Technology | *Book Review*

Letting go

Why nations win or lose the technological race

By Philip Zelikow

PREVIOUS ARTICLE
Growing together

The 400 metre race at the World Humanoid Robot Games, Beijing, August 23, 2026 | © Pedro PARDO/AFP via Getty Images

September 4, 2026
Read this issue ›

IN THIS REVIEW

How Progress Ends
Technology, innovation, and the fate of nations
552pp. Princeton University Press. £30 (US $35).
Carl Benedikt Frey

AI for Good
How real people are using artificial intelligence to fix things that matter
272pp. Simon and Schuster. £20.
Josh Tyrangiel

The public argument about artificial intelligence tends to create a binary of machines and humans. On one side are prophets who urge humans to accept "machines of loving grace" (the Anthropic co-founder Dario Amodei's phrase). On another side are prophets such as the journalist Garrison Lovely, whose forthcoming book is entitled *Obsolete: The AI industry's trillion-dollar race to replace us - and how to stop it*. This binary misleads. The wide use of AI is already a fact. Yet its fate and meaning will be derived from the uses that humans design. We are still driving the story.

"About 1760 a wave of gadgets swept over England." Thus the economic historian T. S. Ashton recalled how "a schoolboy began his answer to a question on the industrial revolution" (in *The Industrial Revolution, 1760-1830*, 1948). As a wave of AI models now sweeps over the world, it helps to observe that progress is not a pattern of some great invention followed by simple adoption. It is more useful to think of an evolutionary process. An initial innovation is followed by a succession of improvements, some of which may become transformational. The later innovations may straddle fields and disciplines in unexpected ways. When someone thought to pair lasers with glass fibres, this revolutionized communications. Biotech exploded when it combined the discovery of genetic code with digital technologies to splice genes and process vast genomic data.

Indeed, the computer itself was once thought to have highly limited uses, even by its original innovators. Reflecting on this in a celebrated essay, "Uncertainty and Technological Change" (1996), Nathan Rosenberg noted that "new technologies typically come into the world in a very primitive condition. Their eventual uses turn upon an extended improvement process that vastly expands their practical applications". And this process is not only technical.

Here the prophets on both sides share an assumption so large that it usually passes unexamined: that the technology itself will determine the outcome. The historical record says otherwise. Between the invention of a powerful technology and the dispersal of its benefits stands a long, contingent, frequently botched process in which societies do or do not reorganize themselves – their tooling, their businesses, their occupations, their laws. That process, more than the technology itself, is what shapes the future.

Two fine new books examine that process with diametrically different instruments. Using an ultra-wide-angle lens, the scholar Carl Benedikt Frey conducts a masterly sweep across two millennia to ask why progress starts and why it stops. Using a powerful microscope, the journalist Josh Tyrangiel trains his gaze on the experimenters themselves, reporting from schools, a healthcare system, municipal waste facilities, the American federal tax agency (the IRS) and the nerve centre of a giant vaccine distribution effort to consider how these bodies have put artificial intelligence to work. Read together, these books frame the core issue. Theirs is less a story about wondrous technology than a drama of the often frustrating, often inspiring "modernization journey" – a phrase used by the chief technology officer of the IRS, Kaschit Pandya – of those, like Pandya, struggling to transform the sagging or broken systems of everyday life.

Frey's previous book, *The Technology Trap* (2019), examined why workers have sometimes resisted machines and why their resistance has sometimes prevailed. His latest book, *How Progress Ends*, generalizes the inquiry, his argument resting on the distinction between exploration and exploitation. Frey shows how decentralized systems – such as those in most free, pluralistic societies – are good at exploration. They excel when the frontier must be extended into the unknown, which requires many concurrent experiments, most of which will fail. In contrast, centralized, bureaucratic systems – think of Imperial China or today's People's Republic – are good at exploitation. They excel when the task is to absorb known technologies, mobilize resources and catch up to a visible frontier.

Neither form of governance is best everywhere and always. Each has, in Frey's phrase, its ecological niche. The tragedy is that a system optimal for one phase can become, almost inevitably, ill suited for the next. The author's conclusion may seem surprising. It is that, “while it has become popular to pitch China and the United States as locked in an economic and technological race, the truth is that both are trending toward stagnation, and further headwinds lie ahead as Chimerica starts to disintegrate”.

Frey works his thesis through an impressive range of cases: from Chinese mandarins to European guilds, Prussia’s “visible hand”, Soviet planning, postwar Japan and the American ascendancy. He shows how, while the British kicked off the Industrial Revolution in the eighteenth century, assisted by an environment of decentralization, weakened guilds and an ascendant merchant class, it was highly centralized Prussia that, a century later, most profitably exploited the new technologies. The author’s account of the Soviet Union, building on the work of scholars such as Chris Miller, is exemplary. Nikita Khrushchev grasped that computers were the future and built Zelenograd, a whole city for semiconductors, staffed by some of the world’s finest scientists. It failed anyway, because its efforts were orchestrated from a single ministerial desk in Moscow, while California’s inventors job-hopped from start-up to start-up (assisted by local laws that eschewed non-compete agreements, opening up labour markets), running experiments no committee had approved. The problem was with institutions designed to protect established powers and ways of doing things. Frey’s epigraph for a late chapter comes from the sociologist Jane Jacobs, who argued that the primary economic conflict is not between owners and workers, but between those invested in established activities and those whose interests lie with emerging ones – a conflict, Jacobs noted, that can be put to rest only by stagnation.

Frey’s book about the past hands down a blunt verdict on the present. The reason he judges that the US and China are both trending towards stagnation is that Xi Jinping is currently recentralizing an economic system whose growth came from letting go, while in the US, corporate consolidation, procedural handcuffs and politicized institutions hamper exploitation and corrode exploration. Europe fares worse still: Frey calculates the costs of its internal barriers and regulatory reflexes with the grim relish of a sardonic coroner. His conclusion is that progress, as he puts it, is “neither automatic nor natural”. If it were, the Industrial Revolution would have come centuries sooner – and it would have landed first in China.

Where does AI fit in with this? Frey resists the fashionable inference that a technology of intelligence will dissolve the old constraints. Progress, he argues, is constrained less by intelligence – human or artificial – than by the people in power and the incentives that shape their decisions. Yes, the machinery may become smarter, but that does not mean the guilds will therefore become weaker. Tyrangiel interviewed two people who worked hard at innovating inside the US government. One, Jennifer Pahlka (the executive director of the non-profit Recoding America), deals with discouraged innovators all the time. “Most of what I do on a daily basis is like therapy”, she says. The other, Eric Schmidt, is the former CEO of Google and now leader of three important newer tech companies. Asked about his years on the US Defense Innovation Board, he “took a breath and paused for a moment”, Tyrangiel recounts. “‘I’m not going to make a more emotional argument, I’m just going to tell you the following: Government will perform suboptimally until it adopts the software practices of the industry.’” To Tyrangiel, Schmidt “sounded pretty emotional”. What were those “software practices of the industry”? Schmidt boiled them down to: “Rapid prototyping with a small group of engineers, lots of user surveys, a continuous loop of refinement, iteration, and deployment”. He continued, “Every single thing I just told you is illegal in the federal government”, mainly due to its rigid procurement rules.

The strength of Frey’s magnificently informative book is also its limit. The units of analysis in *How Progress Ends* are nations and systems of governance. But adaptation happens inside particular hospitals, agencies, firms and schools, one redesigned job at a time. Frey can tell us why civilizations fail to adapt. He has less to say about what adapting actually looks like on a Tuesday.

That is Tyrangiel’s subject. *AI for Good* is a work of reportage, not of theory. The author spent years inside institutions that are actually trying to use the new tools against problems that matter: school districts across the US experimenting with the non-profit Khan Academy’s Khanmigo tool as a personalized tutor or teacher’s assistant; clinicians at the Cleveland Clinic using a tool called Bayesian Health to monitor patient data, detect signs of sepsis and guide care; trash collectors in East Lansing, Michigan using a company called Prairie Robotics to detect recycling contamination and help households to do better; the enforcement and taxpayer-service machinery of the IRS; or the leaders of the American Covid vaccine programme, Operation Warp Speed, using Palantir to help manage their daunting logistics.

Tyrangiel wants to know what happens when the frontier models meet the real world, where the condition or value of the data is messy, the incentives are misaligned and everything that matters is charged with human emotion. The Khanmigo struggle, for example, has taken years of refinement and adaptation. It is a story featuring some quite admirable people who discover that, rather than replacing teachers, their AI-powered tool turns out to be a powerful partner that makes good teachers much better. “Khanmigo was still glitchy”, Tyrangiel writes, but “teachers began to trust that they could send kids at various degrees of mastery to [it] during a lesson, and that the bot would meet them with the right content at the right pace.” Student engagement “started strong and never dropped”. The adaptation is still evolving, but for one veteran chemistry teacher who became a convert, “what’s striking isn’t how fast her teaching changed. It’s how easily it could have stayed exactly the same”, because the old customs were so deeply entrenched.

In Tyrangiel’s pages we see that the enabling power of AI might be roughly constant, but the results are not constant at all. So much is about the people. When the work is redesigned around the tool – when someone rethinks who does what and the surrounding institutions permit the rethinking – the gains can be startling. When the tool is treated as a panacea, or bolted onto the old workflow, the gains are little more than the stuff of press releases, or vapourware. Economic historians such as Paul David and Alexander Field long ago observed that American factories captured the productivity of electricity only after a generation's delay. This was not because the dynamo was immature, but because its gains required rebuilding the factory and reorganizing the workforce around it. That took new managers with new conceptions – and quite a lot of investment. And even with this lag America still led the world in the speed of that conversion.

Frey's and Tyrangiel's books need each other. And here a middle body of evidence, with which neither engages at length, helps to fill out the picture. A decade ago, in work I conducted with the Markle Foundation for a group calledRework America, I drafted a book called *America's Moment* (2015). My colleagues and I argued that the digital revolution's technologies could be used to replace people or to empower people, and that the empowering path was both available and better. The question we posed still seems right, and it has a Socratic rhythm. Can the new technology empower ordinary workers? Yes. Will it? It depends. Depends on what? Partly on markets. But only partly.

The rest depends on creative managers, on finance to support fresh business ideas, and on law. Consider the managers: Tyrangiel's protagonists are, in effect, the variance. The same tool in different hands produces different institutions. Managerial imagination is a scarcer input than the AI compute, which powers artificial intelligence. But managers operate inside rules, and this is where Frey's history and the present moment converge with real force.

Healthcare is a sector in which the gap between what the technology permits and what the rules allow is now very wide – and a quarrel currently running through every corner of the English-speaking world. Be fair to the rules first. Occupational licensing answers a real problem: patients cannot judge the quality of medical care in advance, so someone must vouch for it. This vouching is important. One of the strongest studies on the subject looks at quasi-random variation in nurse/doctor staffing in American veterans' hospitals, where a directive issued in 2016 let nurse practitioners treat emergency patients without physician supervision: a rare example in which the two professions could be studied independently carrying out genuinely overlapping work. It was found that, while thirty-day mortality did not differ significantly overall in cases handled by the different professionals, the nurses tended to use more resources than the physicians (hospital stays were a little longer; costs a little higher), with the gap widening in more severe cases. But the same study found something more interesting: the variation within each profession exceeded the differences between them. Nurse practitioners outperformed physicians in nearly four out of ten random pairings. And the mechanism behind the average gap was human capital – accumulated training and experience.

That is precisely the deficit that the evidence says AI-enabled tools are best at closing. So it is not that the professional boundary doesn't exist; it is that it is wrongly thought to be static. It is drawn as a permanent status frontier around occupations, rather than as an adjustable frontier around tasks, with no mechanism for updating as augmentation changes what a well-supported practitioner at the edge can safely do. Scope-of-practice statutes, licensure regimes and payment rules that reimburse credentialled inputs rather than outcomes all encode the pre-AI division of labour into law. Adaptation here is not merely difficult. In much of the US, and not only there, it is unlawful.

The UK is currently running two similar experiments inside a single health service. The Leng Review (2025) examined physician associates in the NHS. Many patient deaths had been linked to misdiagnosis by associates, so the "associate" role was renamed as "assistant" (making clear to patients that these practitioners were not replacements for doctors) and its holders were barred from seeing "undifferentiated" patients (those who have not yet been diagnosed). Yet the same NHS is moving a different boundary outwards: from this month, every newly qualified pharmacist in the UK will register as an independent prescriber, able to prescribe within their competence from the first day of practice. The boundaries, we see, can move. The question a well-governed system must answer is whether they move on evidence or on incumbency.

Frey's readers will recognize such stories at once, because his book is full of them. The medieval guild claimed jurisdiction over who might practise a craft, in the name of quality, and used that jurisdiction to strangle techniques that threatened its members. (As we have seen, guilds in England were weaker, which aided its industrial revolution.) Sheilagh Ogilvie's great study *The European Guilds* (2019), on which Frey draws, could be reissued with the chapter headings changed to the names of the modern professional bodies on either side of the Atlantic. The rhyme runs back a thousand years, and it is the evidence for Frey's darkest claim: the default outcome when a general-purpose technology confronts entrenched interests is that the interests win and the progress quietly ends.

Education and workforce development show the same structure in different garb: degree requirements screening workers out of jobs for which the degree does not predict performance; teacher certification regimes that decades of evidence show barely predict classroom effectiveness; accreditation systems that make it nearly impossible to charter new kinds of institutions.

The EU has managed the remarkable feat of building such boundaries in advance of the technology. Frey's chapters on contemporary Europe are the angriest in his book. He looks at data protection rules with compliance costs that entrench the very incumbents they were meant to discipline; an AI Act that classifies as "high risk" a suite of AI educational tools; and an environment of internal regulatory and non-tariff barriers that the IMF prices as equivalent to a 44 per cent tariff on goods and a 110 per cent one on services, in an EU that supposedly has a single market. A continent can regulate the improvement process before it has begun.

Yet Europe also supplies instructive counterexamples. The Swedish labour market made a different bargain half a century ago. Today in Sweden, employer-financed job security councils, created by collective agreement in the 1970s, retrain and replace workers whom technology displaces, and they work well enough that Swedish unions have become famous for welcoming automation.

This is where Tyrangiel's ground-level optimism earns its place alongside Frey's high-altitude pessimism. Tyrangiel's protagonists did not wait for a national decision. They are adapting, or failing, one at a time. Those who succeed are demonstrating, in the most persuasive currency there is, what works. Throughout history, the most influential ideology has been the ideology of what works. The demonstrations accumulate, and eventually they can become politically usable, as in some of the great deregulatory and liberalizing struggles in finance, trade, containerization and telecommunications that fuelled an American and European revival during the 1980s and 1990s – one that almost nobody had predicted during the depths of the pessimism of the 1970s.

So the fundamental question that both books circle is not whether the machines will master the humans or the humans will master the machines. And this is because whether AI becomes an instrument of broad human empowerment is not a property of the technology and will not be settled by it. Instead, it will be settled in regulatory fights in American statehouses and royal colleges, in reimbursement rules and collective agreements, in Brussels directorates and procurement offices, in the choices of managers nobody has heard of. Frey shows how often such contests have been lost, and how quietly. Tyrangiel shows that they are being contested now, by people who mostly do not know they are re-enacting a very old drama. What is different this time is the scale and speed of this drama's arrival. The record of the past thousand years says we often fail. The record also says that the exceptions built the modern world.

Philip Zelikow is the Botha-Chan Senior Fellow at Stanford University's Hoover Institution. His books include America's Moment: Creating opportunity in the connected age*, drafted for Rework America in 2015.*

Keep reading

Technology | *Book Review*

Machinations

The history of government by technology

By Carl Benedikt Frey

Technology | *Book Review*

Battle for the soul of the web

Three assessments of the online world

By Joanna Kavenna

Technology | *Book Review*

What the public wants?

Writing novels by AI - and committee

By Gordon Fraser

Technology | *Book Review*

Authentic intelligence

The maverick founder of the world's first $4 trillion company

By Emily Jones

Mediterranean Yacht Voyages

Unwind in expansive residential suites with warm, intuitive Four Seasons hospitality

Four Seasons Yachts

Learn More >

The TLS Newsletter

Join 40,000 readers to enjoy a regular dose of inspiration and motivation, delivered to your inbox every Thursday.

You information will be used in accordance with our Privacy Policy

Your email address

Send

Copyright © The Times Literary Supplement Limited 2026. The Times Literary Supplement Limited: 1 London Bridge Street, London SE1 9GF. Registered in England.

Company registration number: 935240. VAT no: GB 243 8054 69.