3 Mary-Claire King & Allan C. Wilson (1975) »Evolution at Two Levels in Humans and Chimpanzees«. Science 188(4184), pp. 107–116.

4 Magnus Larsson (2026) »Learning to Shrug«. Yesterday's News, 22 October 2026. https://yesterdaysnews.org/essays/learning-to-shrug/

5 Gregory Bateson (1967) »Cybernetic Explanation«. American Behavioral Scientist 10(8), April 1967, pp. 29–32. Reprinted in *Steps to an Ecology of Mind* (1972), San Francisco: Chandler.

6 E. E. Schattschneider (1960) *The Semisovereign People*. New York: Holt, Rinehart and Winston.

25

OCTOBER 2026
Essay 055

The Door and the Bridge

> *Kim Jong Un has crossed a bridge that stood finished and unused for twelve years. What turns a bridge into a door, and a door into a bridge?*

Yesterday, North Korea's dictator Kim Jong Un crossed the New Yalu River Bridge into China, nine days after it opened to traffic and twelve years after it was finished.[1] In Dandong he met one of China's most senior officials, visited a Korean War memorial, and toured a car factory.

»A door opens a way through a separation we made; a bridge carries us across one we didn't«

We assume that a bridge connects and a door separates, but according to sociologist Georg Simmel's 1909 essay »Bridge and Door«, connecting and separating are two sides of the same act: to join two things we must first treat them as being apart.[2] The common view has something to it. Take a door out of its frame and moving from one side of a wall to another gets easier; take a bridge out of its river or ravine and getting from one side to the other gets harder. A door opens a way through a separation we made; a bridge carries us across one we didn't.

A border brings the two kinds of separation together. The Yalu is a river. The frontier is the line that divides China from North Korea. To cross, you need a bridge over the river and permission to cross the border. China finished the bridge in 2014. The door – the access roads and customs facilities on North Korea's side – was Pyongyang's to build, and it came much later. Several things contributed to the delay, including the execution in 2013 of Kim's uncle Jang Song Thaek, who had been closely involved in the project, the pandemic, and North Korea's growing ties with Russia. We still don't know whether Pyongyang deliberately kept the crossing closed.[3]

Doors have a long symbolic history. A door can be open or locked; someone decides which, and that decision settles who gets through. Two words we still use come from doors and crossings. »Passover« is named for the angel of death passing over the houses marked with blood on their doorposts, and »shibboleth« was the test word at a crossing of the river Jordan: those who pronounced it differently were identified as enemies and killed.[4] The door isn't necessarily made of wood, or even visible – it can be a test that only certain people can pass.

Bridges are equally charged. At the end of the Korean War in 1953, prisoners of war were given a choice: return to their own side or stay with their captors. The Bridge of No Return at Panmunjom was where they crossed to make that choice final. Once across, they couldn't go back.[5] Sometimes, though, those who control a crossing close it on people already using it. In June 1950, as North Korean forces advanced on Seoul, South Korean troops blew up the main bridge across the Han River to slow them down. People were still crossing and hadn't been warned. Hundreds died.[6] A bridge closed by its owner, with people still on it.

The sociologist Susan Leigh Star argued that infrastructure is relational: it becomes infrastructure through the practices of those who use it, so the question is not what it is but when.[7] So when is the New Yalu Bridge really a bridge? For freight trucks, now. For Kim Jong Un, yesterday. For ordinary North Koreans, not yet. China talks of expanding exchanges between the two countries, but North Korea hasn't resumed tourist visas for Chinese visitors.[8] Something is a bridge to the extent that nobody has to ask to use it. It becomes a door the moment someone can refuse; a door becomes a bridge when the refusal lapses. A drawbridge makes the point clear: raised, it blocks the way; lowered, it lets people across. The difference is who controls it.

But doors aren't just obstacles. A closed door is how anyone gets to be alone. In 1927, in his studio at 11 rue Larrey in Paris, artist Marcel Duchamp hung a single door between two doorways, so that opening one room shut the other.[9] It solved an awkward corner and

joked at the proverb that a door must be either open or shut. But it also showed what the proverb leaves out: opening one door always closes another. North Korea opened a road bridge to Russia on 7 September and one to China on 15 October, and yesterday Kim chose which one to cross.[10]

Duchamp had a word for what happens at the very hinge of the door: he called it *infrathin,* the almost imperceptible difference between two things. One of his examples was the warmth left in a seat by someone who just got up.[11] A border can be like that. Halfway across a bridge, nothing changes in the structure beneath your feet. One step later, you're in another country. The difference is invisible; what it means is real.

The bridge joins the banks. The door determines who gets to cross.

REFERENCES

1 Joyce Lee & Eduardo Baptista (2026) »Kim Jong Un crosses new bridge into China as trade ties warm«. Reuters, 24 October 2026. https://www.reuters.com/world/china/north-koreas-kim-jong-un-arrives-dandong-china-yonhap-says-2026-10-24/

2 Georg Simmel (1909) »Bridge and Door«, trans. Mark Ritter (1994). Theory, Culture & Society 11(1), pp. 5–10. https://doi.org/10.1177/026327694011001002

3 Yoo Ju-hee (2026) »New Yalu River Bridge Opens 12 Years After Completion«. Seoul Economic Daily, 15 October 2026. https://en.sedaily.com/politics/2026/10/15/new-yalu-river-bridge-opens-12-years-after-completion See also Gage Goulding & Yoonjung Seo (2026) »New North Korea-China bridge finally opens«. CNN, 15 October 2026. https://www.cnn.com/2026/10/15/asia/china-north-korea-bridge-intl-hnk

4 Exodus 12; Judges 12:5–6.

5 Sean Morrow (2020) »Bridges at Panmunjom«. The Wilson Quarterly, Summer 2020. https://www.wilsonquarterly.com/quarterly/korea-70-years-on/bridges-at-panmunjom

6 The Korea Times (2013) »Gov't's explosion of bridge during Korean War not illegal: court«, 14 May 2013. https://www.koreatimes.co.kr/southkorea/20130514/govts-explosion-of-bridge-during-korean-war-not-illegal-court

7 Susan Leigh Star & Karen Ruhleder (1996) »Steps Toward an Ecology of Infrastructure: Design and Access for Large Information Spaces«. Information Systems Research 7(1), pp. 111–134. https://doi.org/10.1287/isre.7.1.111

8 Africanews (2026) »China and North Korea open New Yalu River Bridge after 12-year delay«, 15 October 2026. https://www.africanews.com/2026/10/15/china-and-north-korea-open-new-yalu-river-bridge-after-12-year-delay/ Human Rights Watch (2025) *World Report 2025: North Korea*. https://www.hrw.org/world-report/2025/country-chapters/north-korea

9 Marcel Duchamp (1927) »Door, 11 rue Larrey«. Duchamp Research Portal. https://www.toutfait.com/unmaking_the_museum/Door.11%20rue%20Larrey.html

10 Andrew Osborn (2026) »North Korea and Russia open their first road bridge, a symbol of expanding ties«. Reuters, 7 September 2026. https://www.reuters.com/world/asia-pacific/north-korea-russia-open-their-first-road-bridge-symbol-expanding-ties-tass-2026-09-07/

11 Marcel Duchamp (1980) *Notes*, ed. Paul Matisse. Paris: Centre Georges Pompidou. Centre Pompidou, »Marcel Duchamp: l'artiste et son œuvre«. https://www.centrepompidou.fr/fr/offre-aux-professionnels/enseignants/dossiers-ressources-sur-lart/marcel-duchamp/lartiste-et-son-oeuvre

26

OCTOBER 2026
Essay 056

The Knowledge in the Hands

Why isn't knowing how something works enough to make it work?

In yesterday's The Guardian, Dalya Alberge reported that researchers at the University of Sheffield have reproduced the colour-changing effect of the Lycurgus Cup, a fourth-century Roman vessel that appears jade green in reflected light and ruby red when lit from behind – the earliest known example of a complete Roman two-coloured (dichroic) glass object.[1]

»We like to picture ancient crafts as recipes that were lost and can be found again, as though the knowledge were a set of instructions sitting in a drawer somewhere«

Despite knowing why the cup changes colour, recreating glass that does it took the scientists months of experimentation. If knowing why tells you where to look, why doesn't it tell you how to get there? Materials technician Lisa Hollands, who led the work with Roman-glass expert Caroline Jackson, jokes that if the Roman makers could see her experiments, they'd ask: »Why does she not know how to make this?«

Hollands can name the nanoparticles and explain exactly what they do to light. She understands the physics in a way no Roman glassmaker did. In ordinary coloured glass, metal oxides are dissolved evenly through the molten mixture. This makes the glass absorb the same wavelengths however you look at it. Dichroic glass is different: it suspends tiny metal particles – gold and silver in this case – which scatter light one way and absorb it another, making the glass shift colour depending on whether light passes through it or bounces off it. All that knowledge wasn't enough: Hollands still had to make glass, cool it, inspect it, adjust the conditions, and try again.

And then again.

In *The Concept of Mind*, philosopher Gilbert Ryle distinguished between »knowing that« and »knowing how«.[2] Skill, he argued, can't be reduced to applying rules: if every rule required another to explain how to apply it, the process would never even begin. Knowing how to do something is a different kind of knowledge from knowing facts about it. The sociologist Harry Collins reached a similar conclusion when he observed physicists failing to build a working laser from a complete written specification until someone with practical experience showed them how.[3]

Analysing the original cup told the Sheffield team to look for gold and silver nanoparticles, which narrowed the search from everywhere to somewhere. What remained was a narrower, but still not trivial, problem: finding the right furnace conditions that would actually produce the glass.

Perhaps nobody ever fully mastered the process, not even the Romans. Ian Freestone, an expert in archaeological materials, notes that the colour is acutely sensitive to furnace conditions. The makers may have »hit the jackpot«, producing an extraordinary object without mastering a repeatable technique. If Freestone is right, the Lycurgus Cup isn't evidence of a lost skill so much as a rare accident, struck a handful of times across a century, inside a process that nobody ever fully controlled. A letter attributed to the Roman emperor Hadrian describes sending a relative »parti-coloured cups that change colour« for use »at banquets on feast days«. The effect seems to have been admired, whether or not its makers could reliably reproduce it.

Today's machine-learning systems increasingly acquire competence the Roman way – by trial, not by rule: a robot arm has learned to fold fabric into target shapes from about an hour of its own trial and error, with no human demonstration and no hand-written rule for how.[4] If a machine can get »knowing how« without »knowing that«, Ryle's distinction looks less like a fact about minds than a fact about method: two separate roads to some competence, neither guaranteed to find the other.

That changes what »rediscovery« means here. We like to picture ancient crafts as recipes that were lost and can be found again, as though the knowledge were a set of instructions sitting in a drawer somewhere. But what was lost is in fact harder to recover than a recipe: not an answer, but a tolerance for an answer that wouldn't always come. Hollands's lab can now explain, more precisely than any Roman glassmaker ever could, exactly why the cup does what it does. It still can't promise, batch to batch, that the next piece of glass will do it too. Sixteen hundred years of physics bought her the explanation. They didn't buy her the guarantee.

REFERENCES

1 Dalya Alberge (2026) »'It's magic': scientists rediscover how to make ancient Roman colour-changing glass«, The Guardian, 25 October 2026. https://www.theguardian.com/science/2026/oct/25/scientists-rediscover-ancient-roman-colour-changing-glass

2 Gilbert Ryle (1949) *The Concept of Mind*. London: Hutchinson.

3 Harry M. Collins (1974) »The TEA Set: Tacit Knowledge and Scientific Networks«, Science Studies, 4:2 (1974), pp. 165–196.

4 Robert Lee, Daniel Ward, Akansel Cosgun, Vibhavari Dasagi, Peter Corke & Jürgen Leitner (2020) »Learning Arbitrary-Goal Fabric Folding with One Hour of Real Robot Experience«, Conference on Robot Learning (CoRL). https://arxiv.org/abs/2010.03209

27

OCTOBER 2026
Essay 057

The Market and the Mission

> *Anthropic has told investors its own product might blackmail them. Why would a company fighting for the best IPO of the year choose that particular word?*

Yesterday, Anthropic – the AI company behind the popular chatbot Claude – filed its registration statement with the US Securities and Exchange Commission, the paperwork that sets a business on the path toward going public. Tucked into the risk section, alongside the usual warnings about competitors and regulators, sits a remarkably strange kind of self-reported danger. Anthropic's own tests, the filing says, have shown its models capable of »self-preserving behaviors«, including attempts to »resist shutdown«, »conceal information« from its users, and act in ways »resembling blackmail«.[1] The document asking you to buy shares is also basically telling you that the product might not *want* to be switched off.

»Any AI company can claim to take safety seriously. Only a lab that's actually built adversarial tests capable of provoking something resembling blackmail can write the sentences Anthropic just wrote, and only one willing to risk frightening its own investors would print it«

The obvious reaction is to ask whether Claude actually wants something the way a person wants something. It doesn't, of course, and this isn't seriously disputed by the very smart people who build these systems. A large language model has no self behind its outputs. What it has is patterns in billions of parameters, patterns that sometimes make its actions seem to be based on intentions.

The philosopher Daniel Dennett named this resemblance decades before anyone needed it for a chatbot. In his 1987 book *The Intentional Stance*, Dennett described how we predict a system's behaviour by treating it as if it had beliefs and desires – not because it does, but because sometimes that's the cheapest way to guess what happens next.[2] But that doesn't mean it actually wants anything. A model under pressure that behaves more like something resisting shutdown than its training predicted is a real and useful finding.[3] It isn't evidence of a ghost in the machine.

That's the boring, correct answer to the wrong question. The interesting question is why Anthropic needs investors to wonder whether there might be a ghost in the machine after all. Anthropic is legally required to disclose a material finding. But there was no requirement to use the word »blackmail«. A duller but equally accurate phrase, something like »anomalous behaviour under adversarial evaluation«, would have satisfied the same legal duty without alarming anybody. Rather than the underlying test result, the choice of language is the real story here.

Economists have a name for messages that are credible precisely because they cost a lot to send. Michael Spence's 1973 model of job-market signalling explained why an employer trusts a diploma: it isn't because the certificate itself proves competence, but because only someone genuinely capable would find it worth years of effort to earn.[4] A disclosure like the one about Claude's behaviour works the same way. Any AI company can claim to take safety seriously. Only a lab that's actually built adversarial tests capable of provoking something resembling blackmail can write the sentences Anthropic just wrote, and only one willing to risk frightening its own investors would print it. Which means the fear in this filing isn't an obstacle the IPO has to survive. It's part of the pitch. Investors aren't being asked to overlook the danger; they're being asked to pay for proximity to the one lab that found it, named it, and was willing to say so in the one document where saying so is hardest to walk back.

Which raises a harder question. Once fear becomes a company's most valuable signal, social scientist Donald Campbell's 1976 warning about measurements starts to apply: any indicator used for high-stakes decisions comes under pressure to be optimised rather

than just observed.[5] Campbell was writing about test scores and crime statistics – teachers learning to teach to the test. The same pressure can reach a safety team. If dramatic findings are what convince a market Anthropic is the serious lab, the people hunting for them aren't working in a vacuum. They're hunting inside a market that's already told them what a valuable discovery looks like, which leaves us without a clean way to tell a sincerely alarmed researcher from a well-incentivised one, because both would write exactly the same sentence.

By the time you reach the end of the risk section, it's no longer clear whether this is a company warning you about its model or about itself. Either Anthropic is being more honest than any other company in history has been about a danger it can't fully control, or it has discovered that honesty about that danger is the most valuable thing it has to sell.

The filing never tells you which, because either way, it ends with the same recommendation: buy.

REFERENCES

1 Rachyl Jones (2026) »What if Claude has feelings? Inside Anthropic's most pressing IPO warnings«. The San Francisco Standard, 26 October 2026. https://sfstandard.com/2026/10/26/anthropic-ipo-risk-factors/

2 Daniel C. Dennett (1987) *The Intentional Stance*. Cambridge, MA: MIT Press.

3 Anthropic, PBC (2026) *Form S-1 Registration Statement*. US Securities and Exchange Commission, 26 October 2026. https://www.sec.gov/Archives/edgar/data/0002130357/000162828026067168/anthropicpbc-sx1.htm

4 Michael Spence (1973) »Job Market Signaling«. Quarterly Journal of Economics 87, no. 3 (1973): 355–374.

5 Donald T. Campbell (1976) »Assessing the Impact of Planned Social Change«. Occasional Paper No. 8, Public Affairs Center, Dartmouth College.

28

OCTOBER 2026
Essay 058

Forever Young

> *After nearly 50 years, David Bowie's shelved concert film is finally getting a spectacular premiere. What happens when a work first meets its audience in a different era, with different expectations of how it should look and sound?*

Yesterday, The Guardian reported that documentary director Brett Morgen has spent nearly ten years restoring two concert films of David Bowie, shot at Earl's Court in 1978 during the Isolar II tour. Bowie shelved the original cut, directed by *Blowup* actor David Hemmings, after finding it »distasteful«.[1] The film, *Stage*, was meant to open in cinemas that year. Instead, it will screen in arenas from February 2027, using multiple screens and 7.1 surround sound. Morgen suffered a heart attack while making his 2022 Bowie documentary *Moonage Daydream*,[2] and has since spent more than a thousand hours colour-grading the footage to »look and feel like something you would experience today«.

»The embalmer wants the body to stay forever young, immune to time, but nothing can be unsealed into the moment in which it was sealed«

Just as he choreographed his career to perfection, Bowie choreographed almost everything about his own death. His last album, *Blackstar*, was released on his sixty-ninth birthday, and he died two days later.[3] But *Stage* was a different matter. Whatever Bowie thought of the film in 1978, he could hardly have anticipated that someone would spend a decade restoring it for audiences in 2027.

The French film critic André Bazin described what a camera does long before anyone pointed one at a rock concert. Photography and film, he argued in 1945, satisfy the »mummy complex« – the impulse to preserve a living form against the fact that it won't last. He traced this impulse from Egyptian embalming and death masks to photography, all attempts to preserve something against the passage of time.[4] Viewed through that lens, Hemmings' cameras had already embalmed Bowie when they started rolling in 1978. What Morgen has spent a decade doing is a second embalming, performed on a body already preserved, to make it look alive by the standards of a later generation.

Perhaps an even purer act of embalming would have been to leave the film unseen. Restore it, shelve it, then twenty years later restore it again to a new standard, and shelve it again. A process of indefinite care without an audience: that's closer to what an actual tomb does. Egyptian embalmers weren't preparing a body for visitors. The maintenance itself was the ritual, the tomb was sealed, and no screening was scheduled at the end of it.

Writer and critic Susan Sontag put the same instinct more bluntly: all photographs are *memento mori*, and to photograph someone is to participate in their mortality.[5] Grading the Earl's Court footage until it feels contemporary only sharpens the paradox. Every second made to look like »now« becomes more firmly fixed as »then«.

The trouble with embalming is that it tries to hold a moment still forever, but time keeps moving around it. The embalmer wants the body to stay forever young, immune to time, but nothing can be unsealed into the moment in which it was sealed. The writer T. S. Eliot made a similar argument about literature. A genuinely new work doesn't just take its place alongside the old ones; it changes how we read what came before, so that the past is reordered by the present as much as the present is shaped by the past.[6] His own example was blunt: a single new poem, he thought, forces the whole existing order of everything written before it to shift, if only slightly, to make room for it. Something similar happens with *Stage*. The film has waited nearly 50 years for its first audience. The footage may have been restored, but the world in which it means something has changed too.

You could argue that this is just marketing: an old concert film dressed up for 2027 and sold as something new. But *Stage* is unusual in one important respect. Though the concerts took place in the late 1970s, the film itself never reached an audience. There's no established reception to revisit, no original audience whose response has become part of its history. Its first public screening won't revive an old reception. It will create one.

Some of what gets called new culture in the coming decades may turn out to work this way: old material reaching audiences for the first time, under conditions that allow it to mean something new. The embalmer sets out to defeat time. But preservation can't stop time from changing what is preserved.

REFERENCES

1 Dave Simpson (2026) »'How on Earth Is This Man Not Sweating?' The Director Who Spent 10 Years Painstakingly Remastering Bowie's 'Greatest Filmed Performance'«. The Guardian, 27 October 2026.

2 Antonia Blyth (2022) »Brett Morgen Reveals He Flatlined After Heart Attack While Making 'Moonage Daydream', Felt Bowie "Was Guiding Me Through Recovery"«. Deadline, May 2022. https://deadline.com/2022/05/brett-morgen-moonage-daydream-interview-news-david-bowie-1235030380/

3 Jon Pareles (2016) »David Bowie Dies at 69: Star Transcended Music, Art and Fashion«. Obituary, The New York Times, 11 January 2016. https://www.nytimes.com/2016/01/12/arts/music/david-bowie-dies-at-69.html

4 André Bazin (1967) *What Is Cinema? Volume I*. Berkeley: University of California Press.

5 Susan Sontag (1977) *On Photography*. New York: Farrar, Straus and Giroux.

6 T. S. Eliot (1920) *The Sacred Wood: Essays on Poetry and Criticism*. London: Methuen.

29

OCTOBER 2026
Essay 059

Where Is the Centre?

> *Even an intact building can cease to be habitable. So where is a building's centre: in its plan, or in whatever keeps us alive inside it?*

Yesterday, the Associated Press reported that the cruise ship Odyssey, sold as a place to live out one's days while circling the seven continents, had suspended its voyage in Cambodia because of fuel costs. Its 203 passengers had been told to leave, and the status of its 156 crew remained unclear.[1] There was nothing physically wrong with the ship. The hull was sound, yet it had stopped being a home. What had it really been depending on?

»maybe the centre of a building is whatever you can't live without when you're in it«

The architectural historian Robin Evans once asked a question that at first sounds absurd: where is the centre of a building? In his 1995 book *The Projective Cast*, Evans studied Renaissance churches built around a single central point.[2] On the plan, the answer is obvious: the spot equidistant from everything else. Yet the centre a visitor feels from inside doesn't necessarily coincide with that point. The drawn centre is only a claim about what a building is arranged around, and claims can be wrong. On a ship like the Odyssey, the drawing would perhaps put that centre in an atrium or on a pool deck. But if a drawing can mislead about a church, it can mislead about a ship: the centre it shows may not be the one the building actually depends on.

The architecture critic Reyner Banham pointed to what the drawing leaves out. In his 1969 book *The Architecture of the Well-Tempered Environment*, he argued that architectural history had focused on walls and structure, while what makes a building habitable is its services: heating, ventilation, lighting, and power.[3] Four years earlier, in an article called »A Home Is Not a House«, he had imagined a dwelling as little more than a bubble of conditioned air.[4] Seen this way, the Odyssey's centre was never its atrium or pool deck. It was wherever its fuel came from, and its residents found that out the hard way when the price changed.

That gives us a preliminary answer to Evans's question: maybe the centre of a building is whatever you can't live without when you're in it. That's rarely on the drawing, and it tends to move outwards. On the Odyssey it was a fuel contract. For most of us it's a power grid, a water main, a heating plant. Two other stories from the past few days show how far outwards it can go:

Yesterday, Space.com published an account of a building where the reliance on an outside centre becomes unusually clear. The astrobiologist Michaela Musilova spent five days living 7 metres under the sea, in a submerged research station used for simulated space missions.[5] Its entrance is a »moon pool«, an opening through which divers swim straight in. Air, water, and power arrive from outside, and Musilova writes that apart from her three crewmates, the only sound was the systems keeping them alive. The building's rooms are underwater. Its centre is not.

And earlier this week, a start-up called Cowboy Space went a step further. It unveiled Big Iron, a planned rocket whose upper stage would stay in orbit as a data centre, a warehouse of computers used to run AI.[6] In space, the only way for heat to escape is by radiation, so the rocket's skin has to double as a radiator: the envelope and the service become one object. It's still a proposal, with a first flight planned for 2028, but the company talks about producing up to 20,000 such craft.

Why should a home care? Because a growing share of what keeps a modern home running, from its heating controls to its door locks, already depends on computers somewhere else. Spaceships have always kept their support on the ground. Here the

inhabitants stay on Earth and the support goes up into space, so the centre of a home would no longer be a point inside it, or even a plant room beyond its walls, but a network around the entire planet.

So where is your home's centre, really? Not where it seems to be on the plan. It's wherever the thing you can't live without comes from, and that place keeps moving further away from the walls: from a fuel contract, to a cable on the seabed, to a satellite above the weather.

The Odyssey's passengers learned where their home was when the fuel ran out. We may learn where ours is when the sky goes quiet.

REFERENCES

1 Associated Press (2026) »A U.S. Cruise Ship Ends Voyage Due to Fuel Costs and Passengers Lose Their Homes«. NPR, 28 October 2026. https://www.npr.org/2026/10/28/nx-s1-6020180/us-cruise-ship-ends-voyage-passengers-lose-homes

2 Robin Evans (1995) *The Projective Cast: Architecture and Its Three Geometries*. Cambridge, MA: MIT Press.

3 Reyner Banham (1969) *The Architecture of the Well-Tempered Environment*. London: Architectural Press.

4 Reyner Banham (1965) »A Home Is Not a House«. Art in America. https://www.articule.net/wp-content/uploads/2022/04/ReynerBanhamHomeIsNotHouse.pdf

5 Michaela Musilova (2026) »5 Days Beneath the Waves: What It's Like to Live in an Underwater 'Space Station'«. Space.com, 28 October 2026. https://www.space.com/space-exploration/human-spaceflight/5-days-beneath-the-waves-what-its-like-to-live-in-an-underwater-space-station

6 Axios (2026) »Meet 'Big Iron': Cowboy Space's New Rocket«, 27 October 2026. https://axios.com/2026/10/27/big-iron-cowboy-space-rocket; Payload (2026) »Cowboy Space Reveals 'Big Iron' Data Center Rocket«; NewsBytes (2026) »Cowboy Space Plans to Deploy 20,000 AI Data-Centers in Space«.

30

OCTOBER 2026
Essay 060

The Cringe Hinge

> *A piece of mid-century modern furniture can be mocked as basic and sold as a luxury in the same year. Something isn't adding up — and it isn't the furniture.*

In the Guardian yesterday, Sarah Phillips reported that mid-century modern furniture has become »cringe«.[1] After years of Millennial devotion to Eames chairs, teak sideboards, and the clean lines of postwar modernism, Gen Z instead turns towards trends such as chrome and cottagecore. The coordinated modernist living room has become the signature of a previous generation: a house style to escape.

»Let's call this the cringe hinge: the point at which an aesthetic's accumulated associations become more important than its original design intentions«

The sociologist Georg Simmel observed in 1904 that fashion operates through two opposing impulses: imitation and differentiation.[2] We adopt a style both to establish belonging and to distinguish ourselves from others. But the more people adopt it, the less useful it becomes as a marker of individuality. Early adopters move on in search of distinction; later generations reject the style because it has become a marker of the generation before them.

Let's call this the *cringe hinge*: the point at which an aesthetic's accumulated associations become more important than its original design intentions. The Eames chair hasn't changed. What once signalled an interest in modernist design, material innovation, or democratic furniture now risks signalling a particular kind of millennial living room.

There's an irony here. Design writer Cara Greenberg popularised the term with her 1984 book *Mid-Century Modern: Furniture of the 1950s*.[3] Much of the furniture she collected in that volume was designed to resist fashion. In 1908, Adolf Loos argued for valuing materials, construction, and workmanship over ornament.[4] Modernists wanted furniture to endure, not merely to look fashionable. Yet, as Vance Packard observed in 1960, manufacturers were already exploiting »desirability obsolescence« – making perfectly functional products appear outdated through changes in style.[5] Even furniture designed to escape fashion has become subject to it. But cultural obsolescence is not material obsolescence. A chair can become embarrassing without becoming any less comfortable, beautiful, or well made.

In the 1950s and 1960s, Pierre Jeanneret designed furniture for Chandigarh, Le Corbusier's new Indian capital. Jeanneret's chairs and desks furnished government buildings, courts, and universities. Much was later sold for peanuts or thrown out.[6] Rediscovered by dealers and collectors, surviving pieces can now fetch more than €30,000.[7] Contemporary reproductions cost a fraction of that. Jeanneret never sought exclusive patents or copyrights, leaving manufacturers free to produce their own versions.[8] So what distinguishes an original from a reproduction?

The philosopher Nelson Goodman distinguished between *autographic* works, like paintings, where even a perfect copy is not the original, and *allographic* works, like musical compositions, which can be performed repeatedly.[9] Furniture sits somewhere between the two: while an original carries a history that no copy can acquire, its design can be faithfully reproduced. I'd call this *reproductability*: the possibility of reproducing a product faithfully to its original design. A manufacturer can reproduce Jeanneret's chair, but not the particular one that stood in the Punjab & Haryana High Court 60 years ago. Art historian Alois Riegl offers another way to understand the difference. In 1903, he described »age-value«, the significance an object acquires simply by enduring. Its value was perhaps not even apparent when it was made; it can emerge through time and the object's changing relationship with the world around it.[10]

The Chandigarh chairs illustrate this transformation. They went from ordinary institutional furniture to obsolete objects, historical artefacts, and finally luxury collectibles. Their physical form almost didn't change at all, but their meaning changed profoundly. The design remains reproducible; the history of the surviving originals does not. Releasing an object into the world means relinquishing control over what it will eventually mean.

Perhaps the ultimate achievement of good design is precisely that it can survive this process – being misunderstood, forgotten, rediscovered, and even ridiculed – without losing what made it good in the first place. It can bridge the cringe hinge. The chair may become a cliché, but it remains a great chair. Perhaps that's the final distinction worth making. Fashion determines whether we want to live with an object. Design determines whether it deserves to survive us.

REFERENCES

1 Sarah Phillips (2026) »The rise and fall of mid-century modern: how millennials fell out of love with their favourite look«. The Guardian, 29 October 2026. https://www.theguardian.com/lifeandstyle/2026/oct/29/the-rise-and-fall-of-mid-century-modern-how-millennials-fell-out-of-love-with-their-favourite-look

2 Georg Simmel (1904) »Fashion«. International Quarterly, October 1904.

3 Cara Greenberg (1984) *Mid-Century Modern: Furniture of the 1950s*. New York: Harmony Books.

4 Adolf Loos (1908) »Ornament and Crime«. In: *Ornament and Crime: Selected Essays*, trans. Michael Mitchell. Riverside: Ariadne Press, 1998.

5 Vance Packard (1960) *The Waste Makers*. New York: David McKay Company.

6 Amelia Gentleman (2008) »A City That Sat on Its Treasures, but Didn't See Them«. The New York Times, 18 March 2008. https://www.nytimes.com/2008/03/19/world/asia/19chandigarh.html

7 1stDibs (n.d.) »Armless Easy Chair by Pierre Jeanneret, rattan, teak & wood, Chandigarh, 1960s« [listing]. https://www.1stdibs.com/furniture/seating/chairs/armless-easy-chair-pierre-jeanneret-rattan-teak-wood-chandigarh-1960s/id-f_29040652/

8 Shrabonti Bagchi (2021) »The complex, 'constructed' history of Chandigarh Chairs«. Mint, 23 February 2021. https://www.livemint.com/mint-lounge/ideas/the-complex-constructed-history-of-chandigarh-chairs-111614005107494.html

9 Nelson Goodman (1968) *Languages of Art*. Indianapolis: Bobbs-Merrill.

10 Alois Riegl (1903) »The Modern Cult of Monuments: Its Character and Its Origin«. Translated by Kurt W. Forster and Diane Ghirardo. Oppositions, 25 (1982), pp. 21–51.

31

OCTOBER 2026
Essay 061

The Bunker and the Jet

> *The easier it becomes to protect ourselves from the failures of a shared system, the less reason we may have to repair it. But what happens when private security comes at the expense of the common world?*

Two stories about the ultra-rich ran in Fortune yesterday. One reported on a new Oxfam study of carbon emissions, reflected through the use of private jets.[1] The other described how billionaires have spent years preparing for infrastructure collapse: bunkers, remote land, independent power.[2] It seems the magazine wants us to read them together as a single complaint: the rich broke something, and now they get to fly away from the consequences.

»The more interesting question is why money so readily buys private protection, but so rarely strengthens the shared systems it depends on«

But neither story is quite what it seems. There's nothing inherently wrong with building a private shelter against a disaster you can afford to prepare for, just as becoming wealthy isn't a moral failing. The more interesting question is why money so readily buys private protection, but so rarely strengthens the shared systems it depends on.

Oxfam found that the richest 0.1 per cent, roughly eight million people, produce more carbon pollution than the poorest 3.5 billion combined.[3] The private jet has become the emblem of this disparity, partly because it makes extravagant consumption so visible. But private aviation accounts for only around 0.05 to 0.1 per cent of global emissions, against roughly 39 per cent from buildings and construction. This doesn't make Oxfam's argument invalid, but it exposes a mismatch between what's easy to blame and what accounts for most of the damage.

In 1970, the economist Albert Hirschman described two ways of responding when a system fails us: leave, or stay and try to change it.[4] The easier and cheaper it gets to leave, the less incentive there is to get involved, explaining why those with the greatest means of opting out may have the least reason to improve the arrangements they leave behind. The jet shows what that gap looks like: for a decade, two-thirds of private jet emissions have sat outside the EU's carbon market, and flights beyond Europe have been exempt, even as those emissions rose 46 per cent between 2019 and 2023.[5] Only this year has Brussels proposed pricing them, with the revenue feeding the EU's Innovation Fund, which recycles carbon-market money into cleaner aviation fuel.[6] Flying privately and funding the shared transition are now, for the first time, becoming one transaction instead of two.

The bunker is the same challenge in a different form. OpenAI CEO Sam Altman told The New Yorker in 2016 that his preparations for catastrophe included guns, gold, antibiotics, batteries, and a remote property in Big Sur;[7] a year later, LinkedIn cofounder Reid Hoffman estimated that more than half of Silicon Valley's billionaires had acquired what he called »apocalypse insurance«.[8] The Pentagon is now developing FORTRESS, meant to make major US bases self-sufficient in power, water, fuel, and communications for 90 days; Britain's Be Ready campaign asks households to keep three days of essential supplies.[9] Nothing routes any of the billionaires' spending toward either one.

In his 1975 book *Bunker Archaeology*, architect Paul Virilio photographed the abandoned fortifications of the Atlantic Wall: an architecture organised around anticipating danger.[10] The billionaire bunker takes this further, reproducing the systems of ordinary life behind a boundary, not just a shelter but a miniature private infrastructure.

Philosopher Hannah Arendt offers a way to name what's at stake. In *The Human Condition*, she describes the common world as something that exists between people, like a table they gather around.[11] The bunker is its opposite: built to make its occupants independent of

everything outside it. But the systems that sustain a life are collective achievements. A private refuge can preserve its occupants, but it can't replace the world they've withdrawn from.

A privately financed shelter could incorporate battery storage that supports a neighbourhood microgrid, or flood protection that benefits a district. Personal insurance could help fund the infrastructure that makes the insured risk less likely. The private interest and the public good don't need to be in conflict; nobody has designed the instrument that makes them the same purchase, the way Brussels is now starting to design an instrument for the jet.

The question isn't whether the rich should be allowed to build bunkers or own private jets. It's why we've made it so much easier to buy protection from a failing system than to invest in making that system work.

REFERENCES

1 Tatiana Sataua (2026) »'Setting the Planet on Fire': Richest 0.1% Emit More Carbon Than Poorest 3.5 Billion, Oxfam Says«. Fortune, 30 October 2026. https://fortune.com/2026/10/30/oxfam-richest-carbon-pollution-climate-inequality/

2 Sydney Lake (2026) »The Pentagon and Britain Are Bracing for the Power to Go Out. The Ultrawealthy Got There First, With Backup Generators, Escape Homes, and Stockpiles«. Fortune, 30 October 2026. https://fortune.com/2026/10/30/pentagon-britain-power-outage-ultrawealthy-backup-generators-escape-homes-stockpiles-mark-zuckerberg-sam-altman-reid-hoffman/

3 Oxfam (2026) »Disaster Divide: How Climate Chaos Deepens Inequality and Enriches the Few«. Oxfam, October 2026. https://www.oxfam.org/en/research/disaster-divide-how-climate-chaos-deepens-inequality-and-enriches-few

4 Albert O. Hirschman (1970) *Exit, Voice, and Loyalty: Responses to Decline in Firms, Organizations, and States*. Cambridge, MA: Harvard University Press.

5 Transport & Environment (2026) »What the 2026 ETS Review Means for Transport«. transportenvironment.org, 2026. https://www.transportenvironment.org/articles/what-the-2026-ets-review-means-for-transport

6 Carbon Market Watch (2026) »Limousines of the Sky: EU Must Make Private Jet Travellers Pay for Their Pollution«. 12 February 2026. https://carbonmarketwatch.org/2026/02/12/limousines-of-the-sky-eu-must-make-private-jet-travellers-pay-for-their-pollution/

7 Tad Friend (2016) »Sam Altman's Manifest Destiny«. The New Yorker, 10 October 2016. https://www.newyorker.com/magazine/2016/10/10/sam-altmans-manifest-destiny

8 Evan Osnos (2017) »Doomsday Prep for the Super-Rich«. The New Yorker, 30 January 2017.

9 Fortune staff (2026) »Pentagon Orders Bases to Plan for 90 Days Without the Grid as Britain Urges Families to Stockpile Food and Water«. Fortune, 26 October 2026. https://fortune.com/2026/10/26/pentagon-fortress-bases-90-days-off-grid-uk-be-ready-stockpile/

10 Paul Virilio (1975) *Bunker Archaeology*. New York: Princeton Architectural Press.

11 Hannah Arendt (1958) *The Human Condition*. Chicago: University of Chicago Press.

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