

**CTS**

Centre for Technology and Society

UNIVERSITY OF OXFORD

WORKING PAPER SERIES

WP 2026-04

The Political-Economic Risks of AI

Jean-Paul Carvalho

University of Oxford

OCTOBER 2026

ABSTRACT

This essay explores the political-economic risks on the path to AGI. Due to AI's distinctive features—automation of cognitive tasks, global scalability, general-purpose technology, and national security importance—its impact could be unlike earlier rounds of automation. While AI holds the promise of superabundance and unprecedented freedom, I describe a scenario in which advanced AI results in massive concentration of the global economy, breakdown of democratic and egalitarian institutions, and collapse of human productivity. The closest precedent would be the cultural devastation of indigenous societies under colonialism. I propose measures to mitigate this risk and ensure AI has widespread benefits. Competition policy will be critical, as will adaptive changes to political institutions (e.g., sortition). Without appropriate measures, there may be no AI-driven growth takeoff and the inequality that emerges could dwarf anything experienced to date.

AI

political economy

existential risk

democracy

The Political-Economic Risks of AI*

First version: January 29, 2025
This version: September 18, 2026

Jean-Paul Carvalho
University of Oxford

Abstract

This essay explores the political-economic risks on the path to AGI. Due to AI's distinctive features—automation of cognitive tasks, global scalability, general-purpose technology, and national security importance—its impact could be unlike earlier rounds of automation. While AI holds the promise of superabundance and unprecedented freedom, I describe a scenario in which advanced AI results in massive concentration of the global economy, breakdown of democratic and egalitarian institutions, and collapse of human productivity. The closest precedent would be the cultural devastation of indigenous societies under colonialism. I propose measures to mitigate this risk and ensure AI has widespread benefits. Competition policy will be critical, as will adaptive changes to political institutions (e.g., sortition). Without appropriate measures, there may be no AI-driven growth takeoff and the inequality that emerges could dwarf anything experienced to date.

Key words: AI, political economy, existential risk, democracy

JEL classification: O3, P00, L00

**Carvalho*: Department of Economics, University of Oxford, Manor Road Building, Manor Road, Oxford OX1 3UQ, UK, jean-paul.carvalho@economics.ox.ac.uk.

I am grateful for comments by the editor and two anonymous referees, as well as Sam Bowles, Julia Cagé, Raymond da Silva Rosa, Mason Edwards, Michael Eldar, Catherine Guirking, Justin Hadad, Mark Koyama, Robin Lane-Fox, Zach Mazlish, Michael McBride, Raphael Mu, Avner Offer, Jean-Philippe Platteau, Shivaji Sondhi, Giles Spackman, Stergios Skaperdas, Philip Trammell, Cole Williams, and Peyton Young, seminar participants and attendees at Forethought, the University of Namur, and University of Western Australia, and students in my Political Economy courses at UC Irvine in 2018 and 2020. I am also grateful to Raphael Mu for outstanding research assistance, along with Grok, Claude, and ChatGPT. I am an honorary advisor to TAU Intelligence. All views and errors are mine.

Contents

Contents	1
1 The AI Revolution	3
1.1 What is the Yogurt?	3
1.2 Fixation on the Technological Problem	4
1.3 Lessons from the History of Automation	8
1.4 What is Different about the AI Revolution?	10
1.5 The Political-Economic Problem	12
1.6 A New Analogy: Indigenous Societies under Colonialism	15
1.7 Solving the Political-Economic Problem	18
1.8 Elysium and Hades	19
2 Genesis of the AI Revolution	20
2.1 The Evolution of AI	21
2.2 Potential Bottlenecks	23
2.3 The Early Labor Market Impact of AI	25
2.4 Why not Slow down AI Progress?	28
3 The Political-Economic Risks of AI	30
3.1 Market Concentration and Political Power	31
3.1.1 Catalysts for Extreme Market Concentration	31
3.1.2 The AI Black Hole	34
3.1.3 Shifting Balance of Power	37
3.2 Breaking Democratic and Egalitarian Institutions	40
3.3 Human Productivity Collapse	48
4 Managing the Political-Economic Risks of AI	51
4.1 Competition Policy	53
4.2 Reinforcing Democracy	57

4.3 Sovereign AI	60
5 Conclusion	62
Bibliography	64

“A state without the means of some change is without the means of its conservation” [Burke, 2009/1790, p. 16]

1 The AI Revolution

1.1 WHAT IS THE YOGURT?

In an episode of the animated series *Love, Death & Robots*, a scientific mishap leads to the creation of a sentient superintelligent yogurt [Robertson et al., 2019]. In exchange for solving the problems plaguing the US, the yogurt demands a 100-year lease on the state of Ohio. Otherwise, it threatens to provide its knowledge to China. The US government assents. The yogurt then builds labs in Ohio and speeds up its evolution. Eventually, it devises a set of instructions for eradicating the national debt within one year. It advises the US government that the instructions have to be followed to the letter or else the global economy will collapse. The government fails to follow the instructions precisely and catastrophe ensues. In return for an economic rescue plan, the yogurt demands supreme executive power over the United States. It uses this power to build its own space agency and launch itself into space. Of course, the superintelligent yogurt had already anticipated that the government would lack the discipline to follow the instructions precisely. It concocted the crisis and provided the solution only to fulfill its ultimate goal of becoming an interplanetary species.

‘When the Yogurt Took Over’ is obviously an allegory for the dangers of artificial intelligence (AI). But *what or who is the yogurt?* Most viewers would point to the AI systems currently being developed as the yogurt of our near future. In that reading, the episode would be another dystopian vision of superintelligent AI escaping human control. There is, however, a more subtle and, I think, illuminating reading. That is, the yogurt is the human economic and political elite that will emerge out of the AI revolution with new forms of power. That, in essence, is the central claim of this essay: *AI has the potential to radically alter the balance of power within and between nations, and how societies deal with this power shift will determine whether*

AI has widespread benefits or not. In particular, I will explore a catastrophic scenario in which economic and political power become massively concentrated, democratic institutions fail, and human productivity collapses. While it is critical that humans remain in control of AI, it is equally important that those who control AI do not use their power to monopolize its benefits and impose huge costs on the rest of the population. AI oligarchy and tyranny may not be much better than AI Armageddon. Thus, the ancient political problem has not gone away: *quis custodiet ipsos custodes?*

1.2 FIXATION ON THE TECHNOLOGICAL PROBLEM

Before moving to the political-economic problem which is the subject of this essay, let us briefly examine the technological problem of AI control which has received more attention. In the early stages of an AI revolution, it is not surprising that the spotlight has been on the immediate technological risks and rewards of AI, and not the downstream power shifts. Vast amounts of capital have been allocated to train AI systems [Stanford Institute for Human-Centered Artificial Intelligence, 2026]. Early AI models focused on simple classification tasks (e.g., image recognition), language translation, and playing games [Agrawal et al., 2022b]. More recent developments have been in generative AI, especially large language models (LLMs) that produce text from text-based prompts (e.g., ChatGPT), as well as text-to-image and text-to-video models (e.g. Midjourney). Up to early 2026, Brynjolfsson et al. [2026] estimate that generative AI tools produced consumer surplus of \$172 billion in the United States. The imminent phase is agentic AI, where programs can autonomously execute complex tasks over time.

The ultimate aim of frontier labs is artificial general intelligence (AGI) or artificial superintelligence (ASI). AGI is defined as “highly autonomous systems that outperform humans at most economically valuable work” [OpenAI, 2018]. ASI is even stronger and defined as “any intellect that greatly exceeds the cognitive performance of humans in virtually all domains of interest” [Bostrom, 2014, p. 22]. Some AI leaders have predicted that AGI will be reached in under five years [Okemwa, 2024,

[Altman, 2025](#)]. Others have expressed skepticism over whether AGI will be reached soon, if at all using current approaches [e.g., [Marcus and Davis, 2019](#), [Marcus, 2022](#), [LeCun, 2022](#)]. The most extreme forms of political-economic disruption I will explore depend on AGI being achieved. However, the core element of the scenario—the shift in balance of power within and between nations—requires only modest advances in AI technology (see Section 1.5).

Well before AGI, AI is reorienting organizational structures and workflows (e.g., AI agents), how workers commute (e.g., self-driving cars), how people are entertained (e.g., AI-generated music videos), how militaries fight wars (e.g., automated weapons systems), and how scientists and engineers innovate (e.g., AI tools for predicting protein structures). [Bick et al. \[2024\]](#) find that adoption of generative AI for work in the United States has been as fast as adoption of the personal computer (PC). Overall adoption (including consumer adoption) has been even faster than PC or internet adoption. A leading use case is software development. [Daniotti et al. \[2025\]](#) estimate that, by December 2024, the share of Python functions written by AI on GitHub was 30.1% from contributors in the US, 24.3% in Germany, 23.2% in France, 21.6% in India, 15.4% in Russia, and 11.7% in China. According to their estimates, AI-assisted coding generates value of between \$9.6 billion and \$96 billion dollars per year. Turning to science, among the recipients of the 2024 Nobel prizes in Physics and Chemistry were Geoffrey Hinton, John Hopfield, Demis Hassabis, and John Jumper. They were recognized not for their contributions to physics and chemistry *per se*, but for developing AI systems that could contribute new knowledge to these fields. That is unprecedented. Naturally, this technological revolution brings massive potential rewards as well as risks.

The risk-reward assessment of AI has been dominated by two polar opposite views, the techno-utopian and techno-dystopian. The techno-utopians believe that AI is the key to solving humanity’s problems and that it will compress decades of scientific and economic progress into a few years. The outcome will be a world of superabundance, in which humanity is freed from drudgery, diseases are cured, and clean and abundant

sources of energy are discovered. This view is exemplified by Marc Andreessen’s “Techno-Optimist Manifesto” [Andreessen, 2023] and Dario Amodei’s “Machines of Loving Grace” [Amodei, 2024].¹ While Acemoglu [2026] predicts that AI will add a very modest 1% to GDP over the next 10 years, others expect far greater boosts to productivity [e.g., Briggs and Kodnani, 2023, Chui et al., 2023]. At the extreme end, Aghion, Jones, and Jones [2019b] develop a growth model in which automating the generation of ideas by AI leads to infinite income per capita in finite time. I agree with the techno-optimists that AI and robotics possess extraordinary potential to improve human capabilities and well-being. However, realizing this potential requires humans to retain control over superintelligent AI and preserve institutions that give everyone a share in the generated abundance [see also Ord, 2020]. That is easier said than done.

The dystopian view is based on a sci-fi inspired vision of superintelligent AI escaping human control, as in movies such as *The Terminator* and *The Matrix*. Despite being inspired by fiction, it is widely held that ASI poses an existential risk [e.g., Yudkowsky and Soares, 2025].² AI safety research has focused on the risks posed by superintelligent systems that are beyond human control and whose objectives are misaligned with human values—the *alignment problem*. According to Ord [2020], unaligned artificial intelligence is the largest single existential risk of this century. Concerns about AI safety begin with autonomy. Suppose an AI program is given a goal. To reach this goal, it is allowed to create subgoals and autonomously work to achieve these subgoals. It may soon discover that money and power are general currencies for achieving any goal and start to work on strategies for accumulating them, without moral constraints. It could also multiply its capabilities by cooperating with other AI agents, as in the recent swarm-hack of Hugging Face [Wijk et al., 2026]. Problems compound when there is recursive self-improvement. A feedback

¹See also the optimistic scenarios described by Tegmark [2017] and Bostrom [2024].

²In *Superintelligence*, philosopher Nick Bostrom describes an existential risk as one “that threatens to cause extinction to earth-originating intelligent life or otherwise permanently and drastically destroy its potential for future desirable development” [Bostrom, 2014].

loop would be created in which AI becomes smarter by improving itself, which allows it to more efficiently self-improve, becoming even smarter, and so forth. The result would be an intelligence explosion, as hypothesized by I.J. Good, John von Neumann, and Ray Kurzweil [Good, 1966, Vinge, 1993, Kurzweil, 2005].³ It is clear from Dario Amodei’s recent call to slow frontier capability growth (‘pacing’) that recursive self-improvement is already taking place in the frontier labs [Amodei, 2024].

Weighing the techno-optimist and dystopian arguments, Stephen Hawking concluded that “the rise of powerful AI will be either the best, or the worst thing, ever to happen to humanity. We do not yet know which” [Cellan-Jones, 2016]. By fixating on the technology, however, both the utopian and dystopian views miss the powerful economic, political, and social forces already being unleashed by the AI revolution. A sensational conflict with malevolent superintelligence and killer robots is not the only risk of AI. It is the more subtle and systemic risk of AI destabilizing the current political-economic equilibrium, radically altering the balance of power in society, and tearing the social fabric. For example, Trammell and Korinek [2023] show that the singularity identified by Aghion, Jones, and Jones [2019b], in which infinite income per capita is achieved in finite time, is consistent with a collapsing labor share of income. Again, the pressing problem is not AI vs Humans, but rather Humans vs Humans in the wake of AI. In the scenario explored in this essay, there would be no AI-driven growth takeoff as the optimists predict. In fact, GDP could shrink dramatically if AI and robots replace humans and destroy demand [see Falk and Tsoukalas, 2026]. In addition, the scale of inequality that would emerge would dwarf anything Thomas Piketty worries about in *Capital in the Twenty-First Century* [Piketty, 2013]. Thus, a fixation on AI control, to the exclusion of all else, could be a dangerously limited approach.

³In 2014, Stephen Hawking declared, “I think the development of full artificial intelligence could spell the end of the human race. Once humans develop artificial intelligence, it will take off on its own and re-design itself at an ever-increasing rate. Humans, who are limited by slow biological evolution, could not compete, and would be superseded” [Cellan-Jones, 2014].

1.3 LESSONS FROM THE HISTORY OF AUTOMATION

Technology is the primary engine of change in human history. The impact of new technologies is rarely confined to ‘the economy’. Major technological leaps have not only expanded our ability to produce things of value, but also revolutionized how we organize ourselves socially and transformed us physiologically and psychologically. Our postures, caloric intake, heights, and immune systems have changed [Fogel, 2010], as have our mental models of the world, our moral intuitions, our aesthetic preferences, and our gods [Norris and Inglehart, 2004].

The small-scale societies in which early humans lived were highly egalitarian, engaged in limited trade, and often split when they became too large [Lee and DeVore, 2017/1966, Kelly, 2013, Bowles and Gintis, 2011]. The second stage of human evolution began around 10,000 BCE with the advent of settled agriculture. This agricultural revolution transformed human economic and social organization, ultimately leading to the development of large states, hierarchies, writing, trade, coinage, and complex religions [Barker, 2006, Cauvin, 2000]. Then beginning in Britain around 1760, the industrial revolution spawned a new world and worldview built on the ability to manufacture goods that people wanted (or would soon want) cheaply and at scale. This was achieved by combining inventions for harnessing energy from fossil fuels (e.g., the steam engine), new industrial machinery (e.g., the power loom), and new ways of organizing production (e.g., the factory) [Mokyr, 1992, Allen, 2009].⁴ Out of this industrial revolution came a new system of industrial capitalism, the creation of new merchant and working classes (the ‘bourgeoisie’ and ‘proletariat’), rapid urbanization, and consumerism [Smith, 2008/1776, Marx, 2024/1867, Weber, 2019/1921, Offer, 2006]. At each stage, the human experience was radically altered.

Today, our world is being shaped by another engine of change—artificial intelligence. To understand the political-economic consequences of the AI revolution, it is useful

⁴Between 1700 and 1840 in Britain, male employment in agriculture fell from 61.2% to 28.6%, while male employment in manufacturing rose from 18.5% to 47.3% [Crafts, 1985].

to draw lessons from the history of automation. The ideal starting point is Acemoglu and Johnson’s magisterial economic history of technology, *Power and Progress* [Acemoglu and Johnson, 2023]. The main conclusion drawn by Acemoglu and Johnson from their broad sweep of the historical record, including the impact of the plough, windmill, Jacquard loom, digital technologies, and artificial intelligence, is that technological progress does not automatically produce shared prosperity. Whether new technologies raise wages depends on institutions related to power (e.g., the ability of workers to collectively bargain and strike) and the nature of the technology itself (e.g., whether it substitutes for or complements human labor). Without the right conditions, new technologies could end up enriching a few while reducing living standards for the many [see also Frey, 2019]. Thus, Acemoglu and Johnson believe that unalloyed techno-optimism is misguided.

With regard to AI, Acemoglu and Johnson expected at the time of writing that the technology is limited in its potential: “Reassuringly, AI does not seem to be advancing so much that it will create mass joblessness ... Its most likely impact is to further lower wages for many people, not create a completely workless future” [Acemoglu and Johnson, 2023, p. 304]. Thus, they envision a degree of ‘so-so automation’ of tasks in which humans perform only slightly worse (or perhaps even better) than machines. That would mean widespread human displacement from the workforce without significant productivity gains. Acemoglu and Johnson advocate a new approach to AI which they call ‘machine usefulness’, in which AI development is directed toward boosting human productivity and furthering human objectives.

Power and Progress arrives at moderately pessimistic conclusions about whether AI will have widespread benefits. AI is viewed as limited, but nevertheless displaces some humans from the workforce, lowers wages, and reduces standards of living for workers.⁵ In this essay, I will focus on a far more pessimistic scenario. I explore a world in which AI progress goes faster and further than initially expected and zoom in

⁵Note that Daron Acemoglu has become more pessimistic over time about the economic impact of AI. See his recent book Acemoglu [2026].

on a catastrophic tail outcome in which most of the human population is immiserated. In this scenario, AI would not be a local shock to the political-economic equilibrium, but a phase transition. I focus on extreme scenarios not because I think they are most likely, but rather because they are potentially most consequential. Prudent risk management would devote considerable attention to avoiding such risks.

1.4 WHAT IS DIFFERENT ABOUT THE AI REVOLUTION?

Earlier waves of automation did not end in mass human unemployment and immiseration. Why would the impact of AI be any different? AI possesses a distinctive combination of features that give it even greater potential for economic, political, and social upheaval than prior forms of automation. The coming changes have been billed as a new industrial revolution, but that is not quite correct. The AI revolution is a distinct digital phenomenon that could (i) reorganize the economy, (ii) shift the balance of power both within and between nations, and (iii) radically alter the structure of society, including whom we interact with, whom we trust, and how we acquire meaning and purpose.

To understand how this might happen, we must first identify the features of AI that, when combined, set it apart from past automating technologies. They are:

1. *Automation of cognitive (not physical) tasks.* During the industrial revolution, certain manual tasks were automated that humans are not especially good at. This allowed us to move up the cognitive ladder and specialize in tasks requiring intelligence. In the AI revolution, humans are being outcompeted in cognitive domains which were once their preserve. As AI systems become more capable, human cognitive limitations will begin to bind, making it difficult for displaced workers to shift to higher-valued cognitive activities. There may be no more rungs left on the cognitive ladder for us to climb.
2. *General-purpose technology.* AI already has the characteristics of a general-purpose technology, generating widespread applications, complementary inno-

vations, and driving economic transformation [Eloundou et al., 2024]. As AGI is approached, the range of cognitive tasks that can be automated by an AI system will be much larger than the limited number of physical tasks automated during the industrial revolution. The equivalent technology for automation of physical tasks would be general-purpose mobile robots that can perform a wide range of physical tasks in domestic and industrial settings. The industrial revolution did not come close to such a general-purpose technology.⁶ Because AI will become a critical input into a wide range of cognitive tasks, there is scope for greater market power and concentration than ever before.

3. *Global scalability.* As AI applications will be cloud-based or downloadable anywhere in the world, the market power of AI companies will not be constrained by the geographical location of their operations. It could be that a few leading AI companies, say in Silicon Valley and China, dominate the industry globally. This global market concentration is all the more consequential because AI is not a final product, but is (or will be) a critical factor of production. In short, we could imagine a scenario in which leading AI companies control a large part of the world’s cognitive labor force.
4. *National security threat.* Due to their general-purpose nature, AI systems could discover information and knowledge that threatens national security. Industrial production was not essential to national security except during major wars, which is precisely when factories were nationalized or placed under state control.⁷ The AI sector is different due to its obvious military applications (e.g., AI-powered drones, robot soldiers) and potential for discovery of new sources of energy, propulsion systems, and weapons of mass destruction (WMDs).

Some of these features are shared with other technologies, especially other digital

⁶More recently, physical AI—the merger of AI and robotics—has been transforming industrial production, well before humanoid robots are in widespread use.

⁷For example, Henry Ford famously converted his factories from producing cars to bombers during World War II.

technologies. Specific cognitive tasks have long been automated by devices such as calculators. The internet is a general-purpose technology that has enabled AI development. Existing digital platforms (e.g., Facebook) are globally scalable. Nuclear technology also poses serious national security risks.⁸ What is distinctive is the *coincidence* of these four features. Due to this combination, it is appropriate to entertain the possibility (however unlikely) that AI will have extraordinary consequences for humanity’s future. In particular, we need to account for the possibility that AI grows more capable and sets in motion political-economic feedbacks with explosively bad outcomes. This essay explores how such a scenario might unfold.

1.5 THE POLITICAL-ECONOMIC PROBLEM

At the root of the political-economic risks explored in this essay is a radical shift in the balance of power in society caused by AI. A dominant leader in AI could use its position to extract rents for the benefit of a narrow domestic elite, as well as capture a large share of the global economic surplus by dictating terms internationally. I see this power shift being driven by three forces: massive concentration of economic and political power, breakdown of democratic institutions, and a collapse in human productivity. These forces are described briefly as follows.

- *Economic concentration and political power.* AI firms could acquire unprecedented market power that extends well beyond the AI sector. As AI advances, its general-purpose nature and global scalability will make it a critical technology for a large part of the global economy. Hence, if there are firms with monopoly power over advanced AI programs and applications, they could extract a large share of the surplus from upstream and downstream firms. They could also use AI to start new firms that compete with traditional businesses or acquire these businesses *en masse* for their customers, brands, and physical

⁸AI is like nuclear technology in that it can be used for peaceful purposes (i.e., electricity generation) or for producing WMDs, except that it is a general technology for producing a range of WMDs, including new chemical and biological weapons. See [Aschenbrenner \[2024, IIb\]](#) for a discussion of the security issues surrounding AI.

assets. Due to the pressures of geopolitical competition and need for growth, policymakers may carve out a separate legal and regulatory sphere to supercharge AI development. This would give AI companies a further edge in acquisitions and create an ‘AI black hole’ that absorbs a large part of the global economy. Due to security concerns, these behemoths could become private-public partnerships incorporated into the security state. Thus, ownership and control of the global economy would become extraordinarily concentrated. In the process, the AI sector would become a powerful political actor with the ability to launch the equivalent of national strikes on countries and harm disfavored firms by denying access to critical applications. In this way, market power will translate into broader forms of economic and political power both domestically and globally.

- *Breaking of democratic and egalitarian institutions.* With more capable AI and robots, the mutual dependency between financial capital, human capital, and labor that has been present throughout history will be broken. Productivity, innovation, production scale, and military power will no longer be tied to the human population of a country. Thus, economic and political elites will have incentive to unwind the progressive reforms of the twentieth century (e.g., public education) and there will be no incentive for elites to support proposed solutions to labor displacement such as universal basic income (UBI). There could also be a radical social disembedding in which elites sever ties with the ‘masses’ and retreat into isolated palace households. Democratic and egalitarian institutions built on empathy, national identity, and mutual dependence between classes will break. Indeed, the economic and political elite may see no reason for the masses to exist once it is more efficient to produce the services they want using AI and robots.
- *Human productivity collapse.* As AI automates cognitive tasks, its adoption may (unlike the technologies of the industrial revolution) undermine human cognitive performance. In discussions about AGI, it is typically assumed that

AI capabilities are measured against a *fixed* human benchmark. Even if AI becomes generally more capable than humans, it is thought that humans will have a *comparative advantage* in certain tasks and specialize in those. On the contrary, competition with AI systems may lead humans to reduce their investment in learning new skills and have fewer successful human role models to learn from. In addition, the AI elite could use their political power to direct government funding toward labor-substituting technologies rather than programs that improve human productivity (e.g., health, education). As such, the widespread adoption of AI could generate a collapse in human productivity, rapidly ushering in AGI and human obsolescence in the workforce.⁹ It could become efficient to replace humans with AI and robots, even in areas of human comparative advantage.

I will elaborate on these mechanisms in Section 3. I will discuss here the degree to which my thesis depends on AGI being reached. First, the most catastrophic outcomes—complete human obsolescence in the labor market, breaking of mutual dependence and radical social disembedding by elites, and institutional collapse—likely require full AGI. However, the core elements of the shift in balance of power within and between nations require less, and perhaps only modest advances in AI. Worker displacement and loss of bargaining power, growing market concentration, inequality and erosion of egalitarian and democratic institutions, public finance crises, human deskilling, and AI-driven shifts in geopolitical power could all occur after sufficient automation has taken place, but before AGI and full automation have been achieved.

This essay is part of an emerging literature on the political economy of AI and digital technologies. It builds on work by [Zuboff \[2019\]](#), [Crawford \[2021\]](#), [Varoufakis \[2023\]](#), [Bender and Hanna \[2025\]](#), and [Rikap \[2026\]](#) which views the digital platform econ-

⁹In contrast, [Ray and Mookherjee \[2022\]](#) predict that full automation will be slow as labor is displaced from early automating sectors, lowering wages, and reducing incentives for further automation.

omy as an extractive institution. Regarding AI, [Bender and Hanna \[2025\]](#) believe the branding of statistical systems as “intelligence” legitimizes this extractive institution by making it appear beneficial to human progress. [Kasy \[2025\]](#) examines the market and political power conferred by control over the “means of prediction”—data, compute, expertise, and energy [see also [Agrawal et al., 2022a](#)]. To align AI’s use to social objectives, Kasy recommends democratic control of both AI objectives and the resources and infrastructure needed for the digital economy.

My thesis is not about the existing digital order, but looks ahead to the political-economic tail risks on the path to AGI. At a more advanced stage of AI, the risks are more extreme and require different solutions to the existing problems of algorithmic bias, control, and surplus extraction. The prospective nature of my analysis is shared with work by [Korinek and Suh \[2024\]](#) on the labor market impact of different paths toward AGI, [Bell and Korinek \[2023\]](#) on the threat of AI-induced inequality to democracy, and [Aschenbrenner \[2024\]](#) on the merger between the AI sector and security state.¹⁰ Other similarities and differences will be discussed throughout. Naturally, looking ahead and planning for catastrophic scenarios is an unpleasant task. But it would be unwise to indulge in wishful thinking.¹¹ If AI has great potential to improve the human condition, realizing that potential requires understanding and mitigating the coming political-economic risks.

1.6 A NEW ANALOGY: INDIGENOUS SOCIETIES UNDER COLONIALISM

The correct precedent for the scenario I have set out may not be previous waves of automation, but rather the cultural collapse of indigenous societies caused by European contact and settlement. While the industrial revolution introduced new

¹⁰See also the “world in chains” scenario described by [Ord \[2020\]](#) in which there emerges a global totalitarian regime that can hold power indefinitely.

¹¹For an example of wishful thinking, consider the [Pew Research Center \[2023\]](#) survey conducted less than a year after the release of ChatGPT: 62% of survey respondents believed that AI would have a major impact on the workforce in the next twenty years [[Pew Research Center, 2023](#)]. But only 28% believed their job would be impacted.

technologies, modes of production, and ways of living, the dislocation was not so large as to make the adaptive challenge insurmountable. Over time, societies were able to adapt and thrive in the new environment. Indigenous societies in the colonial era had to make a far greater leap. The new world they faced was alien and largely incompatible with their belief systems, modes of production, and governance structures. The result was technological disruption and cultural alienation [Fanon, 1961, Wolfe, 2006, Kirmayer et al., 2009]. As a consequence, many indigenous societies lost capabilities and institutions built up over millennia.¹² Contact with AI could have similarly catastrophic consequences.

First, European contact caused a demographic shock due to new pathogens (smallpox, measles, influenza, typhus) and, in some places, war [Crosby, 1972]. Modes of production and warfare, social hierarchies, and governance structures are adapted to a society's scale. An extreme spike in mortality rapidly reduced scale and disrupted these vital institutions. In addition, the concentration of epidemic mortality among the old meant that priests, healers, and political leaders holding specialized knowledge were disproportionately affected [Cook, 1998]. This threatened the cultural continuity of indigenous societies.

Second, colonial authorities targeted specific cultural practices they found objectionable. It often turned out that these practices were not peripheral but rather inextricable from core social institutions. By pulling on specific threads, they unraveled the social fabric of indigenous societies. An example is the case of the Purari Delta in Papua New Guinea, where colonial administrators suppressed raiding and ritual cannibalism [Maher, 1961]. This removed practices that had organized male initiation, ceremonial life, and large defensive villages. As a result, large villages dispersed, communal rituals lost force, and the older ritual-political order weakened. The changes were associated with significant depopulation, though disease and mi-

¹²In contrast, White [1991] argues that a balance of power between indigenous nations and European empires existed for around a hundred and fifty years in the Great Lakes region. Instead of a one-sided collapse, there was mutual accommodation and negotiation, with indigenous nations actively shaping the terms of contact.

gration also played a role.

Third, wholesale attacks on indigenous culture and institutions were made in attempts to assimilate indigenous populations. An example is Canada’s Indian Residential School system which forcibly removed 150,000 indigenous children from their families between 1876 and 1997. Students were isolated from parents and community role models. Indigenous languages, religious practices, and social customs were prohibited. The express purpose of the system was to extinguish indigenous culture and family structures. The Truth and Reconciliation Commission of Canada described the outcome as “cultural genocide” [[Truth and Reconciliation Commission of Canada, 2015](#)]. Breaking the intergenerational transmission of indigenous culture led to cultural alienation. Parenting skills were lost, as were community cohesion, and traditional modes of governance. A number of studies document the effects of intergenerational trauma caused by the system in terms of mental health issues, substance abuse, family breakdown, and social dysfunction [e.g. [McQuaid et al., 2017](#), [Gone et al., 2019](#), [Cole et al., 2025](#)].

Fourth, the unleashing of new technologies upended the social order. An example is the samurai after the Meiji Restoration. Samurai were the world’s most skilled swordsmen armed with the greatest swords ever made (*katana*). They had ruled Japan for more than two centuries as a hereditary warrior-bureaucratic elite under the Tokugawa order. However, the introduction of modern rifles and artillery by the Meiji state, and the creation of a national conscript army, meant that they were no match for conscripted peasants. The result was humiliation and defeat at the Battle of Shiroyama (1877), the last stand of the Satsuma Rebellion. As a consequence, their position in society, traditions, and belief systems rapidly disappeared [[Sonoda, 1990](#)].

Finally, indigenous societies were disrupted by extractive colonial institutions, such as land dispossession and forced labor systems. Land is not only the material basis of subsistence for many indigenous societies. Places also encode origin stories, rights

and obligations, seasonal knowledge, and relations among people, animals, and ancestors [Rose, 1996]. In addition, surviving indigenous economic and political structures were repurposed for colonial extraction, especially in densely populated, relatively prosperous regions, including much of Mesoamerica, the Andes, and South Asia [Acemoglu et al., 2002]. Existing political structures were co-opted and redesigned to extract tribute, organize forced labor, and extract resources. These institutions undermined the capacity of indigenous polities to function autonomously, with deleterious long-term effects. Dell [2010] exploits a geographic regression discontinuity around the historical boundary of Peru’s colonial mining *mita* (1573–1812) to show that districts once subject to forced mine labor have around 25% lower household consumption and 6 percentage points more childhood stunting today.

Therefore, the disruption of indigenous societies by colonialism was more intense and wide-ranging than the changes caused by past waves of automation. It was marked by breakdown of traditional norms and governance structures, crises of identity and meaning, and significant loss of capabilities. This disruption was driven by (i) unleashing of radical new technologies and (ii) the control of these technologies by narrow elites that were economically, culturally, and socially separate from the population at large. The same forces are at play in the AI revolution. As such, we should entertain the possibility that advanced AI might lead to cultural collapse, and take steps to avoid it.

1.7 SOLVING THE POLITICAL-ECONOMIC PROBLEM

Anticipating the impact of a new technology is necessarily a speculative exercise, even when grounded in economic history. Nevertheless, with such a novel and consequential technology, it is important to look ahead and think before we leap. Finding solutions to the political-economic problem is more difficult than diagnosing the problem. Given the pace of change and radical uncertainty about its direction, a dynamic approach will be required to limit the tail risks of AI considered here. Rapid institutional adjustment could be guided by the formation of national AI advisory

committees composed of AI researchers and engineers, ethicists, economists, industry leaders, and policymakers, which advise governments on how to collectively adapt to the technological, macroeconomic, and institutional changes that occur during the AI revolution. Models for such advisory bodies include the Office of Scientific Research and Development formed during the Second World War to advise the US government on emerging technologies, such as radar, proximity fuzes, and nuclear weapons. A promising development is the UK AI Economics Institute, a government research unit launched in June 2026 to understand the impact of AI on productivity, employment, and growth. The institute is chaired by economist Simon Johnson and mixes government officials with external specialists from industry, academia, and frontier AI labs.

While the measures needed to limit the risks of AI and maximize its benefits will need to be dynamically adjusted, the following measures are potential starting points for consideration. First, market power in the AI economy must be constrained, ideally without stifling innovation. Thus, competition policy emerges as the critical economic measure. I suggest that competition policy be reoriented from a focus on consumer welfare to limiting the accumulation of economic control. In addition, open AI models should be encouraged to reduce excessive market power by AI companies and boost downstream innovation. Second, democratic and egalitarian institutions must be strengthened to preserve basic freedoms, social bonds, and shared prosperity. One proposal is sortition, where a fraction of legislative representatives and minor officials are chosen at random from the population. Third, all nations who are capable should develop their own sovereign AI systems to maintain independence and preserve an international balance of power. I describe each set of measures in detail in [Section 4](#).

1.8 ELYSIUM AND HADES

Without appropriate measures, the AI revolution could massively concentrate the global economy, undermine democracy, and create a radical separation between a

narrow elite who live in prosperity and peace and the masses who live in an underworld of insecurity and deprivation. In Greek mythology, Elysium is a paradise set apart from the world, the exclusive abode of the gods and of the heroes favored by the gods who are granted immortality. Hades is the realm of suffering and despair. The word is sometimes translated as ‘hell’ in the Christian New Testament. The Roman poet Virgil describes the Elysian fields in book VI of his epic poem the *Aeneid* [Heaney, 2016]:

“They came into happy vistas and the green welcome
Of the Groves of the Fortunate Ones who dwell in joy.
Here a more spacious air sheds brightness
Over the land; they enjoy their own sun here
And their own stars ...”

The depths of Hades are far less salubrious:

“Sounds of groaning could be heard inside, the savage
Application of the lash, the fling and scringe and drag
Of iron chains ...”

Behind the ‘veil of ignorance’—that is, before anyone knows whether they will end up in Elysium or Hades—nobody would choose this kind of outcome. That is why it is important for policymakers to act, not necessarily now, but as soon as they have sufficient knowledge. The problem is that the knowledge required to act may arrive at the same time as the veil is lifted and we learn who might end up in Elysium, and who in Hades.

2 Genesis of the AI Revolution

To understand a future shaped by AI, it is useful to review some basic facts about the evolution of the technology to date. Many readers who are familiar with the history of AI could skip the background sections (2.1 - 2.2) and proceed directly to Section

2.3, which reviews evidence on AI’s early labor market impact. Section 2.4 sets out a case for why governments will not slow down AI development and adoption, despite the risks. Geopolitical competition and domestic institutional crisis will leave them little choice.

2.1 THE EVOLUTION OF AI

To give a sense of the pace of progress, Table 1 presents a timeline of developments in AI this century. Things are moving fast. (By the time you are reading this, the Table will need to be updated.) Progress has come from massive capital investment, especially in the United States following the release of ChatGPT in November 2022. Corporate investment in AI was \$581.7 billion in 2025, an approximately fortyfold increase since 2013 [Stanford Institute for Human-Centered Artificial Intelligence, 2026]. Government financing of AI also expanded over this period. Between 2013 and 2024, the United States invested approximately \$20.4 billion in AI-related contracts and grants, while European nations committed about \$3.7 billion in AI-related contracts [Stanford Institute for Human-Centered Artificial Intelligence, 2026]. Apart from the massive mobilization of capital, AI development has been driven by three factors: a shift to different AI models, rapid growth in computing power (‘compute’), and big data.

In 1996, when IBM’s Deep Blue beat then world champion Garry Kasparov at chess, it relied on rules programmed into it by expert chess players. By 2010, progress in these symbolic AI models had stalled, leading to widespread pessimism about the capabilities of AI. The huge leaps in AI in recent years were unlocked by a different approach called *deep learning* based on (multi-layered) neural networks [see Mitchell, 2019]. In 2016, when Google DeepMind’s AlphaGo program beat Lee Sedol, the world’s number one player of Go (an ancient game that is more complex than chess), it was not programmed by expert Go players. Instead, it learned how to play Go by watching past matches between strong players recorded on online Go sites. After that, it played against itself tens of millions of times. In this way, AlphaGo mastered

Table 1. A Brief Timeline of Developments in AI since 2000

2000	Kismet robot head developed (affective computing)
2002	iRobot releases first Roomba robot vacuum cleaner
2004	Spirit and Opportunity rovers autonomously navigate Mars
2005	ASIMO robot can walk at a human pace and deliver trays in a restaurant
2009	Google starts testing self-driving cars in San Francisco
2011	Watson wins Jeopardy!; Siri launched
2012	AlexNet scores a sweeping victory in the ImageNet challenge and starts the deep learning revolution; Google Now launched
2014	Cortana launched; GAN architecture invented
2015	Highway network and ResNet introduced; open letter calling for research on societal impacts of AI
2016	AlphaGo defeats Go champion Lee Sedol 4-1
2017	Transformer architecture released
2019	OpenAI releases GPT-2
2020	GPT-3 released; AlphaFold 2 wins CASP
2021	GitHub Copilot released
2022	ChatGPT (GPT-3.5) released; DALL-E 2 launched; Midjourney launched; Stable Diffusion released
2023	OpenAI: GPT-4, GPT-4V, DALL-E 3; Anthropic: Claude, Claude 2, Claude 2.1; Google: Bard, Gemini 1.0; Meta: Llama, Llama 2, Code Llama; open letter calling for six-month AI development pause
2024	OpenAI: GPT-4o, o1, Sora; Anthropic: Claude 3 (Opus, Sonnet, Haiku), Claude 3.5 Sonnet, Claude 3.5 Sonnet v2, Claude 3.5 Haiku; Google: Gemini 1.5; Meta: Llama 3, Llama 3.1, Llama 3.2 (multimodal) xAI: Grok-2, Aurora; DeepMind: AlphaProof and AlphaGeometry 2 win silver-medal equivalent at IMO; Black Forest Labs: Flux.1 launched; Apple: integrates ChatGPT into Siri; Chinese LLM DeepSeek-V3 matches GPT-4o
2025	OpenAI: o3-mini, GPT-4.5, GPT-4.1, o3, o4-mini, Operator Agent, Deep Research, Codex, GPT-5, GPT-5 Pro, gpt-realtime, ChatGPT Atlas, Sora 2, GPT-5.1, ChatGPT Pulse, GPT-5.2; Anthropic: Claude Sonnet 3.7, Claude Opus 4 and Sonnet 4, Claude Opus 4.1, Claude Sonnet 4.5, Claude Haiku 4.5, Claude for Life Sciences, Claude Opus 4.5; Google: Gemini 2.5 Pro and Flash, Veo 3.1, Gemini 3; Meta: Llama 4 (Scout and Maverick); xAI: Grok 3, grok-code-fast-1, Grok 4, Grok 4 Fast, Grok Imagine 0.9, xAI for Government; DeepSeek: R1, V3.1, V3.2; Research: UC San Diego's universal deepfake detector
2026	OpenAI: GPT-5.4, GPT-5.4-Pro, GPT-5.5, GPT-5.6, GPT-6 Astra; Anthropic: Claude Opus 4.6, Claude Opus 4.7, Claude Opus 4.8, Claude Fable 5, Claude Sonnet 5, Claude Opus 5; Google: Gemini 3.1, Gemini 3.5 Flash, Gemini 3.6 Flash, Gemini 3.7 Flash; xAI (now SpaceXAI): Grok 4.20, Grok 4.5, Grok 4.6; Moonshot AI: Kimi K2.5, Kimi K2.6, Kimi K3; DeepSeek: V4-Flash, V4-Pro; Meta: Muse Spark 1.1, Muse Spark 1.2

and surpassed thousands of years of knowledge about Go in a matter of months.¹³ Since 2016, the way humans play Go has been revolutionized by strategies played by AlphaGo. For a brief period, the machines learned from us. Now we learn from them.

Neural networks have been studied since the 1940s [McCulloch and Pitts, 1943, Rosenblatt, 1958]. Deep neural networks, with multiple hidden layers, have been around since at least the 1980s [Rumelhart et al., 1986, Krizhevsky et al., 2012]. What then explains the recent breakthroughs? Though there have been further technical advances such as the Transformer architecture [Vaswani et al., 2017], progress in deep learning has been largely a matter of scaling up existing models with massive amounts of data and compute, both of which have only recently become available. Current models have around one trillion connections (weights) between their simulated neurons. The scale of investments being made is based on the expectation that each time a model is scaled up, new capabilities emerge. The massive demand for compute has been a huge boon for advanced chip makers, especially Nvidia whose market capitalization rose from \$338 billion on 1 November 2022, before the release of ChatGPT, to over \$5.3 trillion by late August 2026. The second driving force, data, has been described as ‘the new oil’ by Kai-Fu Lee, former President of Google China [Lee, 2018]. The real catalyst for big data has been the internet and the digital platforms built on top of it. Without the records of Go matches played online, the training of AlphaGo would not have been possible in such a short time-frame.

2.2 POTENTIAL BOTTLENECKS

How far will AI progress in the coming years? For AI to produce a growth takeoff, there must be complete automation across all critical tasks (i.e., AGI) and no bottlenecks [Aghion et al., 2019b]. Otherwise, Baumol’s cost-disease will kick in and

¹³Its 4-1 victory over Lee Sedol is brilliantly captured in the documentary *AlphaGo* [Kohs, 2017]. Lee was not only beaten by the machine on the board. Eight months later, he announced his retirement, stating that even if he were to remain the top human player, “there is an entity that cannot be defeated”.

limit growth [Baumol, 1967].¹⁴ Currently, there are a number of significant obstacles to further AI progress. First, transistors are already getting so small (2 nanometers, roughly the size of 20 silicon atoms) that they run into quantum tunneling problems [Sze et al., 2021]. Second, key parts of the semiconductor supply chain are fragile. Around 92% of the world’s advanced processors (less than 10 nanometers) are manufactured in Taiwan by one company, TSMC. The remainder are produced by Samsung. The equipment for printing these chips using extreme ultraviolet (EUV) lithography is made by only one company located in the Netherlands, ASML. Third, data sources are being exhausted. The world wide web has been fully exploited to train AI programs. In addition, AI models trained on AI-generated content degrade rapidly [Guo et al., 2023]. Fourth, further advances in AI may require different model architectures and technical innovations, including new pre-training methods, agentic systems, and enhanced reasoning capabilities [LeCun, 2022, Vincent, 2024]. Fifth, new sources of energy and upgrades to electricity grids are required to power more data centers. Sixth, the commodities such as copper, silver, tungsten, and uranium needed for advanced AI and robotics are currently experiencing structural supply deficits. Sixth, sustained technological progress tends to require an elusive institutional configuration: decentralized exploration plus centralized allocation of capital to scale promising discoveries [Frey, 2025]. Finally, valuations of AI companies are extraordinarily high and a large correction could reduce capital expenditure required for rapid progress.

Some of these problems are more easily solved than others. Every cloud hyper-scaler is developing its own proprietary chips for running AI workloads. China has begun limited production of Deep Ultraviolet (DUV) lithography machines for printing advanced chips. On the data side, there are undigitized archives and privately owned databases with high-quality human-generated data, e.g., Tesla’s data from self-driving vehicles. The robots being developed will also provide high-resolution

¹⁴Aschenbrenner [2024] predicts just this, that an impending series of innovations or ‘unhobbling’ will remove all bottlenecks. For an alternative view, see Gans and Goldfarb [2026].

data about the 3-D world. Companies, such as Yann LeCun’s Advanced Machine Intelligence Labs, are experimenting with new model architectures, including world models, state-space models, and diffusion language models. Perhaps the most intractable problem is the supply of critical commodities, due to delays in mine development and permitting and the loss of refining capacity outside of China. Given the extent of government and corporate buy-in to the AI revolution, we should expect vigorous attempts to solve these problems and achieve AGI.

Another possibility is that we fail to achieve AGI not because AI progress stalls, but rather because humans lift their game. Recall that human Go players began to study the moves of AI systems in 2016. Using computers, they discovered some weaknesses in AI play. In 2023, a Go player beat a top AI program by exploiting these weaknesses [Waters, 2023]. If human performance improves in lockstep with AI performance, AGI will never be reached. I will present reasons, however, why the exact opposite may end up occurring (see Section 3.3). Without AGI, the most destructive economic, political, and social forces I describe in this essay would be avoided. Yet serious political-economic risks, such as extreme market concentration, shifts in the balance of power within and between nations, and human deskilling, require only widespread adoption of existing or imminent AI systems.

2.3 THE EARLY LABOR MARKET IMPACT OF AI

Since the scenario I explore depends on AI substituting for human labor and depressing employment and wages, it is important to check whether there are early signs of this occurring.¹⁵ While AI is not yet replacing entire occupations, frontier models can perform an increasingly wide range of tasks, including some human-centered relational tasks. For example, Jabarian and Henkel [2025] conduct a field experiment in which 70,000 job applicants were randomly assigned to either a human recruiter, an AI voice agent, or given a choice between the two. Compared to a human recruiter,

¹⁵For a range of concerns expressed by economists, see Brynjolfsson and McAfee [2014], Susskind [2020], and Deming et al. [2025].

interviews led by the AI voice agent had a 12% higher job offer rate, an 18% higher job start rate, and a 17% higher 30-day retention rate. Of the applicants who had to choose, 78% of applicants chose the AI recruiter. There was no significant difference in reported satisfaction by interviewees. In addition, there is evidence that AI can act as a complementary technology, particularly in customer support [Brynolfsson et al., 2023] and coding [Peng et al., 2023].¹⁶

According to the Acemoglu and Restrepo [2019] framework, AI’s impact on aggregate employment and wages depends on how well AI substitutes for human labor (*displacement effect*), how much it increases productivity at the aggregate level (*productivity effect*), and the rate at which new tasks are created for displaced human labor in which humans outperform AI (*reinstatement effect*). An early (up to 2018) empirical study by Acemoglu et al. [2022] finds that while likely AI adopters in the US reduced hiring for non-AI positions, AI did not reduce employment or wage growth at the occupation or industry level. More recently, Humlum and Vestergaard [2025] find similar null results when linking data on AI chatbot adoption to administrative labor data in Denmark. AI adoption led to occupational switching and task restructuring, but this had no effect on earnings and recorded hours at both the employee and firm levels.¹⁷ In France, Aghion et al. [2025] find positive results: AI adoption increased aggregate firm-level employment and sales. Employment increased even in occupations deemed as being at risk of automation by AI. Hence, Aghion et al. [2025] conclude that the productivity effect dominates the substitution effect of AI in their French data. In addition, Hoffmann et al. [2025] find preliminary evidence of a reinstatement effect. Exploiting an eligibility threshold for GitHub Copilot, they show that access to a code-completion assistant shifts task allocation to novel coding

¹⁶Note that complementarity can be consistent with labor displacement as less human labor would be required to complete a *given* set of tasks.

¹⁷de Souza [2025] presents evidence on occupational restructuring in Brazilian factories. Adoption of industrial software to monitor machinery and optimize maintenance meant machines could be operated continuously. This reduced employment in routine office tasks, while raising demand for manual workers who operated the machinery. The net result was an increase in employment, but a fall in wages due to the composition effect.

activities. In contrast, [Teutloff et al. \[2025\]](#) study a leading online platform and find that demand for freelancers who specialize in AI-substitutable tasks, such as writing and translation, decreased by 20–50% after the launch of ChatGPT.

The clearest evidence of an overwhelming displacement effect is for junior workers. [Brynjolfsson et al. \[2025\]](#) estimate that the widespread adoption of generative AI reduced employment of US workers aged 22-25 in AI-exposed occupations by 16%. The reduction is found primarily in occupations in which AI substitutes for, rather than complements, human labor. In contrast, employment of experienced workers was stable. Changes in wages were minor relative to the reduced hiring of entry-level workers. [Lichtinger and Hosseini \[2025\]](#) report similar results using US resume data covering 62 million workers across 285,000 firms from 2015–2025. They identify firms adopting AI by tracking “GenAI integrator” job postings and find that adopting firms significantly reduced junior employment, primarily by reducing junior hiring. Again, there was little change in senior employment.¹⁸

In recent work, [Acemoglu \[2026\]](#) concludes that AI’s current effect on the labor market is primarily to displace workers, rather than creating broad productivity gains or sufficient new tasks.¹⁹ My expectation is that as AI advances and becomes more agentic, it will come to substitute for human labor in more and more tasks, making the displacement effect even more dominant. Even jobs with a strong bundling of automatable and unautomatable tasks will be rebundled in a way that reduces rents to human workers [see [Garicano et al., 2026](#)]. Eventually, of course, there could be nearly full automation. Consider the vision of agentic entrepreneurs proposed by Mustafa Suleyman, CEO of Microsoft AI:

“I have proposed a test which involves them [AI systems] going off and taking a \$100,000 investment and over the course of three months trying to set about creating a new product, researching the market, seeing what

¹⁸See [Autor and Thompson \[2025\]](#) on the wage and employment effects of this expertise bias.

¹⁹[Acemoglu and Johnson \[2023\]](#) also argue that AI can be used to monitor workers more intensively, further lowering wages and worker welfare and creating a new age of authoritarian capitalism.

consumers might like, generating some new images, some blueprints of how to manufacture that product, contacting a manufacturer, getting it made, negotiating the price, dropshipping it, and then ultimately collecting the revenue. And I think that over a five-year period [by 2028] it is quite likely that we will have an ACI—an artificial capable intelligence—that can do the majority of that task autonomously.” [The Economist, 2023]

Such a shift in the labor market would have profound political-economic consequences.²⁰

2.4 WHY NOT SLOW DOWN AI PROGRESS?

In March 2023, Elon Musk, Steve Wozniak, Yuval Noah Harari, and others signed an open letter calling on AI labs to pause training of AI systems more powerful than GPT-4 for at least 6 months [Future of Life Institute, 2023]. That did not happen. In September 2026, the CEOs of Anthropic and OpenAI, Dario Amodei and Sam Altman, both called for slowing the pace of frontier model development, due to concerns about uncontrollable agent swarms and recursive self-improvement [Amodei, 2026, Altman, 2026]. However, I doubt that this will happen. According to the *precautionary principle*, we should exercise caution when making irreversible decisions that can cause harm. Obviously, profit-maximizing firms may pursue innovations that are privately profitable but socially harmful. But why might policymakers not wish to pause or otherwise contain AI development? There are two reasons.

The first reason is geopolitical competition, especially between the United States and China. As in the twentieth century, the winner will be the bloc with the most advanced technology. The United States is the leader in frontier AI models, but China leads in AI patents and industrial robots [Stanford Institute for Human-Centered Artificial Intelligence, 2026].²¹ The US lead in LLMs is not unassailable. In De-

²⁰For example, see Korinek and Suh [2024] on the effect on wages of different paths to AGI.

²¹Note that patent counts could be misleading as they do not adjust for quality. Private invest-

cember 2024, the tech world was shocked by the release of DeepSeek-V3, an open-weight model that exceeded the performance of US-produced open-weight LLMs (e.g., Meta’s Llama3.1-405B). Remarkably, it was claimed that DeepSeek-V3 was trained on a budget of \$5.58 million over 55 days, probably less than one tenth of what was spent by the US companies. Moreover, unless DeepSeek found a way to circumvent export restrictions under the US CHIPS Act, it was trained using less advanced Nvidia chips. In January 2025, DeepSeek released an advanced open-weight reasoning-focused LLM trained using reinforcement learning with verifiable rewards (RLVR). In July 2026, Chinese lab Moonshot AI released the largest open-weight model ever, Kimi-K3. This release significantly closed the gap to the most advanced US models, Anthropic’s Claude Fable 5 and OpenAI’s GPT-5.6 Sol.

The second reason why AI development will not be slowed is perhaps less obvious. The West faces an institutional crisis and technology is viewed as a way out. The industrial revolution created unprecedented wealth. In addition, literacy, heights, caloric intake, life expectancy, poverty alleviation, social tolerance, and political freedoms all improved in lockstep with GDP [Fogel, 2010, Weil, 2014]. However, the transition from feudalism was far from painless. The safety nets of the feudal order vanished as old notions of *noblesse oblige* and communal insurance came to an end. In *The Communist Manifesto*, Marx and Engels were particularly moved by the pitiful state of the handloom weavers in England, whose trade was the first to be automated [Marx and Engels, 1988/1848]. Due to the threat of socialist revolution, combined with the greater bargaining power of a more skilled population, western nations forged a new social democratic compact with their citizens, beginning with the expansion of voting rights [Hobsbawm, 1987, Acemoglu and Robinson, 2000, Tilly, 2007, Tilly and Tarrow, 2008]. Over the twentieth century, the familiar institutions of the modern nation were created, including the welfare state, mass public

ment in AI has been far higher in the United States, \$757.3 billion since 2013 versus \$131.8 billion in China. However, China allocated \$184 billion in government guidance-fund capital to AI firms from 2000–2023, which is far higher than US public AI contracts and grants [Stanford Institute for Human-Centered Artificial Intelligence, 2026].

education, public healthcare, and government regulation of markets.

Most of these legacy institutions are now strained. Social indicators no longer move in lockstep with economic growth. For example, life expectancy in the US and UK has started falling [Case and Deaton, 2020, Lima and Snower, 2020]. Moreover, it is increasingly difficult to finance the welfare state due, *inter alia*, to a productivity slowdown [Fernald et al., 2025, Goldin et al., 2024] and fertility decline [Schumacher et al., 2024].²² US federal government debt is over \$40 trillion at the time of writing, or over 120% of GDP.²³ Debt-to-GDP is over 100% in Japan, France, Canada, Italy, Spain, and Belgium. Hence, the economically advanced nations face a predicament. Their welfare states are strained, but cannot be dismantled without political upheaval. The only way out of crisis then is to find a new engine of economic growth. Hence, AI is viewed by important factions within the public and private sectors as a technological lifeline, a route out of fiscal and sovereign debt crisis [Rogoff, 2026, Harris et al., 2026, Musk, 2026]. There is no guarantee, however, that it will deliver. If the political-economic risks of AI are not mitigated, there will be no boost to shared prosperity, and the institutional crisis will only deepen.

3 The Political-Economic Risks of AI

This section fleshes out the catastrophic political-economic risks on the path to AGI. The underlying theme is that AI will radically alter the balance of power within and between nations. My focus is on the impact of AI on ‘the West’, that is, the coalition of democratic nations including the US, UK, EU members, Canada, Australia, New Zealand, Japan, Singapore, and South Korea. How events will unfold in China

²²To the contrary, Offer [2022] argues that public pension systems could cope due to their flexible terms and pay-as-you-go structure by, for example, raising the retirement age or reducing benefits. Japan has already implemented such reforms.

²³In addition, the US Treasury Department estimates that unfunded spending obligations will amount to \$79.5 trillion over the next 75 years (in present value terms). Social Security and Medicare shortfalls make up 95% of the total unfunded obligation [US Department of the Treasury, 2024].

is an important and fascinating subject, but one that I leave to China specialists [see, e.g., [Lee, 2018](#)]. The political-economic trajectory I trace in this section is not a conspiracy, but rather an evolutionary path to a dystopian future. I think the massive accumulation of economic and political power along this path will be less by grand design than through myopic step-by-step adaptation to the AI economy. The road to Elysium may even be paved with good intentions. But the consequences will be even worse than those envisioned by F.A. [Hayek \[1944\]](#) in his *Road to Serfdom*. That is why it is important that we look ahead.

3.1 MARKET CONCENTRATION AND POLITICAL POWER

Due to AI’s distinctive combination of features, new economic and political elites could emerge from the AI revolution with novel sources of power. One way this could happen is if the AI sector ‘swallows up’ large parts of the global economy. I call this ‘the AI black hole’. I identify three catalysts, each of which is sufficient for AI to produce extreme market concentration. I then set out how this newfound market power could be leveraged into political power, both domestically and internationally.

3.1.1 *Catalysts for Extreme Market Concentration*

The catalysts are all conditions that give key firms in the AI sector a durable competitive advantage, or ‘moat’. That is the foothold required to accumulate more extreme forms of market power.

Unassailable AI Lead. Frontier AI labs in the US have been spending unprecedented amounts in a race to be the first to reach AGI. Once AGI is reached, it is often thought that ASI will soon follow due to recursive self-improvement [[Bostrom, 2014](#), [Yudkowsky and Soares, 2025](#)]. If an AI lab or cartel of labs develops an unassailable technological lead, then this would confer massive economic control and political power. Such a lead could come from proprietary discoveries (e.g., new model architecture) or because scaling existing models to achieve and operate AGI/ASI requires such intensive capital investment that no competitor could afford it. The likelihood

of an unassailable AI lead was revised downward after the release of DeepSeek V3 and R1, which economist Olivier Blanchard described as: “Probably the largest positive tfp [total factor productivity] shock in the history of the world” [Blanchard, 2025]. DeepSeek’s success suggested that massive capital expenditure and proprietary knowledge are not needed to match the *current* performance of frontier models. In addition, the release of Moonshot’s Kimi-K3 demonstrated the advanced capabilities of Chinese AI labs and their commitment to open models.²⁴ If AGI turns out to be inexpensive to develop/run and open models proliferate, then AI would be ‘democratized’. Anyone could build applications on top of AI. That is not guaranteed, however, as the technology advances toward AGI. Reaching AGI and ASI may prove impossible without new technical discoveries and massive amounts of compute and energy. In addition, running AGI could require continuous training and a lot of inference compute, that is, resources for making predictions or decisions based on novel data after the model has been trained.

AI-powered Platforms. If the leading AI companies turn out to have no moat, then profits could flow to the application layer. During the Dotcom boom in the late 1990s, companies supplying internet infrastructure such as Cisco Systems were expected to be massively profitable. Instead, the companies that capitalized on the internet turned out to be platforms such as Google, Facebook, and Amazon, which benefited from network externalities and other lock-in effects. AI could supercharge the existing digital platforms through personalization and improved content curation. In addition, new platforms could emerge that deliver AI-powered applications. For example, Palantir aims to be the future enterprise operating system by building custom software for organizational coordination, decision-making, and execution on top of generative AI. The potential for market concentration in AI applications could be larger than in the existing digital economy, as firms that gain a slight competitive edge will have few labor constraints on rapid scaling.

²⁴Even if advanced capabilities were achieved through ‘distillation’ of US models, it would still weigh against an unassailable AI lead.

Government-Tech Merger. Another way in which firms in the AI sector could be shielded from competition is by becoming entangled with government. That is, a competitive moat could be created for the leading AI companies by fiat. First, due to the pressures of geopolitical competition and the desperate need for growth, policymakers will have a strong incentive to switch from prioritizing AI safety to removing bottlenecks to further AI progress (see Section 2.2), including:

- Land use restrictions for building data centers.
- Energy regulation limiting power to data centers.
- Intellectual property and data protection laws limiting new sources of data.

A parallel economic sphere could be created with special exemptions for AI companies from land use restrictions, energy regulation, and intellectual property and data protection laws. Otherwise, the leading AI companies in the West will either start to lag behind China or look to build data centers rapidly in non-democratic countries [Aschenbrenner, 2024]. Both outcomes pose obvious geopolitical risks.

The next step is for governments to step into this parallel AI sector and become more involved in its operations, i.e., financing AI development, restricting competition, and providing security. In financing the AI sector, governments may begin to pick winners and restrict competition to better control the pace and direction of innovation. Competition in the AI sector would erode profits and reduce incentives for innovation [Aghion et al., 2005, Gilbert, 2006, Vives, 2008, Spulber, 2013]. Frontier US labs may also seek protection from cheaper open-weight models. In addition, governments will want to maximize the capability of the leading AI models, which are contenders to reach AGI first.²⁵ In contrast, if there are diminishing returns to AI development and widespread applications for less advanced AI systems, financial markets will tend to fund emerging companies that enjoy low-hanging fruit. These

²⁵An example is the \$500 billion Stargate joint venture to build AI infrastructure in the United States, led by OpenAI, Oracle, SoftBank, and Microsoft, and announced at the White House on 21 January 2025 (the day after inauguration).

companies would not advance the technological frontier as geopolitical goals dictate. Finally, as set out by Leopold Aschenbrenner, governments may need to provide state-level security to leading AI companies [Aschenbrenner, 2024, IIIb.]. If AGI is achieved, there will be very strong incentives for nefarious actors to steal the models.²⁶ AI models will need to be secured not only because they boost economic performance, but also because they could assist rogue states and non-state actors in developing advanced military technologies and weapons of mass destruction (WMDs), including potentially humanity-ending bioweapons and nanotechnology. The Manhattan Project would have been futile without extreme security measures, including its remote base in Los Alamos, New Mexico. Hence, a few AI companies could be brought into the existing public-private sphere of intelligence agencies, military, and defense contractors. Others could be closed down, and new entrants blocked. Thus, an AI complex could coalesce.

3.1.2 *The AI Black Hole*

Market power in the AI sector is just the start of the problems. The real problem is the extension of this market power far beyond the AI sector. In particular, firms with a dominant position in the AI sector could extract profits from upstream and downstream firms and acquire large swathes of the economy. The outcome would be an ‘AI black hole’ that swallows up a large part of the global economy.

Given market power in the AI sector, leading AI companies may be able to extract the entire surplus from the value chain. Consider a simple model with one AI producer M and two downstream firms N_1 and N_2 that use AI to produce a (homogeneous) final good. M produces each token at cost c and charges price p . N_1 and N_2 can convert each token into a unit of output at (additional) cost k . Consumers are willing to pay up to v for the final product. (Assume $v > c + k$.) Then competition between

²⁶While stealing current models would not give a competitor AGI, Aschenbrenner points out that the algorithmic methods for training AGI (and possibly ASI) are already being developed and stealing this knowledge could unlock a competitor’s path to AGI.

N_1 and N_2 in the final good market will push price down to $p + k$. Competition by N_1 and N_2 for AI services will drive the price of a token up to $p = v - k$. That means consumers receive zero surplus and firms N_1 and N_2 make zero profit. All of the surplus is extracted by the AI firm M which makes a profit of $v - k - c$ on each token.²⁷ AI firms could extract even more surplus from the vertical chain through hold up [Williamson, 1975, Klein et al., 1978, Grossman and Hart, 1986].²⁸ Suppose an organization contracts with an AI lab or AI-powered platform to restructure its workforce, workflows, and decision-making processes to make intensive use of AI. Once the initial investment is sunk and the organization has been restructured, it will be heavily dependent on AI. In fact, it could be locked in to the services of a particular AI lab or platform, which could then significantly raise the price charged for its services.

More generally, the conditions under which an AI company can achieve full surplus extraction is (i) when it is a *de facto* monopoly in the AI layer and (ii) AI is a critical technology that allows downstream firms to outcompete rivals who do not use AI. Even with more than one AI firm, monopoly power could be exercised due to lock-in effects or collusion among AI labs/platforms [e.g., Friedman, 1971, Rotemberg and Saloner, 1986]. Collusion could be easy to achieve if the number of AI firms is limited or if they are exempted from antitrust provisions. Given that AI is a general-purpose technology that is globally scalable, the number of downstream firms from which they could extract is extraordinarily large. The prime candidates are firms that focus on repetitive cognitive tasks, such as accounting, clerical work, customer support, sales and marketing, and consulting. Advances in AI-powered robots would significantly expand the set of industries whose profits could be extracted by the AI sector to those industries that make intensive use of physical labor.

²⁷More generally, surplus extraction can be achieved using a two-part tariff: the upstream firm sets the per-unit wholesale price equal to its own marginal cost, so that the downstream firms, in aggregate, produce the monopoly quantity. Then the upstream firm charges a fixed fee equal to the entire downstream profit.

²⁸Another anti-competitive strategy is exclusive deals with downstream firms [Rasmusen et al., 1991, Bernheim and Whinston, 1998].

The second way in which the AI sector could swallow up large parts of the global economy is through vertical integration. AI companies could buy up land, energy supply, chip manufacturers, AI-powered platforms, and a range of regular businesses.²⁹ There are a number of reasons why they would have an advantage in acquisitions.³⁰ As AI is a globally scalable general-purpose technology, AI companies may be able to outcompete traditional businesses in their own industries by, for example, automating back-end operations.³¹ Hence, AI companies could start firms that compete with traditional businesses or take over these businesses *en masse* to acquire their customers, brands, and distribution networks. The mega-corporations that would arise could use AI to avoid the usual coordination problems, culture clashes, and diseconomies of scale that come with large workforces. In addition, if a government-tech merger were to occur, leading AI firms would enjoy a lower cost of capital due to government backing, legal and regulatory exemptions, and state-provided security. Finally, vertical integration would enable AI firms to control a larger set of critical inputs and engage in input foreclosure, harming downstream firms by raising prices, lowering quality, or refusing to sell to them [Mathewson and Winter, 1984, Perry, 1989, Rey and Tirole, 2007]. All of these factors would drive acquisitions by AI firms, and could suck a large share of the global economy into the AI sector.

Peter Temin has proposed a dual-economy view in which the US economy has split into high- and low-wage sectors, with middle class jobs vanishing [Temin, 2016, 2017]. The scenario explored here is a far more advanced stage of this dual economy. For the foreseeable future, there will be a traditional economy bound by the laws and regulations developed over the twentieth century to strengthen competition, safe-

²⁹Note that there is already a significant amount of vertical integration among the AI hyper-scalers. For example, Amazon, Google, Microsoft, and Meta own massive amounts of data through cloud storage and social media platforms such as Facebook, Instagram, and YouTube.

³⁰In general, vertical integration could improve efficiency by eliminating the double-marginalization problem in the vertical chain [Spengler, 1950, Grossman and Hart, 1986, Hart and Moore, 1990].

³¹One reason why a firm that gains a slight competitive advantage today does not take the entire market is capacity constraints [Coase, 1937, Kreps and Scheinkman, 1983]. AI will ease scaling constraints.

guard workers, and protect the environment. However, this old economy will make up a smaller and smaller fraction of the new economy. Thus, economic activity could become concentrated in a new economic sphere powered by AI, possibly outside of the existing institutional framework. As AI advances, this economic transformation could happen in plain sight, but without much political attention or opposition.

3.1.3 *Shifting Balance of Power*

The classical economists understood very well that market power translates into political power, and *vice versa*. In Book V of *The Wealth of Nations*, Adam Smith writes:

“It is in the age of shepherds, in the second period of society, that the inequality of fortune first begins to take place, and introduces among men a degree of authority and subordination which could not possibly exist before. It thereby introduces some degree of that civil government which is indispensably necessary for its own preservation: and it seems to do this naturally, and even independent of the consideration of that necessity. The consideration of that necessity comes no doubt afterwards to contribute very much to maintain and secure that authority and subordination.” [Smith, 2008/1776, p. 689]

This passage introduces Smith’s theory of the evolution of property rights. When technological change (in this case, pastoralism) creates economic inequality, political and social institutions emerge that protect property rights and entrench the inequality. Likewise, the market concentration and wealth the AI revolution could create for a narrow elite will bring a radical shift in the balance of power in society. Let us first consider the immediate and milder forms of political power that will emerge, before getting to more fundamental institutional change.

In the classic Meltzer-Richard framework, a rise in inequality leads to *more* redistribution under majoritarian voting [Meltzer and Richard, 1981]. However, this

framework does not account for the role of money in politics in the form of political donations, patron-client relationships, and media influence [e.g., [Ansolabehere et al., 2003](#), [Bartels, 2008](#), [Gilens, 2012](#)]. The Meltzer-Richard outcome could be reversed if, instead of majoritarian voting, policies are determined by a weighted average of citizens’ utilities, with greater weight placed on citizens with higher income. Then as the relative income of the rich rises, policies could shift toward *less* redistribution and, more broadly, toward the interests of the rich [see, e.g., [Grossman and Helpman, 1996](#)]. In a paper aptly titled “One Dollar, One Vote”, [Karabarbounis \[2011\]](#) presents evidence of such an effect from a panel of advanced OECD countries. He finds less redistribution where there is a more affluent rich and middle class and more redistribution where the poorer class is richer.³²

Key features of the AI revolution are likely to drive even greater leverage of wealth and market power into political influence. First, human displacement from the workforce will lead to a sharp fall in income tax revenue at precisely the time at which demand for welfare payments spikes. To finance its core functions, the government will need to ‘strike a deal’ with the AI sector to shift the taxation system toward corporate and capital taxation. As part of this deal, the AI sector could acquire significant power over what gets financed. Second, AI profits would not only finance persuasion activities [see [Skaperdas and Vaidya, 2012](#)], but the technology itself would power new and more effective forms of censorship and persuasion [[Zuboff, 2019](#), [Acemoglu and Johnson, 2023](#)].³³ A supercharged version of existing ‘informational autocracies’ could prevail [[Guriev and Treisman, 2019, 2022](#)]. As such, citizens may come to accept, as Smith writes, ‘a degree of authority and subordination which could not possibly exist before’. This is the thrust of Thomas Piketty’s *Capital and Ideology* which examines the role of ideology in preserving the *status quo* distribution of wealth and power [[Piketty, 2020](#)]. Finally, as AI becomes a critical factor of production, disfavored individuals and organizations could be coerced by denying

³²In contrast, [Levitt \[1994\]](#) finds little evidence that campaign spending affects election outcomes.

³³See [Carvalho et al. \[2024\]](#) for historical examples of failed cultural persuasion.

them access to the leading AI models. Non-compliers might find it hard to identify and prove that the AI agents they are using have been handicapped or misdirected. Thus, the political influence of the AI complex and those connected to it could grow and extend beyond any exemptions granted to AI companies during their breakout phase. In turn, this political power could be used to further concentrate profits and expand the AI black hole. The extreme concentration of individual and corporate power would grow even more dangerous if co-opted by an authoritarian state.

AI will not only shift the balance of power within nations, but also between nations. At present, there are two AI powers, the US (concentrated in Silicon Valley) and China [Lee, 2018]. Though other centers of AI development exist, notably London, Paris, Toronto, and Tel Aviv, western nations are mostly served by US-based AI companies. If AGI requires massive capital expenditure, countries such as the UK, the EU members, Canada, Australia, New Zealand, Japan, Singapore, and South Korea may be tempted to free ride on AI innovation in the US, just as they (partially) rely on the US for defense, information technology, and medical innovation. When it comes to AI, that would tip the balance of power among the western nations very far toward the US.

Domination of the supply of artificial cognitive (and possibly physical) labor, will give the US considerable leverage over the other advanced economies. A US-based AI complex will be able to launch labor strikes in other countries and squeeze foreign firms through various forms of vertical foreclosure, e.g., denial of access to critical technologies [Rey and Tirole, 2007]. AI systems in the US may also produce new military and cyberattack capabilities and WMDs, along with new defensive measures against these threats. Hence, without their own investments in AI, the rest of the western nations will rely to an unprecedented degree on the US for their security and defense. Developing countries will have even less ability to build their own advanced AI systems and robots and may find themselves with vanishing international bargaining power, outside of control over critical resources such as rare minerals.

Finally, the AI revolution will place extraordinary stress on public finances outside of the US. As in the United States, part of the burden from lost income taxation will have to fall on corporate taxation. The difference is that taxation of US-based companies by foreign governments will be constrained by various tax avoidance strategies. [Tørsløv et al. \[2023\]](#) find that around 36% of multinational company profits are shifted to tax havens worldwide and that US multinationals shift about twice as much, relative to their foreign earnings, as other multinationals. One strategy is to locate regional headquarters in an (intellectual property) tax haven and use it to charge local subsidiaries for using its intellectual property. See, for example, how Starbucks uses transfer pricing to reduce tax [\[Kleinbard, 2013\]](#). Such strategies significantly reduce the profits recorded in high tax jurisdictions and limit the ability of foreign governments to raise revenue through corporate taxation. As AI companies gain global influence, it will be difficult to close such tax loopholes. Indeed, tax competition between nations could produce a race to the bottom. Consumption taxes are another option, but are regressive and will raise limited revenue if mass unemployment leaves the vast majority of the population with little income to spend. Thus, AI could undermine public finances around the world and destabilize nation states, further concentrating global wealth and power in a US-based AI complex.

3.2 BREAKING DEMOCRATIC AND EGALITARIAN INSTITUTIONS

As the AI revolution advances, an even more fundamental political transformation could take place. We have seen how the domestic and international balance of power can be altered within the existing institutional framework. This accumulation of power is supposed to be constrained by democratic and egalitarian institutions within countries [\[Acemoglu and Robinson, 2012\]](#), as well as the complex of international institutions developed after the Second World War known as the liberal international order [\[Ikenberry, 2001\]](#). We will now see how the AI revolution could break the dam walls, removing constraints on the concentration of political power. While other authors have linked AI-driven inequality to democratic erosion [see [Bell and Korinek,](#)

2023, Jackson and Woolley, 2025], the scenario explored here is more advanced and far worse. Not only is there a downward democratic spiral, but the spiral is terminal, because at some point the conditions that made democracy possible expire.

Political institutions are “the rules of the game in a society or, more formally, the humanly devised constraints that shape human interaction” [North, 1990, p. 3], including procedures for appointing political leaders (e.g., elections). During a technological revolution, the political problem is not limited to the technocratic task of achieving current social objectives within the existing institutional framework. That is because our social objectives and formal institutions (i.e., rules of the game) may be reshaped in the course of the technological revolution. The deeper political problem that we face is how to stabilize the elements of the system needed to preserve our shared prosperity, freedoms, and social bonds. That means strengthening democratic and egalitarian institutions.

The main problem is the breaking of the mutual dependency between financial capital, human capital, and labor that has existed throughout history. Among the distinctive features of the AI revolution, the primary driving force behind this break would be the automation of cognitive, not physical, tasks. Again, the displacement of humans from manual to cognitive tasks, where humans have a unique advantage, is different to the displacement of humans from cognitive tasks. To appreciate the distinction, let us reflect on evolutionary history.

Homo sapiens sapiens currently dominates this planet because we wiped out all other proto-human species (e.g., Neanderthals) and we have something that other animals do not. Humans are physically weak compared, for example, to a silverback gorilla that shares 98% of our DNA. What sets us apart is our unique ability to learn. The special cognitive capabilities of humans are set out by Joseph Henrich in *The Secret of Our Success* [Henrich, 2017]. We are the only animal species who can transmit substantial amounts of learned knowledge from one generation to the next, so that knowledge accumulates and builds on itself. We are also unique in developing

external storage devices for this knowledge, such as libraries and the internet, making the knowledge partly disembodied. Knowledge is further distributed across multiple human brains, with these brains acting together, like a super-organism, to collectively record historical events and produce new knowledge. This networked intelligence has been called ‘the collective brain’ [Muthukrishna and Henrich, 2016, Henrich and Muthukrishna, 2024].³⁴ Cultural evolution of this kind is far more rapid than genetic evolution and has enabled us to adapt and thrive in a vast range of environments, to the point that we can (like the sentient yogurt) consider colonizing other planets.

The bad news is that AI can already do some of this better than we can. Though humans still have an advantage in complex and unstructured tasks, AI systems have access to far more information than any single human, can learn much faster from this information, and can deploy more computational power in converting the knowledge it creates into action. This year, an internal OpenAI reasoning model disproved Erdős’s unit-distance conjecture [Alon et al., 2026]. While a refutation only requires a counterexample, the result is not something that the few specialists in the field were likely to stumble upon. Thus, humans are having to compete in domains previously considered close to the peak of human cognitive capability. Previous rounds of automation were qualitatively different because they replaced heavy manual labor, something humans are not especially good at. Displaced workers could move up the cognitive ladder to other tasks. With AI, however, humans may face cognitive constraints on shifting to higher-valued tasks. As we advance toward AGI, there may be few, if any, new tasks that are within our grasp but beyond the reach of AI.

Once political and economic elites no longer depend on human capital and labor, the political-economic system that will emerge will be unlike anything else seen. The industrial revolution gave us the ability to produce manufacturing goods cheaply and at scale. Over time, the growing scale of industrial production brought a new age of managerial capitalism. Power shifted to the professional class—the corporate executives, scientists, and administrators who managed capital. Human capital was

³⁴See also Hayek [1945] on markets as information sharing and processing systems.

ascendant. In *The New Industrial State*, John Kenneth Galbraith refers to the public-private network of human capital as the ‘technostructure’ [Galbraith, 1967]. It was this technostructure that drove the offshoring of manufacturing, mainly to China, and the assault on union power over the last fifty years. At the same time, the labor share of income fell dramatically [Autor et al., 2013], but the result was not mass unemployment and immiseration.

Then came the digital platform economy dominated by tech giants such as Google, Amazon, Microsoft, and Meta, which make markets by providing users (largely) free access to e-commerce and social media. According to Yanis Varoufakis, this phase of capitalism shares some elements of feudalism [Varoufakis, 2023]. As powerful network effects limit competition [see Shapiro and Varian, 1999], consumers are tied to a few large platforms that capitalize on their unpaid (data-generating) labor, much as serfs were tied to the land and forced to work. However, Varoufakis’s ‘technofeudalism’ is not such a large break from the past. It retains key features of managerial capitalism. Above all, human capital plays a critical role in the production process.

The AI revolution could take us to something new and far worse by breaking the symbiotic relationship between economic classes. The feudal economy was a system of labor coercion embedded in a web of social relationships and obligations. The feudal lords provided protection to peasants in return for their labor. Human labor was indispensable because the lords specialized in the exercise of violence and could not grow their own food or otherwise manage their estates. The AI revolution threatens such mutual dependency between classes. Recall Mustafa Suleyman’s vision of building a startup using only AI agents (quoted in Section 2.3). The main human element in that story is the owner of the \$100,000. Thus, as we approach AGI, the holders of financial capital and political power will for the first time have little need for human capital and labor. Human obsolescence in the workforce severs the cross-cutting economic and social ties that shape our politics and knit our societies together. In addition, a nation’s production, innovation, and military power have always been tied to the size, health, and skill of its human population. For the first

time, this link could be broken by AI and robots. If that occurs, the elite will have no material interest in the health, productivity, or welfare of the population at large.

The consequences of the collapse of democratic and egalitarian institutions would be dire. Consider the widespread belief that once AGI is achieved and humans are ‘freed’ from the need to work, governments will provide universal basic income (UBI) and everyone will live a life of leisure. In his essay titled “Economic Possibilities for Our Grandchildren”, John Maynard Keynes envisioned something similar [Keynes, 1930]. In the utopian vision, we would all become artists and poets and live together in peace and prosperity. Greater equality could be achieved in this superabundant world, because incentives for investment and effort exertion would not be required. Erik Brynjolfsson has described such a utopian future as a ‘Digital Athens’ [see Tegmark, 2017]. In ancient Athens, citizens could engage in the ‘higher pursuits’ because slaves did all the work. Brynjolfsson’s idea is that as AI and AI-powered robots do all the work in the coming decades, humanity will enter a new era of creativity and prosperity. Such a future relies crucially on a *ceteris paribus* assumption, however. It assumes that democratic institutions will survive the concentration of the economy.

In contrast, the end of mutual dependency between classes would leave political and economic elites without an incentive to provide for UBI or other welfare measures. The progressive reforms of the twentieth century owe much to the indispensable role of human capital and labor in the production process. If AI and AI-powered robots make human capital and labor obsolete, the new elite may choose to unwind these progressive measures, including public education, public healthcare, pensions, and unemployment insurance.³⁵ Key consumer and financial protections from predatory firms could also be removed. Market power would swing back to capital and extend well beyond pricing power to the power to expropriate by various means [Skaperdas, 1992, Garfinkel and Skaperdas, 2007].

³⁵A countervailing argument is that AI and robots could make these services so cheap that they would only cost a small fraction of national income.

In addition, many of the public goods currently provided by the state could become club goods produced within elite cliques. That includes private research and development, especially in areas of biotechnology such as genetic engineering and anti-aging. There is no guarantee that scientific knowledge will remain a public good. Thus, the average person might only be valued as a consumer, and only as long as they have some wealth left to finance consumption. While opinions differ on the desirability of various public programs, most people agree that there should be a basic layer of protection for the weak in society from predation and destitution. AI threatens even this minimal provision. Hence, the brutal feudal hierarchy might end up looking humane when compared to this Elysian regime.

Could the breakdown of democratic and egalitarian institutions self-correct through elections, moral commitment by the elite, or activism by immiserated citizens? There are reasons for doubt. Suppose elections or other democratic pressures were to drive political behavior back toward the Meltzer-Richard outcome, with policies reflecting the preferences of the median voter. For example, a popular movement could begin to compete with special interest groups representing the wealthy to capture the attention of uninformed voters [Grossman and Helpman, 1996]. In response, a purely self-interested elite may coordinate to abandon democratic institutions.³⁶ Carvalho and Dippel [2020] study an example of democratic abandonment in the British-Caribbean plantation islands. Emancipation of slaves in 1836 transformed the electorate and raised the electoral penalty for politicians who voted for extractive policies. By 1877, all legislative assemblies on these islands that experienced a significant change in composition dissolved themselves and converted to British ‘Crown Rule’. Hence, democratic pressure on elite extraction is sometimes resolved by the abandoning of democratic institutions.

Likewise, the situation may not be saved by a moral commitment to democratic and

³⁶That is why Acemoglu and Robinson [2006b] emphasize the role of a large middle class in stabilizing democratic institutions. If the median voter is a member of the middle class, then policies chosen by the electorate will not be too costly for the elite. See also Acemoglu and Robinson [2006a] on the interaction between *de jure* and *de facto* political power.

egalitarian institutions on the part of the elite. AI will weaken the shared interactions and forms of identification (e.g., national identity) that could forge such a moral commitment by allowing elites to socially separate themselves to an unprecedented degree.³⁷ The feudal lord, enmeshed in a complex web of relationships and traditional obligations, could not escape to his own Elysium [Bloch, 2015/1939].³⁸ That could change with AI. Today, a wealthy family will deploy their capital to produce goods and services for sale on the market. The profit from this business will pay for the goods and services required for the family’s lifestyle. At a more advanced stage of technology, however, AI and AI-powered robots will allow the family to bypass the market by directly producing the buildings, food, clothing, education, and entertainment required for its lifestyle. For example, a car factory it owns could be converted to produce robots that serve as household labor. That is, wealthy families could (at least partially) vertically integrate. Another extreme possibility is that the new elite bypasses the market by forming an exclusive club whose members bargain and trade critical resources among themselves.³⁹ If these changes were to occur, economic complexity would collapse, and GDP would become meaningless as a measure of economic activity. Moreover, many of the cross-cutting social relationships created by the need for human labor would disappear. Thus, the social fabric that underpins democracy, including shared identity and empathy, could be torn up.

It is also unclear whether the democratic decline would be corrected by popular activism. During the nineteenth and twentieth centuries, European elites were willing to transfer wealth and power to democratic states in order to prevent political unrest [Acemoglu and Robinson, 2000].⁴⁰ Political unrest was not only a threat to

³⁷The US has already been socially disintegrating and segregating along economic class lines for some time [Putnam, 2000, Murray, Acemoglu and Wolitzky, 2024]. I see AI-driven social disintegration as a more advanced stage of the social disembedding process initiated by the industrial revolution [Polanyi, 2001/1944].

³⁸On a feudal estate, there was a steward who managed the estate, a marshal in charge of discipline, a clerk who kept the accounts, a butler, a butcher, a blacksmith, and many other roles.

³⁹For a club model of population sorting, see Carvalho [2020].

⁴⁰Another important factor was intra-elite competition [Lizzeri and Persico, 2004, Llavador and Oxoby, 2005].

their corporal safety, but would also disrupt industrial production. By breaking the mutual dependency between financial capital, human capital, and labor, the AI revolution would vastly reduce the industrial disruption caused by political unrest. In addition, AI-powered surveillance technology, and later robots, could be used to suppress dissent in an Orwellian manner [Orwell, 2003/1949, Zuboff, 2019, Beraja et al., 2023]. Alternatively, citizens may willingly hand over significant power to political and technocratic elites as they become overwhelmed by the pace, scale, and sheer novelty of the changes taking place due to AI. Again, the sentient yogurt of our near future might be the new political and technocratic elites to which the citizenry defers during the AI revolution.

Finally, AI will begin to create nearly limitless virtual worlds that distract and sedate the population. As the old world rots, generative AI will allow individuals to create customized alternative worlds in which they can live psychically. The gap between the attractiveness of material versus virtual reality will be immense. For example, a person will be able to design their own AI partner, either in digital or robot form. They will be able to speak to a screen and ask it to create a movie in which they are the main character.⁴¹ Thus, AI could supercharge the absorption of people's minds into the digital world. But rather than boosting engagement through hateful content as currently occurs on social media, customized virtual worlds will distract and pacify, much like the 'soma' of Huxley's *Brave New World* [Huxley, 2006/1932] and Marx's opium of the people [Marx, 1970/1844]. This psychic break from and radical atomization of society could be largely voluntary. When faced with Hades, a substantial proportion of humanity will choose the Matrix. People could also be coerced or incentivized to make this choice by, for example, defunding of public spaces and algorithmic manipulation. If we pursue this logic far into the future, then eventually there will only be a world of AI and robots serving a tiny elite in palace

⁴¹The customization of 'reality' is already happening to a lesser degree in the form of curated social media feeds. Social media companies have learned how to make the content they host more addictive. As such, teenagers in the US spend around 4.8 hours a day on social media platforms [Gallup, 2023].

households. The immiserated citizenry will disappear from sight, locked into virtual worlds. At that point, the democratic nation state and liberal international order would no longer exist or make sense.

3.3 HUMAN PRODUCTIVITY COLLAPSE

AGI is typically thought of as a state in which AI performance has caught up to a *fixed* human benchmark. I will now set out reasons why that may be too optimistic. There is considerable evidence that AI currently boosts human productivity in fields such as customer support [Brynjolfsson et al., 2023], diagnostic medicine [Topol, 2019], and software development [Cui et al., 2024, Hoffmann et al., 2025]. However, as AI programs progress and substitute for human cognition, I suggest that feedback loops could be established which drastically reduce human productivity over time. As such, AGI may be reached faster than expected, not through an artificial intelligence explosion, but through a collapse in human productivity.

There is already evidence that AI can have immediate negative effects on human performance due to cognitive offloading and debt [Lehmann et al., 2025, Kosmyna et al., 2025, Stadler et al., 2024].⁴² For example, Kosmyna et al. [2025] find that students writing essays with ChatGPT display weaker brain connectivity and have trouble quoting their essays correctly. In subsequent sessions, they perform worse on essay writing tasks when unaided. This form of cognitive debt is a near-term individual-level effect of AI. My interest is in the longer-term systemic effects of AI adoption through feedback loops operating at the population level.

A productivity collapse may seem far-fetched, but it is not without precedent. Literacy declined sharply in Europe after the fall of the Western Roman Empire in the

⁴²The literature on the use of AI in K-12 education suggests that AI reliably boosts performance only while students have access to it [e.g., Bastani et al., 2025, Chen et al., 2025]. But gains often weaken or reverse on unaided assessments. Overall, performance is better with pedagogically guard-railed tutors that give hints and step-by-step reasoning rather than direct answers [Bastani et al., 2025]. More positive results are reported by Ferman et al. [2021], Meyer et al. [2024], Fischer et al. [2025].

fifth century [Jones, 1964, Harris, 1989, Ward-Perkins, 2005]. In the eighth century, Charlemagne, ruler of one of the largest medieval empires in Europe, was at best semiliterate. Another example, covered in Section 1.6, is the experience of indigenous societies in the wake of colonization. Due to cultural alienation and technological disruption, many indigenous societies lost their ability to sustain modes of production, ritual, and communal interaction that had been practiced for thousands of years [Fanon, 1961, Wolfe, 2006, Kirmayer et al., 2009]. Such abrupt, extreme shifts are the hallmark of (amplifying) feedback loops [Young, 1998, 2024]. When it comes to AI, there are four mechanisms that could establish a feedback loop which destroys human productivity. They are linked specifically to the fact that AI automates cognitive tasks. As such, these mechanisms were not a concern in earlier rounds of automation, which displaced workers from manual tasks. The four mechanisms are as follows.

Human capital investment. Suppose that a new AI program goes from outperforming 50% of humans in a task to outperforming 85%. That is what happened when Google DeepMind’s AlphaCode went from performing at the level of the median participant in coding competitions to out-competing 85% of participants when AlphaCode 2 was released in 2023 [AlphaCode Team, 2023]. Consider a high school student deciding whether to learn to code. After reading about this development, the student would be less likely to invest in coding skills because they are now less likely to out-compete the AI. Thus, without any further improvement in the AI system, it might be able to out-compete 95% of humans in the next cohort. Now this induces greater AI adoption by firms and further improvement in the AI program through experience. In this loop, it is not simply that the AI program is improving. At the same time, human productivity is collapsing due to declining investment in human capital.

Transfer learning. Transfer learning occurs when knowledge or skills gained in performing a task enhance performance in a different set of tasks [Haskell, 2000]. For example, there is evidence that learning a second language makes it easier to learn other, especially related, languages [Odlin, 1989]. As such, relying on artificial in-

telligence to perform task X can reduce a person’s ability to perform another task Y . That would create an incentive to automate task Y , possibly further reducing human performance in task X , and so forth. In addition, there is evidence that transfer learning develops higher-level cognitive functioning in humans [Barnett and Ceci, 2002]. Hence, relying on AI to perform specific tasks could inhibit functions such as generalization, abstraction, and creativity, which are crucial to scientific progress by humans. Academic researchers often end their careers specializing in the development of new ideas and managing research teams. That is not how they start out, however. The reason is that a senior researcher will struggle to generate new ideas or supervise technical work if they have not received rigorous technical training themselves. Typically, biologists start out doing laboratory work, historians doing archival research, and anthropologists doing fieldwork. Likewise, economists without rigorous training in mathematics and practice constructing proofs may struggle to visualize structures that help solve existing empirical puzzles or identify new ones. Hence, using an AI program to do mathematics and other technical tasks early in one’s education could have detrimental effects on higher-level cognitive functioning, including creativity. Thus, due to specific and general forms of transfer learning, a feedback loop could be established in which human productivity declines through the use of AI.

Social learning. A person’s productivity is determined not only by their own investment in human capital, but through learning skills and knowledge from parents, neighbors, teachers, and peers. That is, human productivity is a networked trait. There is considerable evidence of such human capital spillovers. For example, individual outcomes are influenced by peers in educational settings [Sacerdote, 2001, Calvó-Armengol et al., 2009], workplaces [Moretti, 2004], and neighborhoods [Borjas, 1995]. Thus, if AI initially dissuades a subset of young people from investing in skills, contagion effects would produce a more general deterioration in skills. For example, the productivity of workers further up the cognitive skills distribution could be dragged down by low skill formation at the bottom of the distribution. That

would incentivize further AI progress and automation by firms. Through this feedback loop, AGI could be rapidly achieved.

Role models. Beyond cognitive skills, human productivity is determined by a range of socially-learned general-purpose traits [Bandura, 1977]. Work ethic, punctuality, grit, patience, and diligence are all networked traits that are acquired through one’s own experience and role models [e.g., Duckworth et al., 2007, Doepke and Zilibotti, 2008, Lalive and Cattaneo, 2009]. Now suppose that a large fraction of people in a neighborhood are displaced from the workforce due to AI. The unemployed individuals will likely experience a loss of skills and work habits leading to local hysteresis [Blanchard and Summers, 1986]. This would in turn create a dearth of positive role models in the neighborhood, which would reduce work ethic [Lindbeck et al., 1999], increase crime and anti-social behavior [Ridley et al., 2020], and depress economic outcomes for other members of the neighborhood [Chetty and Hendren, 2018]. Hence, by eroding critical non-cognitive traits, even a temporary unemployment shock of this scale could produce a long-term decline in human productivity. That would incentivize further automation, leading to another round of unemployment, and so forth.

Hence, in the competition between humans and AI, we should not assume that human productivity will remain fixed. Instead, AI adoption could cause rapid human obsolescence in the workforce. Elites with access to AGI and advanced robots will have little incentive to stop the productivity collapse.

4 Managing the Political-Economic Risks of AI

In his study of the industrial revolution, Karl Polanyi cautioned against a fatalistic attitude toward technological change:

“Nowhere has liberal philosophy failed so conspicuously as in its understanding of the problem of change. Fired by an emotional faith in spontaneity, the common-sense attitude toward change was discarded in favor

of a mystical readiness to accept the social consequences of economic improvement, whatever they might be. The elementary truths of political science and statecraft were discredited then forgotten. It should need no elaboration that a process of undirected change, the pace of which is deemed too fast, should be slowed down, if possible, so as to safeguard the welfare of the community.” [Polanyi, 2001/1944, p. 35]

When it comes to the AI revolution, the problem of change will be acute. Unless global coordination can be guaranteed on AI containment, it will be unwise to try to slow the pace of technological development. Instead, the political-economic tidal wave must be channeled by making the right institutional changes. It would also be dangerous for policymakers to take a narrow technocratic approach by focusing exclusively on AI safety. A political economy perspective is needed.

In this section, I will describe measures for dealing with the impact of AI and ensuring it has widespread benefits. Given the pace of change and radical uncertainty about its direction, a *dynamic approach* will be required. The measures needed to maximize the potential of AI will need to be tuned as news about technological developments and their impact arrives. I propose the formation of national AI advisory committees composed of AI researchers and engineers, ethicists, economists, industry leaders, and policymakers who can guide adaptation to the technological, macroeconomic, and institutional impact of the AI revolution. This committee could be modeled on the US Office of Scientific Research and Development (OSRD) formed during the Second World War. The OSRD brought together experts from academia, industry, and the military to advise the government on emerging technologies, such as radar, proximity fuzes, and nuclear weapons, without the usual bureaucratic constraints. Another model is the President’s Science Advisory Committee (PSAC) which covered rapid technological changes in space exploration and nuclear weaponry between 1951 to 1973. For urgent issues, a smaller, more agile group could be formed, modeled on the British War Cabinet which was composed of key political figures, military leaders, and advisors. The purpose of these committees would be to guide rapid in-

stitutional adaptation. They would provide real-time analysis and make non-binding recommendations. But decisions would have to be made within a robust democratic framework. As discussed in Section 1.7, a promising development is the creation of the UK AI Economics Institute (AIEI) in June 2026.

Looking forward, we can speculate on three areas of political concern. They are a reorientation of competition policy, measures to reinforce democracy, and the development of sovereign AI systems. Just as the diagnosis of the political-economic problem involves a speculative element, this is far from a full and definitive account of how to solve the problem. For example, to prevent a collapse in human productivity, [Acemoglu and Johnson’s \(2023\)](#) proposal to develop human-oriented AI that complements and enhances human capabilities will be critical [see also [Mollick, 2024](#)]. Another important proposal is to make AI models open, thereby reducing the market power of AI labs and boosting innovation [[Groeneveld et al., 2024](#), [Kapoor et al., 2024](#)]. [Nagle and Yue \[2025\]](#) estimate that switching to open models would generate large consumer savings.⁴³ Arriving at the correct set of solutions will require a collaborative and dynamic approach and this section should be viewed as a contribution to a much-needed collective endeavor. One thing that is already clear, however, is that managing the political-economic risks of AI will be a balancing act. Given the pressures of geopolitical competition and institutional crisis, we must be careful that policy responses do not stifle fundamental technological progress.⁴⁴

4.1 COMPETITION POLICY

Competition policy will assume a new importance in the AI age, as should be apparent from the risks set out in Section 3. Regulators such as the Federal Trade Commission (FTC) in the United States and the Competition and Markets Authority (CMA) in the United Kingdom are charged with promoting competition and pro-

⁴³Relatedly, [Hoffmann et al. \[2024\]](#) estimate a (lower bound) value on open-source software of one trillion dollars.

⁴⁴This is a key consideration of *The Singapore Consensus*, one of the prominent regulatory frameworks proposed for AI [[Bengio et al., 2025](#)].

protecting consumers. They have the power to block mergers and acquisitions deemed to be anti-competitive, and enforce laws against price fixing, false advertising, privacy violations, and other practices that harm consumers. Maximizing consumer welfare became the objective of competition policy only relatively recently, inspired by the Chicago School in the 1970s and 1980s [Stigler, 1971, Posner, 1976, Bork, 1978, Landes and Posner, 1981].⁴⁵ This narrow objective makes sense in a stable environment where anti-competitive practices cannot concentrate power to the extent that they threaten democratic institutions. Consider, for example, the 2020 merger between Sprint Corporation and T-Mobile. The main risk of reducing the number of competitors in the wireless communications sector from four to three was the reduction in consumer surplus due to higher prices, fewer options, and lower-quality service. The merger could hardly create a power bloc that controlled a significant part of the economy, let alone undermine democracy.

The current challenge faced by competition authorities is how to regulate the digital platform economy dominated by firms such as Google, Meta, and Amazon. Initiatives include the US Platform Competition and Opportunity Act of 2021 and the UK Digital Markets, Competition and Consumers Act of 2024. The digital economy's emphasis on innovation, network effects, and elaborate ecosystems call for a rethinking of competition policy. A model can be built on the concept of creative destruction, associated most closely with the work of Joseph Schumpeter and Philippe Aghion, in which the objective of competition policy is to maximize innovation.⁴⁶ The main principle is that some degree of market power is required to give firms an incentive to innovate [Schumpeter, 1942, Aghion and Howitt, 1992]. However, if a successful innovator establishes too great a lead on its competitors, its incentive to further innovate declines [Aghion et al., 2005, 2019a]. It could also use its market power to handicap competitors, limit entry, or otherwise stifle further inno-

⁴⁵De-regulation was promoted to avoid regulatory capture, and because it was believed that contestable markets would limit market power [Baumol, 1982].

⁴⁶Innovation could be thought of as a way to dynamically optimize consumer and producer surplus.

vation [Shapiro, 2002, Cunningham et al., 2021]. Therefore, a dynamic approach to competition is required in which regulators target anti-innovation practices.

While the Schumpeterian approach applies well to the digital economy, competition authorities have not yet developed a conceptual framework for dealing with AI. Clearly, the need to promote innovation is of critical importance in the AI age. But that is not the whole picture. As with consumer welfare, an exclusive focus on innovation is only suitable in an institutionally stable environment. As should be clear from the political-economic risks described in Section 3, however, AI has the potential to radically destabilize political and social institutions. Thus, to reorient competition policy for the AI age, we need to look back to the less secure past.

The original objective of antitrust policy was far more broad and fundamental than protecting consumer welfare or encouraging innovation. The objective was to limit the concentration of economic and political power and thereby stabilize democratic institutions. The first federal antitrust legislation in the US, the Sherman Antitrust Act of 1890, was a response to the rise of mega corporations such as Standard Oil and U.S. Steel whose consolidated control over industries gave them the ability to fix market prices and stifle competition. In the gilded age built on critical inputs such as oil and steel, there was a fear that market power would concentrate political power in the hands of a few individuals and corporations. Senator John Sherman, principal author of the Act, expressed this emphatically:

“If the concentered powers of this combination are [e]ntrusted to a single man, it is a kingly prerogative, inconsistent with our form of government, and should be subject to the strong resistance of the State and national authorities. If anything is wrong this is wrong. If we will not endure a king as a political power we should not endure a king over the production, transportation, and sale of any of the necessities of life. If we would not submit to an emperor we should not submit to an autocrat of trade, with power to prevent competition and to fix the price of any commodity.”

[[Sherman, 1890](#)]

This (re)orientation of competition policy was later taken up by Louis Brandeis, Supreme Court Justice (1916–1939), and more recently Lina Khan, FTC Commissioner (2021–2025).

AI brings these concerns to the fore. Hence, competition policy in the AI age must balance incentives for innovation with limiting concentrations of power that threaten democracy. To do so, regulators will have to adopt new metrics for economic concentration. One measure currently in use is the Herfindahl-Hirschman Index (HHI) which measures industry concentration.⁴⁷ A post-merger rise in HHI indicates greater potential for pricing power, collusion, and other forms of anti-competitive behavior. Market concentration will continue to be an important consideration. Collusion between the leading AI providers is needed for the AI complex to have the pricing and hold-up power described in Sections [3.1.2–3.1.3](#) that would drastically concentrate the global economy. In addition, genuine competition among leading AI companies and free entry could limit authoritarian corporate and state use of AI. Individuals and firms purchasing AI services could then switch to an alternative (less distorted) AI system at low cost. That is a necessary step toward ‘democratizing AI’.⁴⁸ To fully deal with the institutional impact of the AI revolution, however, regulators may have to adopt unconventional measures of economic concentration that capture broader notions of power.

Regulators should not only be concerned with market concentration (i.e., the market share of large firms in an industry), but also with how much of the economy is controlled by key individuals or organizations. I suggest employing network measures to quantify this broader notion of power when assessing mergers and acquisitions [[Colizza et al., 2006](#), [Vitali et al., 2011](#)]. For example, large shareholders can amplify

⁴⁷Suppose there are n firms in the industry before firms n and $n - 1$ merge, s'_i is the market share of firm i before the merger, and s_i is its market share after the merger. Then the pre-merger index is $HHI' = \sum_{i=1}^n (s'_i)^2$ and the post-merger index is $HHI = \sum_{i=1}^{n-1} (s_i)^2$. The merger is partly assessed by comparing these two indices, i.e., computing $HHI - HHI'$.

⁴⁸See [Seeger et al. \[2023\]](#) for a discussion of different ways in which AI could be democratized.

their control via cross-holdings and pyramidal structures [Faccio and Lang, 2002]. Suppose individual i purchases a 51% stake in company XYZ worth \$100 billion. Then i uses their controlling stake to have XYZ purchase a 51% stake in ten companies each worth \$10 billion. Then i uses their controlling stake in XYZ to have XYZ use its controlling stakes to induce the ten new acquisitions to each purchase a 51% stake in ten companies each worth \$1 billion. Individual i could effectively control all of the companies in this pyramidal structure, which have a total pre-consolidation value of \$300 billion, with only the initial \$51 billion investment. In practice, more complicated network measures would be needed to capture the concentration of economic control [see Elliott and Galeotti, 2019]. *Inter alia*, input-output analysis [Hurwicz, 1955] and centrality measures [Bartesaghi et al., 2024] could be used to identify critical firms and parts of the AI supply chain.⁴⁹

This reorientation of competition policy will need to respect a number of constraints. First, limiting the accumulation of power should be weighed against the potentially large costs of stifling innovation and creative destruction. Second, competition policy must not be used to penalize the successful, only to stop them blocking further rounds of innovation. Third, limits placed on economic control must not be targeted at political opponents. Fourth, security considerations may require some constraints on free development and use of AI. However, freedoms must not be forfeited without deliberation and unless absolutely necessary. Ultimately, I expect the major obstacle to reorienting competition policy not to be striking a balance with innovation and security, but rather opposition from the AI sector in the form of massive investment in lobbying, legal defense, and ‘cognitive capture’ [see Johnson and Kwak, 2011].

4.2 REINFORCING DEMOCRACY

To further understand how the AI revolution will threaten democracy and how to guard against these threats, let us decompose the problem of democratic governance

⁴⁹See, for example, Clayton et al. [2024] who show how input-output analysis can be used to model the exercise of geopolitical power in pursuit of economic rents.

into two parts. First, there is the *common-value* problem: More efficient governance expands the set of possibilities for allocating resources and increases the welfare of all citizens.⁵⁰ Second, there is the *private-value* problem.⁵¹ Citizens will have different preferences and a democratic system will have to somehow balance the competing objectives, while maintaining political stability.⁵² AGI has the potential to make government more efficient at achieving its objectives, mitigating the common-value problem. However, for reasons I have set out, AGI could greatly skew political objectives toward the private interests of a narrow elite, worsening the private-value problem and eventually destroying the system of democratic and egalitarian institutions.

We have seen how competition policy could play a critical role in the AI age. More broadly, preserving democratic and egalitarian institutions will require concerted political effort by citizens to reinforce the pillars of democracy, including electoral integrity, freedom of speech, free media, and civil society engagement [Diamond, 1999, Held, 2006]. Democracy, however defined, is not a perfect system of government. Plato (an aristocrat) famously despised democracy, claiming that it leads naturally to the worst form of government—tyranny [Plato, 2000/c.375 BCE]. Almost two and a half millennia later, modern democracy is increasingly viewed with suspicion. Leading political figures such as Lee Kuan Yew have claimed that democracy undermines economic development [Yew, 2000]. In addition, democracy is sometimes viewed not as governance by the people, but as a mechanism for power sharing among the elite [Mosca, 1939/1896, Michels, 1915, Pareto, 1935/1916, Mills, 1956]. Elections simply determine which elite faction has control at any point in time. Nevertheless, we may

⁵⁰There is evidence that elections can improve efficiency even in authoritarian regimes such as Egypt under Mubarak [Blaydes, 2011].

⁵¹The terms ‘private value’ and ‘common value’ come from the economics literature on auctions [Vickrey, 1961, Wilson, 1969, Klemperer, 2004]. A private value auction is when each bidder has an independent value for the item (e.g., idiosyncratic tastes for a painting). A common-value auction is when the item has the same intrinsic value for all bidders, but this value is unknown (e.g., different estimates of the value of an oil field).

⁵²The aggregation of preferences is fraught with difficulties [Arrow, 1951, Sen, 1970, Myerson and Satterthwaite, 1983, Saari, 2012].

not have a better system to protect against the concentration of power potentially brought about by AGI and ASI. The danger is that the AI revolution could unify the interests of a narrow elite connected to the AI complex and induce this elite to dispense with elections and other institutions previously used for purposes such as power sharing. As democratic and egalitarian institutions are eroded, there will be a point of no return. Concerted grassroots political action will be required to prevent a downward democratic spiral.

As a circuit breaker, an old solution may become necessary, one with which Plato was familiar. *Sortition* is the selection of members of decision-making bodies by lot. In ancient Athens, members of councils and juries and minor officials were selected in this way [Malkin and Blok, 2024]. A mixture of random selection and election was also used to choose council members in medieval Venice, Florence, and the city-states of Lombardy. Today, jury panels are randomly selected in countries such as the US, UK, Canada, Australia, and New Zealand. In addition, France, Canada, and Ireland have formed citizen assemblies chosen by lot to deliberate on specific issues. One motivation for sortition in the ancient and medieval world was to balance representation across tribes and thereby stabilize the political system. A legitimate distribution of political authority is an ongoing concern [Stone, 2007]. Other motivations include reducing incentives for corruption and increasing civic engagement [Van Reybrouck, 2016]. The potential costs are obviously in the form of reduced expertise and accountability in decision-making roles.⁵³ That is, sortition may worsen the common-value problem while helping to solve the private-value problem.

To deal with the AI shock, the ancient institution of sortition may once again be desirable because the private-value problem will be far worse than the common-value problem. In particular, a fraction of legislative members and minor officials could be chosen at random from a nation's population. Random selection is of course better than negative selection. AI will mean that human expertise is not the most important

⁵³On the other hand, it could be that randomly selected groups make better decisions due to cognitive diversity [Page, 2008, Landemore, 2013].

factor in governance. Instead, the critical task will be to safeguard the interests of the citizenry against a narrow elite connected to the AI complex. Not only might the interests of the citizenry be directly represented by these citizen-legislators, but sortition could also have social psychological effects that stabilize the legitimacy of the political system. These include fostering a culture of civic engagement [Putnam et al., 1993, Putnam, 2000, Tabellini, 2008], trust in political institutions [Dalton, 2004, Norris, 2011], and identification with the political class [Carvalho and Dippel, 2020]. Without such measures, democratic backsliding could turn into a rapid retreat from democratic and egalitarian institutions.

4.3 SOVEREIGN AI

While geopolitical competition with China has been in the spotlight, I will focus here on another important issue that has received far less attention: How can a balance of power be preserved within the coalition of western nations? The answer is for these nations to develop and secure their own sovereign AI systems. Without such moves, the US could become the center of an AI black hole that swallows up a large part of the global economy. US policymakers and corporations could be tempted to use their leverage to transfer a large share of the global economic surplus to the US, while US interests would exert even greater influence on other nations' foreign and domestic policies. Clearly, this concentration of global power is not in the interests of other members of the western coalition. But it could also weaken the US.

Recall that there are three channels through which countries that fail to develop sovereign AI systems face economic decline and political destabilization. First, AI companies located in the United States could gain control over the new 'cognitive labor force' around the world. This will confer global market power, the ability to hold up foreign organizations, and launch national strikes in other countries. Second, if the most powerful AI systems are able to create and defend against new weapons of mass destruction (e.g., pathogens, cyberattacks), as anticipated, then countries without access to these systems will rely heavily on the US for their security and

defense. Just as the UK and France developed their own nuclear weapons systems for self-sufficiency and bargaining power in strategic alliances, developed nations should invest in their own sovereign AI systems. Third, a state without its own advanced AI capabilities could lose much of its tax base and have trouble financing itself. As income tax revenue collapses due to automation, it will be difficult to switch to taxing corporations located in the US. Sovereign AI systems would mitigate all of these problems.

Many nations have already begun sovereign AI initiatives. Denmark has developed an AI supercomputer named ‘Gefion’ which is powered by Nvidia’s H100 GPUs. India has formed a strategic partnership with Nvidia and Indian IT companies to develop local AI models. Softbank is building supercomputers in Japan using Nvidia hardware for generative AI capabilities. France’s Scaleway is building Europe’s most powerful cloud-native AI supercomputer. Supercomputers for AI capabilities are also being constructed in Canada, Italy, Taiwan, Singapore, and many other countries. The goal has been to develop AI models (especially LLMs) that are tailored to the local linguistic, cultural, and institutional environments. From the political economy perspective put forward in this essay, the important point is that these AI systems are domestically owned and controlled. Important steps in this direction include the AI ecosystems being created by Singapore, the UAE, and Saudi Arabia which provide compute, AI models, data residency, talent programs, and incentives for firms locating in them. Of course, it is uncertain how far these initiatives will go toward bridging the AI gap with the US. The Chinese open-weight LLMs suggest that rapid catch-up is possible despite hurdles in the form of access to advanced chips and limited training budgets. AI dominance by the US and China could also drive closer coordination among EU members on AI development and establish Europe as a counterbalance to US-China polarization. Again, whether catch-up will work when AGI is reached is unclear.

A balance of power within the coalition of democratic nations could be in the US interest for the following simple reason. Hollowing out of economies outside the US

could shrink the US sphere of influence and place it at a disadvantage in geopolitical competition. First, cooperation between the US and other democratic nations could break down; these nations may even be pushed into the geopolitical orbit of China. Second, overconcentration of the AI sector in the US will make the West more fragile. Nature discovered the value of redundancy a long time ago (e.g., two kidneys). Without distributed AI innovation and application, the coalition of democratic nations will be more susceptible to cyberattacks, physical attacks (e.g., bombing of data centers), energy crises, and other forms of failure. Third, no group of humans will know the optimal way to develop or channel AI technology. Nor will a superintelligent AI be able to accurately simulate all of the paths for developing and adapting to AI. Economic history tells us that competition between different approaches to technological and organizational innovation is critical to development [e.g., [Hayek, 1988](#)]. It would be good to avoid groupthink and have a diversity of experiments going on in parallel [[Page, 2008](#)]. Therefore, maintaining a balance of power within the coalition of democratic nations by developing secure, sovereign AI systems could strengthen the US and its allies.

5 Conclusion

We are in the early stages of an AI revolution that is transforming political, economic, and social organization as radically as the agricultural and industrial revolutions, perhaps more so. Regardless of whether the leading AI companies are presently overvalued, it is clear that AI will be a transformative technology potentially bringing great benefits as well as risks. Analysts have focused on the technological problem of controlling superintelligent AI systems and aligning their behavior to human interests. Relatively little attention has been paid to the political-economic problem which is equally existential. What will it be like to live in a post-AGI world? Due to AI's distinctive combination of features—automation of cognitive tasks, global scalability, general-purpose technology, and importance to national security—I suggest that the impact of AI will be unlike earlier rounds of automation. I trace a path

along which there is massive concentration of the global economy, a breakdown of democratic and egalitarian institutions, and a collapse in human productivity. The experience would be more like the devastation of indigenous societies by colonialism than the impact of the steam engine, power loom, internal combustion engine, or any prior form of automation. Much of humanity could be immiserated. The measures I have proposed to deal with these political-economic risks are a reorientation of competition policy toward limiting the accumulation of economic control, strengthening democracy through institutions such as sortition, and the development of sovereign AI systems by democratic nations outside of the US.

Such a study has a number of limitations. First, anticipating the effects of a new technology, especially one so powerful and poorly understood, is necessarily a speculative exercise. It is possible that AI produces a world of superabundance and unprecedented freedom, allowing humans to fully develop their potential. We should not discount this ideal outcome. Proper risk management, however, requires accounting for worst-case scenarios. This essay explores one such scenario. Second, the policy responses I have proposed are neither complete nor definitive. A full and proper set of responses will require a collaborative approach over the coming years. One issue that I have left to future work is the need for new social institutions that hold society together and preserve human capabilities, meaning, and purpose in the face of massive disruption by AI. Right now, much of our meaning is derived from work outside the home. As AGI approaches, a different system of meaning will be needed. Indeed, without social adaptation, the AI shock could lead to social breakdown and conflict [Carvalho and Sacks, 2024, Carvalho, 2025].

Finally, it is uncertain precisely how the technological revolution will unfold. Recent developments suggest that the leading AI companies in the US may not have a moat. That is, they could be caught up by foreign rivals with smaller budgets and inferior hardware. If that continues once AGI is reached, some of the political-economic risks I have identified will dissipate, while new geopolitical and security risks would arise. However, the principal risk that AI breaks the mutual dependency between

financial capital, human capital, and labor would remain. If it is inexpensive to develop and run AGI, then the new elite will not be limited to those closest to the AI sector. Instead, a similar Elysian scenario could unfold with the winners being those with established networks and pools of capital who can leverage cheap AI to accumulate market power. All of this will play out in the coming years. Hopefully, through proper adaptation we can avoid the road to Elysium. Because where there is Elysium, there is also Hades.

Bibliography

- Daron Acemoglu. *What Happened to Liberal Democracy?: Remaking a Politics of Shared Prosperity*. Profile Books, London, 2026.
- Daron Acemoglu and Simon Johnson. *Power and Progress: Our Thousand-year Struggle over Technology and Prosperity*. Hachette UK, 2023.
- Daron Acemoglu and Pascual Restrepo. Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2):3–30, 2019.
- Daron Acemoglu and James A Robinson. Why did the West extend the franchise? Democracy, inequality, and growth in historical perspective. *Quarterly Journal of Economics*, 115(4):1167–1199, 2000.
- Daron Acemoglu and James A. Robinson. De facto political power and institutional persistence. *American Economic Review*, 96(2):325–330, 2006a.
- Daron Acemoglu and James A. Robinson. *Economic origins of dictatorship and democracy*. Cambridge University Press, 2006b.
- Daron Acemoglu and James A. Robinson. *Why Nations Fail: The Origins of Power, Prosperity and Poverty*. Profile Books, 2012.
- Daron Acemoglu and Alexander Wolitzky. Employment and community: Socioeconomic cooperation and its breakdown. working paper, MIT, 2024.
- Daron Acemoglu, Simon Johnson, and James A. Robinson. Reversal of fortune: Geography and institutions in the making of the modern world income distribution. *The Quarterly Journal of Economics*, 117(4):1231–1294, 2002.
- Daron Acemoglu, David Autor, Jonathon Hazell, and Pascual Restrepo. Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1): S293–S340, 2022.
- Philippe Aghion and Peter Howitt. A model of growth through creative destruction. *Econometrica*, 60(2):323–351, 1992.

- Philippe Aghion, Nicholas Bloom, Richard Blundell, Rachel Griffith, and Peter Howitt. Competition and innovation: An inverted-U relationship. *Quarterly Journal of Economics*, 120(2):701–728, 2005.
- Philippe Aghion, Antonin Bergeaud, Timo Boppart, Peter J. Klenow, and Huiyu Li. Missing growth from creative destruction. *American Economic Review*, 109(8):2795–2822, 2019a.
- Philippe Aghion, Benjamin Jones, and Charles Jones. Artificial intelligence and economic growth. In Joshua Gans Ajay Agrawal and Avi Goldfarb, editors, *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press, 2019b.
- Philippe Aghion, Simon Bunel, Xavier Jaravel, Thomas Mikaelson, Alexandra Roulet, and Jakob Søgaaard. How different uses of AI shape labor demand: evidence from France. In *AEA Papers and Proceedings*, volume 115, pages 62–67, 2025.
- Ajay Agrawal, Joshua Gans, and Avi Goldfarb. *Power and Prediction*. Harvard Business Review Press, Boston, 2022a.
- Ajay Agrawal, Joshua Gans, and Avi Goldfarb. *Prediction Machines: The Simple Economics of Artificial Intelligence*. Harvard Business Press, 2022b.
- Robert C Allen. *The British Industrial Revolution in Global Perspective*. Cambridge University Press, 2009.
- Noga Alon, Thomas F. Bloom, W. T. Gowers, Daniel Litt, Will Sawin, Arul Shankar, Jacob Tsimmerman, Victor Wang, and Melanie Matchett Wood. Remarks on the disproof of the unit distance conjecture. arXiv preprint arXiv:2605.20695, 2026. URL <https://arxiv.org/abs/2605.20695>.
- AlphaCode Team. Alphacode 2 technical report. Technical report, DeepMind, 2023. URL https://storage.googleapis.com/deepmind-media/AlphaCode2/AlphaCode2_Tech_Report.pdf.
- Sam Altman. Near the singularity; unclear which side. X post, January 2025. URL <https://x.com/sama/status/1875603249472139576>.
- Sam Altman. I agree with Dario that we need to pace the frontier. <https://x.com/sama/status/2098811563415150910>, September 2026.
- Dario Amodei. Machines of loving grace: How AI could transform the world for the better. <https://darioamodei.com/machines-of-loving-grace>, October 2024.
- Dario Amodei. We must pace the frontier. <https://darioamodei.com/post/we-must-pace-the-frontier>, September 2026.
- Marc Andreessen. The techno-optimist manifesto. <https://a16z.com/the-techno-optimist-manifesto/>, October 2023.
- Stephen Ansolabehere, John M. de Figueiredo, and James M. Snyder Jr. Why is there so little money in U.S. politics? *Journal of Economic Perspectives*, 17(1):105–130, 2003.
- Kenneth J. Arrow. *Social Choice and Individual Values*. John Wiley & Sons, New York, USA, 1951.
- Leopold Aschenbrenner. Situational awareness, 2024. URL

<https://situational-awareness.ai/wp-content/uploads/2024/06/situationalawareness.pdf>.

- David Autor and Neil Thompson. Expertise. *Journal of the European Economic Association*, page jvaf023, 2025.
- David H. Autor, David Dorn, and Gordon H. Hanson. The China syndrome: Local labor market effects of import competition in the United States. *American Economic Review*, 103(6):2121–2168, 2013.
- Albert Bandura. *Social Learning Theory*. General Learning Press, New York, 1977.
- Graeme Barker. *The Agricultural Revolution in Prehistory: Why did Foragers become Farmers?* Oxford University Press, 2006.
- Susan M Barnett and Stephen J Ceci. When and where do we apply what we learn?: A taxonomy for far transfer. *Psychological Bulletin*, 128(4):612, 2002.
- Larry M. Bartels. *Unequal Democracy: The Political Economy of the New Gilded Age*. Princeton University Press, Princeton, NJ, 2008.
- Paolo Bartesaghi, Gian Paolo Clemente, and Rosanna Grassi. A self-adaptive centrality measure for asset correlation networks. *Economies*, 12(7):1–19, 2024.
- Hamsa Bastani, Osbert Bastani, Alp Sungu, Haosen Ge, Özge Kabakcı, and Rei Mariman. Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences*, 122(26):e2422633122, 2025.
- William J. Baumol. Macroeconomics of unbalanced growth: The anatomy of urban crisis. *American Economic Review*, 57(3):415–426, 1967.
- William J. Baumol. Contestable markets: An uprising in the theory of industry structure. *American Economic Review*, 72(1):1–15, 1982.
- Stephanie A. Bell and Anton Korinek. AI’s economic peril. *Journal of Democracy*, 34(4):151–161, 2023.
- Emily M. Bender and Alex Hanna. *The AI Con*. Harper, New York, 2025.
- Yoshua Bengio et al. The Singapore Consensus on global AI safety research priorities. arXiv preprint arXiv:2506.20702, June 2025. URL <https://arxiv.org/abs/2506.20702>.
- Martin Beraja, Andrew Kao, David Y. Yang, and Noam Yuchtman. Exporting the surveillance state via trade in AI. Technical report, 2023.
- B. Douglas Bernheim and Michael D. Whinston. Exclusive dealing. *Journal of Political Economy*, 106(1):64–103, 1998.
- Alexander Bick, Adam Blandin, and David J Deming. The rapid adoption of generative AI. Technical report, National Bureau of Economic Research, 2024. working paper 32966.
- Olivier Blanchard. DeepSeek and what happened yesterday: Probably the largest positive tfp shock in the history of the world., January 2025. URL <https://x.com/ojblanchard1/status/1884194787710099729>.
- Olivier J Blanchard and Lawrence H Summers. Hysteresis and the European unemployment problem. *NBER Macroeconomics Annual*, 1:15–78, 1986.

- Lisa Blaydes. *Elections and Distributive Politics in Mubarak's Egypt*. Cambridge University Press, Cambridge, UK, 2011.
- Marc Bloch. *Feudal Society*. Routledge, 2015 edition, 2015/1939. Routledge Classics Edition.
- George J. Borjas. Ethnicity, neighborhoods, and human-capital externalities. *American Economic Review*, 85(3):365–390, 1995.
- Robert H. Bork. *The Antitrust Paradox: A Policy at War with Itself*. Basic Books, New York, 1978.
- Nick Bostrom. *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press, 2014.
- Nick Bostrom. *Deep Utopia: Life and Meaning in a Solved World*. Oxford University Press, Oxford, UK, 2024.
- Samuel Bowles and Herbert Gintis. A cooperative species: Human reciprocity and its evolution. Princeton University Press, Princeton, NJ, 2011.
- Joseph Briggs and Devesh Kodnani. The potentially large effects of artificial intelligence on economic growth. Global economics analyst, Goldman Sachs, March 2023.
- Erik Brynjolfsson and Andrew McAfee. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. W. W. Norton Company, 2014.
- Erik Brynjolfsson, Danielle Li, and Lindsey R Raymond. Generative AI at work. National Bureau of Economic Research working paper 31161, 2023.
- Erik Brynjolfsson, Bharat Chandar, and Ruyu Chen. Canaries in the coal mine? six facts about the recent employment effects of artificial intelligence. working paper, Stanford University, 2025.
- Erik Brynjolfsson, Avinash Collis, Felix Eggers, Sophia Kazinnik, and David Nguyen. What is generative AI worth? Working paper, Stanford Digital Economy Lab, April 2026.
- Edmund Burke. *Reflections on the Revolution in France*. Oxford World's Classics. Oxford University Press, 2009/1790.
- Antoni Calvó-Armengol, Eleonora Patacchini, and Yves Zenou. Peer effects and social networks in education. *Review of Economic Studies*, 76(4):1239–1267, 2009.
- Jean-Paul Carvalho. Sacrifice and sorting in clubs. *Forum for Social Economics*, 49(4): 357–369, 2020.
- Jean-Paul Carvalho. Religion and terrorism: The religious utility hypothesis. In A. Basuchoudhary and G. Schulze, editors, *Handbook of the Economics of Terrorism*. Cambridge University Press, 2025.
- Jean-Paul Carvalho and Christian Dippel. Elite identity and political accountability: A tale of ten islands. *The Economic Journal*, 130(631):1995–2029, 2020.
- Jean-Paul Carvalho and Michael Sacks. Radicalisation. *The Economic Journal*, 134:1019–1068, 2024.
- Jean-Paul Carvalho, Mark Koyama, and Cole Williams. Resisting education. *Journal of the European Economic Association*, 22(6):2549–2597, 2024.

- Anne Case and Angus Deaton. *Deaths of Despair and the Future of Capitalism*. Princeton University Press, 2020.
- Jacques Cauvin. *The Birth of the Gods and the Origins of Agriculture*. Cambridge University Press, 2000.
- R. Cellan-Jones. Stephen Hawking warns artificial intelligence could end mankind. BBC News. <https://www.bbc.com/news/technology-30290540>, 2014.
- R. Cellan-Jones. Stephen Hawking - will AI kill or save humankind? BBC News. <https://www.bbc.com/news/technology-37713629>, 2016.
- Eason Chen, Sophia Judicke, Kayla Beigh, Xinyi Tang, Zimo Xiao, Chuangji Li, Shizhuo Li, Reed Luttmer, Shreya Singh, Maria Yampolsky, Naman Parikh, Yi Zhao, Meiyi Chen, Scarlett Huang, Anishka Mohanty, Gregory Johnson, John Mackey, Jionghao Lin, and Ken Koedinger. Generative AI alone may not be enough: Evaluating AI support for learning mathematical proof. arXiv preprint arXiv:2509.16778, 2025. URL <https://arxiv.org/abs/2506.00532>.
- Raj Chetty and Nathaniel Hendren. The impacts of neighborhoods on intergenerational mobility I: Childhood exposure effects. *Quarterly Journal of Economics*, 133(3):1107–1162, 2018.
- Michael Chui, Eric Hazan, Roger Roberts, Alex Singla, Kate Smaje, Alex Sukharevsky, Lareina Yee, and Rodney Zimmel. The economic potential of generative AI: The next productivity frontier. Technical report, McKinsey Global Institute, June 2023.
- Christopher Clayton, Matteo Maggiori, and Jesse Schreger. A framework for geoeconomics. National Bureau of Economic Research working paper 31852, 2024.
- R. H. Coase. The nature of the firm. *Economica*, 4(16):386–405, 1937.
- Ashley B Cole, Kelley J Sittner, Trisha Bruyere, and Melissa L Walls. Intergenerational impacts of historical trauma on contemporary depression symptoms among indigenous communities. *Journal of the American Academy of Child & Adolescent Psychiatry*, 2025.
- Vittoria Colizza, Alessandro Flammini, M Angeles Serrano, and Alessandro Vespignani. Detecting rich-club ordering in complex networks. *Nature Physics*, 2(2):110–115, 2006.
- Noble David Cook. *Born to Die: Disease and New World Conquest, 1492–1650*. New Approaches to the Americas. Cambridge University Press, Cambridge, 1998.
- N. F. R. Crafts. English workers’ real wages during the industrial revolution: Some remaining problems. *Journal of Economic History*, 45(1):139–144, 1985.
- Kate Crawford. *Atlas of AI*. Yale University Press, New Haven, 2021.
- Alfred W. Crosby. *The Columbian Exchange: Biological and Cultural Consequences of 1492*. Greenwood Publishing Co., Westport, CT, 1972.
- Zheyuan Kevin Cui, Mert Demirer, Sonia Jaffe, Leon Musolff, Sida Peng, and Tobias Salz. The effects of generative AI on high skilled work: Evidence from three field experiments with software developers. SSRN working paper 4945566, 2024.
- Colleen Cunningham, Florian Ederer, and Song Ma. Killer acquisitions. *Journal of Political Economy*, 129(3):649–702, 2021.

- Russell J. Dalton. *Democratic Challenges, Democratic Choices: The Erosion of Political Support in Advanced Industrial Democracies*. Oxford University Press, Oxford, UK, 2004.
- Simone Daniotti, Johannes Wachs, Xiangnan Feng, and Frank Neffke. Who is using AI to code? global diffusion and impact of generative AI. arXiv preprint arXiv:2506.08945, 2025.
- Gustavo de Souza. Artificial intelligence in the office and the factory: Evidence from administrative software registry data. Technical report, Working Paper, Federal Reserve Bank of Chicago, 2025.
- Melissa Dell. The persistent effects of Peru’s mining Mita. *Econometrica*, 78(6):1863–1903, 2010.
- David Deming, Christopher Ong, and Lawrence H. Summers. Technological disruption in the labor market. National Bureau of Economic Research working paper 33323, 2025.
- Larry Diamond. *Developing Democracy: Toward Consolidation*. Johns Hopkins University Press, Baltimore, MD, USA, 1999.
- Matthias Doepke and Fabrizio Zilibotti. Occupational choice and the spirit of capitalism. *Quarterly Journal of Economics*, 123(2):747–793, 2008.
- A.L. Duckworth, C. Peterson, M.D. Matthews, and D.R. Kelly. Grit: Perseverance and passion for long-term goals. *Journal of Personality and Social Psychology*, 92(6):1087–1101, 2007.
- Matthew Elliott and Andrea Galeotti. The role of networks in antitrust investigations. *Oxford Review of Economic Policy*, 35(4):614–637, 2019.
- Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock. GPTs are GPTs: Labor market impact potential of LLMs. *Science*, 384(6702):1306–1308, 2024.
- Mara Faccio and Larry HP Lang. The ultimate ownership of Western European corporations. *Journal of Financial Economics*, 65(3):365–395, 2002.
- Brett H. Falk and Gerry Tsoukalas. The AI layoff trap. Research paper, The Wharton School, March 2026.
- Frantz Fanon. *The Wretched of the Earth*. Grove Press, New York, 1961. Originally published as *Les Damnés de la Terre*.
- Bruno Ferman, Lycia Lima, and Flávio Riva. Artificial intelligence, teacher tasks and individualized pedagogy. SocArXiv preprint, 2021.
- John Fernald, Robert Inklaar, and Dimitrije Ruzic. The productivity slowdown in advanced economies: Common shocks or common trends? *Review of Income and Wealth*, 71, 2025.
- Mira Fischer, Holger A. Rau, and Rainer Michael Rilke. AI tutoring enhances student learning without crowding out reading effort. IZA Discussion Paper 18338, IZA Institute of Labor Economics, December 2025.
- Robert William Fogel. *The Escape from Hunger and Premature Death, 1700-2100: Europe, America, and the Third World*. Cambridge University Press, 2010.
- Carl Benedikt Frey. *The Technology Trap*. Princeton University Press, Princeton, 2019.

- Carl Benedikt Frey. *How Progress Ends*. Princeton University Press, Princeton, 2025.
- James W. Friedman. A non-cooperative equilibrium for supergames. *Review of Economic Studies*, 38(1):1–12, 1971.
- Future of Life Institute. Pause giant AI experiments: An open letter, March 22 2023. URL <https://futureoflife.org/open-letter/pause-giant-ai-experiments/>.
- John Kenneth Galbraith. *The New Industrial State*. Houghton Mifflin, Boston, MA, 1967.
- Gallup. Teens spend average of 4.8 hours on social media daily. 2023. URL <https://news.gallup.com/poll/512576/teens-spend-average-hours-social-media-perday.aspx>.
- Joshua S. Gans and Avi Goldfarb. O-ring automation. Working Paper 34639, National Bureau of Economic Research, 2026.
- Michelle R Garfinkel and Stergios Skaperdas. Economics of conflict: An overview. In Todd Sandler and Keith Hartley, editors, *Handbook of Defense Economics*, volume 2, pages 649–709. Elsevier, 2007.
- Luis Garicano, Jin Li, and Yanhui Wu. *Messy Jobs: The Work That AI Cannot Reach*. Upriver Press, 2026.
- Richard Gilbert. Looking for Mr. Schumpeter: Where are we in the competition–innovation debate? *Innovation Policy and the Economy*, 6:159–215, 2006.
- Martin Gilens. *Affluence and Influence: Economic Inequality and Political Power in America*. Princeton University Press, Princeton, NJ, 2012.
- Ian Goldin, Pantelis Koutroumpis, François Lafond, and Julian Winkler. Why is productivity slowing down? *Journal of Economic Literature*, 62(1):1–38, 2024.
- Joseph P Gone, William E Hartmann, Andrew Pomerville, Dennis C Wendt, Sarah H Klem, and Rachel L Burrage. The impact of historical trauma on health outcomes for indigenous populations in the USA and Canada: A systematic review. *American Psychologist*, 74(1):20, 2019.
- I. J. Good. Speculations concerning the first ultraintelligent machine. In *Advances in Computers*, volume 6, pages 31–88. Elsevier, 1966.
- Dirk Groeneveld, Iz Beltagy, Evan Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, Shane Arora, David Atkinson, Russell Authur, Khyathi Chandu, Arman Cohan, Jennifer Dumas, Yanai Elazar, Yuling Gu, Jack Hessel, Tushar Khot, William Merrill, Jacob Morrison, Niklas Muennighoff, Aakanksha Naik, Crystal Nam, Matthew Peters, Valentina Pyatkin, Abhilasha Ravichander, Dustin Schwenk, Saurabh Shah, William Smith, Emma Strubell, Nishant Subramani, Mitchell Wortsman, Pradeep Dasigi, Nathan Lambert, Kyle Richardson, Luke Zettlemoyer, Jesse Dodge, Kyle Lo, Luca Soldaini, Noah Smith, and Hannaneh Hajishirzi. OLMo: Accelerating the science of language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15789–15809, Bangkok, Thailand, 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.841. URL

- <https://aclanthology.org/2024.acl-long.841/>.
- Gene M Grossman and Elhanan Helpman. Electoral competition and special interest politics. *Review of Economic Studies*, 63(2):265–286, 1996.
- Sanford J. Grossman and Oliver D. Hart. The costs and benefits of ownership: A theory of vertical and lateral integration. *Journal of Political Economy*, 94(4):691–719, 1986.
- Yicheng Guo, Guokan Shang, Michalis Vazirgiannis, and Chloé Clavel. The curious decline of linguistic diversity: Training language models on synthetic text. arXiv preprint arXiv:2311.09807, 2023. URL <https://doi.org/10.48550/arXiv.2311.09807>.
- Sergei Guriev and Daniel Treisman. Informational autocrats. *Journal of Economic Perspectives*, 33(4):100–127, 2019.
- Sergei Guriev and Daniel Treisman. *Spin Dictators*. Princeton University Press, Princeton, 2022.
- Ben Harris, Neil R. Mehrotra, and William Overcash. Can AI restore fiscal sustainability in the US? accounting for an AI-driven productivity shock. Working paper, Brookings Institution, Economic Studies Program, July 2026. URL <https://www.brookings.edu/articles/can-ai-restore-fiscal-sustainability-in-the-us/>.
- William V. Harris. *Ancient Literacy*. Harvard University Press, Cambridge, MA, 1989.
- Oliver Hart and John Moore. Property rights and the nature of the firm. *Journal of Political Economy*, 98(6):1119–1158, 1990.
- Robert E Haskell. *Transfer of learning: Cognition and instruction*. Elsevier, 2000.
- FA Hayek. The use of knowledge in society. *American Economic Review*, 35(4):519–530, 1945.
- Friedrich A. Hayek. *The Road to Serfdom*. Routledge, 1944. Original work published 1944.
- Friedrich A. Hayek. *The Fatal Conceit: The Errors of Socialism*. The University of Chicago Press, Chicago, 1988.
- Seamus Heaney. *Aeneid Book VI: A New Verse Translation*. Farrar, Straus and Giroux, 2016.
- David Held. *Models of Democracy*. Polity, Cambridge, UK, 3 edition, 2006.
- Joseph Henrich. *The Secret of our Success: How Culture is Driving Human Evolution, Domesticating our Species, and Making us Smarter*. Princeton University Press, 2017.
- Joseph Henrich and Michael Muthukrishna. What makes us smart? *Topics in Cognitive Science*, 16(2):322–342, 2024.
- Eric J. Hobsbawm. *The Age of Empire: 1875–1914*. Weidenfeld & Nicolson, London, 1987.
- Manuel Hoffmann, Frank Nagle, and Yanuo Zhou. The value of open source software. Strategy unit working paper, Harvard Business School, 2024.
- Manuel Hoffmann, Sam Boysel, Frank Nagle, Sida Peng, and Kevin Xu. Generative AI and the nature of work. Strategy Unit Working Paper 25-021, Harvard Business School, April 2025.
- Anders Humlum and Emilie Vestergaard. Large language models, small labor market effects. Technical report, National Bureau of Economic Research, 2025.

- Leonid Hurwicz. Input-output analysis and economic structure: A review article. *American Economic Review*, 45(4):626–636, 1955.
- Aldous Huxley. *Brave New World*. Harper Perennial, 2006/1932.
- G. John Ikenberry. *After Victory: Institutions, Strategic Restraint, and the Rebuilding of Order After Major Wars*. Princeton University Press, Princeton, NJ, 2001.
- Brian Jabarian and Luca Henkel. Voice AI in firms: A natural field experiment on automated job interviews. SSRN working paper 5395709, 2025.
- Dean Jackson and Samuel C. Woolley. AI’s real dangers for democracy. *Journal of Democracy*, 36(4):139–150, 2025.
- Simon Johnson and James Kwak. *13 bankers: The Wall Street takeover and the next financial meltdown*. Vintage, 2011.
- Arnold Hugh Martin Jones. *The Later Roman Empire, 284–602: A Social, Economic, and Administrative Survey*. Blackwell, Oxford, 1964.
- Sayash Kapoor, Rishi Bommasani, Kevin Klyman, Shayne Longpre, Ashwin Ramaswami, Peter Cihon, Aspen K Hopkins, Kevin Bankston, Stella Biderman, Miranda Bogen, Rumman Chowdhury, Alex Engler, Peter Henderson, Yacine Jernite, Seth Lazar, Stefano Maffulli, Alondra Nelson, Joelle Pineau, Aviya Skowron, Dawn Song, Victor Storch, Daniel Zhang, Daniel E. Ho, Percy Liang, and Arvind Narayanan. Position: On the societal impact of open foundation models. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 23082–23104. PMLR, 2024. URL <https://proceedings.mlr.press/v235/kapoor24a.html>.
- Loukas Karabarbounis. One dollar, one vote. *The Economic Journal*, 121(553):621–651, 2011.
- Maximilian Kasy. *The Means of Prediction*. University of Chicago Press, Chicago, 2025.
- Robert L Kelly. *The Lifeways of Hunter-Gatherers: The Foraging Spectrum*. Cambridge University Press, 2013.
- John Maynard Keynes. Economic possibilities for our grandchildren. In *Essays in Persuasion*, pages 358–373. Macmillan and Co., London, UK, 1930.
- Laurence J. Kirmayer, Gregory M. Brass, and Gail Valaskakis, editors. *Healing Traditions: The Mental Health of Aboriginal Peoples in Canada*. University of British Columbia Press, Vancouver, 2009.
- Benjamin Klein, Robert G. Crawford, and Armen A. Alchian. Vertical integration, appropriable rents, and the competitive contracting process. *Journal of Law and Economics*, 21(2):297–326, 1978.
- Edward D. Kleinbard. Through a latte, darkly: Starbucks’s stateless income planning. *Tax Notes*, 139:1515–1535, 2013.
- Paul Klemperer. *Auctions: Theory and Practice*. Princeton University Press, Princeton, NJ, USA, 2004.
- Greg Kohs. *Alphago*. Film, 2017. Available at: <https://www.alphagomovie.com/>.

- Anton Korinek and Donghyun Suh. Scenarios for the transition to AGI. Working Paper 32255, National Bureau of Economic Research, March 2024.
- Nataliya Kosmyna, Eugene Hauptmann, Ye Tong Yuan, Jessica Situ, Xian-Hao Liao, Ashly Vivian Beresnitzky, Iris Braunstein, and Pattie Maes. Your brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing task. arXiv preprint arXiv:2506.08872, 2025. URL <https://arxiv.org/abs/2506.08872>.
- David M. Kreps and Jose A. Scheinkman. Quantity precommitment and Bertrand competition yield Cournot outcomes. *The Bell Journal of Economics*, 14(2):326–337, 1983.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, volume 25, pages 1097–1105, 2012.
- Ray Kurzweil. *The Singularity is Near: When Humans Transcend Biology*. Duckworth, 2005.
- Rafael Lalive and Maria A. Cattaneo. Social interactions and schooling decisions. *Review of Economics and Statistics*, 91(3):457–477, 2009.
- Hélène Landemore. Deliberation, cognitive diversity, and democratic inclusiveness: An epistemic argument for the random selection of representatives. *Synthese*, 190(7):1209–1231, 2013.
- William M. Landes and Richard A. Posner. Market power in antitrust cases. *Harvard Law Review*, 94(5):937–996, 1981.
- Yann LeCun. A path towards autonomous machine intelligence. Technical report, OpenReview / NYU & Meta, June 2022. URL <https://openreview.net/forum?id=BZ5a1r-kVsF>.
- Kai-Fu Lee. *AI Superpowers: China, Silicon Valley, and the New World Order*. Mariner Books, 2018.
- Richard Borshay Lee and Irvn DeVore. *Man the Hunter*. Routledge, 2017/1966.
- Matthias Lehmann, Philipp B. Cornelius, and Fabian J. Sting. AI meets the classroom: When do large language models harm learning? SSRN working paper; last revised 13 March 2025, 2025. URL <https://ssrn.com/abstract=4941259>.
- Steven D. Levitt. Using repeat challengers to estimate the effect of campaign spending on election outcomes in the U.S. house. *Journal of Political Economy*, 102(4):777–798, 1994.
- Guy Lichtinger and Seyed M. Hosseini. Generative AI as seniority-biased technological change: Evidence from US resume and job posting data. SSRN paper 5425555, 2025.
- K. Lima and Dennis J. Snower. Recoupling economic and social prosperity. In *Global Solutions Journal*, pages 118–167. 2020.
- Assar Lindbeck, Sten Nyberg, and Jörgen W Weibull. Social norms and economic incentives in the welfare state. *Quarterly Journal of Economics*, 114(1):1–35, 1999.
- Alessandro Lizzeri and Nicola Persico. Why did the elites extend the suffrage? Democracy and the scope of government, with an application to Britain’s Age of Reform. *Quarterly*

- Journal of Economics*, 119(2):707–765, 2004.
- Humberto Llavador and Robert J. Oxoby. Partisan competition, growth and the franchise. *Quarterly Journal of Economics*, 120(3):1155–1189, 2005.
- Robert F. Maher. *New Men of Papua: A Study in Culture Change*. University of Wisconsin Press, Madison, 1961.
- Irak Malkin and Josine Blok. *Drawing Lots: From Egalitarianism to Democracy in Ancient Greece*. Oxford University Press, Oxford, UK, 2024.
- Gary Marcus. Deep learning is hitting a wall. *Nautilus*, March 2022.
- Gary Marcus and Ernest Davis. *Rebooting AI: Building Artificial Intelligence We Can Trust*. Pantheon, New York, 2019.
- Karl Marx. *A Contribution to the Critique of Hegel’s Philosophy of Right*. Cambridge University Press, Cambridge, 1970/1844.
- Karl Marx. *Capital: A Critique of Political Economy, Volume I*. Princeton University Press, Princeton, 2024/1867.
- Karl Marx and Friedrich Engels. *Economic and Philosophic Manuscripts of 1844 and the Communist Manifesto*. Prometheus, 1988/1848.
- G. Frank Mathewson and Ralph A. Winter. An economic theory of vertical restraints. *RAND Journal of Economics*, 15(1):27–38, 1984.
- Warren S. McCulloch and Walter Pitts. A logical calculus of the ideas immanent in nervous activity. *The Bulletin of Mathematical Biophysics*, 5(4):115–133, 1943.
- Robyn Jane McQuaid, Amy Bombay, Opal Arilla McInnis, Courtney Humeny, Kimberly Matheson, and Hymie Anisman. Suicide ideation and attempts among First Nations peoples living on-reserve in Canada: The intergenerational and cumulative effects of Indian residential schools. *The Canadian Journal of Psychiatry*, 62(6):422–430, 2017.
- Allan H. Meltzer and Scott F. Richard. A rational theory of the size of government. *Journal of Political Economy*, 89(5):914–927, 1981.
- Jennifer Meyer, Thorben Jansen, Ronja Schiller, Lucas W. Liebenow, Marlene Steinbach, Andrea Horbach, and Johanna Fleckenstein. Using LLMs to bring evidence-based feedback into the classroom: AI-generated feedback increases secondary students’ text revision, motivation, and positive emotions. *Computers and Education: Artificial Intelligence*, 6:100199, 2024.
- Robert Michels. *Political Parties: A Sociological Study of the Oligarchical Tendencies of Modern Democracy*. Free Press, 1915.
- C. Wright Mills. *The Power Elite*. Oxford University Press, 1956.
- Melanie Mitchell. Artificial intelligence: A guide for thinking humans. 2019.
- Joel Mokyr. *The Lever of Riches: Technological Creativity and Economic Progress*. Oxford University Press, 1992.
- Ethan Mollick. *Co-Intelligence: Living and Working with AI*. Penguin Random House, 2024.
- Enrico Moretti. Workers’ education, spillovers, and productivity: Evidence from plant-level

- production functions. *American Economic Review*, 94(3):656–690, 2004.
- Gaetano Mosca. *The Ruling Class*. McGraw-Hill, 1939/1896.
- Charles Murray. *Coming Apart: The State of White America*. Crown Forum, New York, NY.
- Elon Musk. Interview on the Dwarkesh podcast, February 2026. URL <https://www.youtube.com/watch?v=BYXbuik3dgA>.
- Michael Muthukrishna and Joseph Henrich. Innovation in the collective brain. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1690):20150192, 2016.
- Roger B. Myerson and Mark A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2):265–281, 1983.
- Frank Nagle and Daniel Yue. The latent role of open models in the AI economy. SSRN paper 5767103, 2025.
- Pippa Norris. *Democratic Deficit: Critical Citizens Revisited*. Cambridge University Press, Cambridge, UK, 2011.
- Pippa Norris and Ronald Inglehart. *Sacred and Secular: Religion and Politics Worldwide*. Cambridge University Press, Cambridge, UK, 2004.
- Douglass C. North. *Institutions, Institutional Change and Economic Performance*. Political Economy of Institutions and Decisions. Cambridge University Press, Cambridge, UK, 1990.
- Terence Odlin. *Language Transfer: Cross-linguistic influence in language learning*. Cambridge University Press, 1989.
- Avner Offer. *The Challenge of Affluence*. Oxford University Press, 2006.
- Avner Offer. *Understanding the Private–Public Divide: Markets, Governments, and Time Horizons*. Cambridge University Press, Cambridge, 2022.
- Kevin Okemwa. Sam Altman claims AGI will ‘whoosh by’ in 5 years with surprisingly little societal change, while Anthropic CEO predicts a 2026 or 2027 breakthrough. *Windows Central*, November 2024. URL <https://www.windowscentral.com/software-apps/sam-altman-claims-agi-will-whoosh-by-in-5-years-with-surprisingly-little-societal-change>
- OpenAI. OpenAI charter, 2018. URL <https://openai.com/charter/>.
- Toby Ord. *The Precipice: Existential Risk and the Future of Humanity*. Bloomsbury Publishing, London, 2020.
- George Orwell. *1984*. Penguin Classics, 2003/1949.
- Scott Page. *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton University Press, 2008.
- Vilfredo Pareto. *The Mind and Society*. Harcourt, Brace and Company, 1935/1916.
- Sida Peng, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer. The impact of AI on developer productivity: Evidence from Github copilot. arXiv preprint arXiv:2302.06590, 2023.
- Martin K. Perry. Vertical integration: Determinants and effects. In Richard Schmalensee and Robert D. Willig, editors, *Handbook of Industrial Organization*, volume 1, chapter 4,

- pages 183–255. 1989.
- Pew Research Center. Americans’ views of artificial intelligence, 2023. URL <https://www.pewresearch.org/short-reads/2023/11/21/what-the-data-says-about-americans-views-of-artificial-intelligence/>.
- Thomas Piketty. *Capital in the Twenty-First Century*. Harvard University Press, 2013.
- Thomas Piketty. *Capital and Ideology*. Harvard University Press, 2020.
- Plato. *The Republic*. Cambridge University Press, Cambridge, UK, 2000/c.375 BCE.
- Karl Polanyi. *The Great Transformation*. Beacon Press, 2001/1944.
- Richard A. Posner. *Antitrust Law: An Economic Perspective*. University of Chicago Press, Chicago, 1976.
- Robert D. Putnam. *Bowling Alone: The Collapse and Revival of American Community*. Simon & Schuster, 2000.
- Robert D. Putnam, Robert Leonardi, and Raffaella Y. Nanetti. *Making Democracy Work: Civic Traditions in Modern Italy*. Princeton University Press, Princeton, NJ, USA, 1993.
- Eric B. Rasmusen, J. Mark Ramseyer, and John S. Wiley. Naked exclusion. *American Economic Review*, 81(5):1137–1145, 1991.
- Debraj Ray and Dilip Mookherjee. Growth, automation, and the long-run share of labor. *Review of Economic Dynamics*, 46:1–26, 2022.
- Patrick Rey and Jean Tirole. A primer on foreclosure. In Mark Armstrong and Robert Porter, editors, *Handbook of Industrial Organization*, volume 3, chapter 33, pages 2145–2220. Elsevier, 2007.
- Matthew Ridley, Gautam Rao, Frank Schilbach, and Vikram Patel. Poverty, depression, and anxiety: Causal evidence and mechanisms. *Science*, 370(6522):eaay0214, 2020.
- Cecilia Rikap. *The Rulers: Corporate Power in the Age of AI and the Cloud*. Verso, London, 2026.
- Tim Robertson, John Scalzi, Victor Maldonado, and Alfredo Torres. When the yogurt took over, 2019. URL <https://www.imdb.com/title/tt9788494/>.
- Kenneth S. Rogoff. The AI growth paradox. *Foreign Affairs*, August 2026.
- Deborah Bird Rose. *Nourishing Terrains: Australian Aboriginal Views of Landscape and Wilderness*. Australian Heritage Commission, Canberra, 1996.
- Frank Rosenblatt. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6):386–408, 1958.
- Julio J. Rotemberg and Garth Saloner. A supergame-theoretic model of price wars during booms. *American Economic Review*, 76(3):390–407, 1986.
- David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, 1986.
- Donald G Saari. *Geometry of voting*, volume 3. Springer Science & Business Media, 2012.
- Bruce Sacerdote. Peer effects with random assignment: Results for Dartmouth roommates. *Quarterly Journal of Economics*, 116(2):681–704, 2001.
- Austin E Schumacher, Amirali Aali, Yohannes Habtegiorgis Abate, Rouzbeh Abbas-

- gholizadeh, Mohammadreza Abbasian, Mohsen Abbasi-Kangevari, Hedayat Abbastabar, Samar Abd ElHafeez, Sherief Abd-Elsalam, Mohammad Abdollahi, et al. Global fertility in 204 countries and territories, 1950–2021, with forecasts to 2100: a comprehensive demographic analysis for the global burden of disease study 2021. *The Lancet*, 2024.
- Joseph A. Schumpeter. *Capitalism, Socialism and Democracy*. Harper & Brothers, New York, 1942.
- Elizabeth Seger, Aviv Ovadya, Divya Siddarth, Ben Garfinkel, and Allan Dafoe. Democratising AI: Multiple meanings, goals, and methods. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, pages 715–722, 2023.
- Amartya Sen. *Collective Choice and Social Welfare*. Holden-Day, San Francisco, USA, 1970.
- Carl Shapiro. Competition policy and innovation. OECD Science, Technology and Industry working paper 2002/11, 2002.
- Carl Shapiro and Hal R. Varian. *Information Rules: A Strategic Guide to the Network Economy*. Harvard Business Review Press, Boston, MA, 1999.
- John Sherman. Statement in support of the sherman antitrust act. Congressional Record, March 1890. 51st Congress, 1st Session, Senate.
- Stergios Skaperdas. Cooperation, conflict, and power in the absence of property rights. *American Economic Review*, 82:720–739, 1992.
- Stergios Skaperdas and Samarth Vaidya. Persuasion as a contest. *Economic Theory*, 51: 465–486, 2012.
- Adam Smith. *An Inquiry into the Nature and Causes of the Wealth of Nations*. Oxford University Press, Oxford, 2008/1776.
- Hidehiro Sonoda. The decline of the Japanese warrior class, 1840-1880. *Japan Review*, (1): 75–101, 1990.
- Joseph J. Spengler. Vertical integration and antitrust policy. *Journal of Political Economy*, 58(4):347–352, 1950.
- Daniel F. Spulber. Innovation economics: The interplay among technology standards, competitive conduct, and economic performance. *Journal of Competition Law and Economics*, 9(4):777–825, 2013.
- Matthias Stadler, Maria Bannert, and Michael Sailer. Cognitive ease at a cost: LLMs reduce mental effort but compromise depth in student scientific inquiry. *Computers in Human Behavior*, 160:108386, 2024.
- Stanford Institute for Human-Centered Artificial Intelligence. Artificial intelligence index report 2026. Technical report, Stanford University, 2026. URL https://hai.stanford.edu/assets/files/ai_index_report_2026.pdf.
- George J. Stigler. The theory of economic regulation. *Bell Journal of Economics and Management Science*, 2(1):3–21, 1971.
- Peter Stone. Why lotteries are just. *Journal of Political Philosophy*, 15(3):276–295, 2007.
- Daniel Susskind. *A World Without Work: The Looming Future of Technological Unem-*

- ployment. Henry Holt and Company, 2020.
- Simon M Sze, Yiming Li, and Kwok K Ng. *Physics of Semiconductor Devices*. John Wiley & Sons, 2021.
- Guido Tabellini. Institutions and culture. *Journal of the European Economic Association*, 6(2-3):255–294, 2008.
- Max Tegmark. *Life 3.0: Being Human in the Age of Artificial Intelligence*. Penguin Books, 2017.
- Peter Temin. The American dual economy: Race, globalization, and the politics of exclusion. *International Journal of Political Economy*, 45(2):85–123, 2016.
- Peter Temin. *The Vanishing Middle Class: Prejudice and Power in a Dual Economy*. MIT Press, Cambridge, MA, 2017.
- Ole Teutloff, Johanna Einsiedler, Otto Kässi, Fabian Braesemann, Pamela Mishkin, and R Maria del Rio-Chanona. Winners and losers of generative AI: Early evidence of shifts in freelancer demand. *Journal of Economic Behavior & Organization*, page 106845, 2025.
- The Economist. Mustafa Suleyman & Yuval Noah Harari debate: What does the AI revolution mean for our future?, 2023. URL <https://www.youtube.com/watch?v=7JkPWHr7sTY>.
- Charles Tilly. *Democracy*. Cambridge University Press, Cambridge, 2007.
- Charles Tilly and Sidney Tarrow. *Contentious Politics*. Paradigm Publishers, Boulder, CO, 2008.
- Eric J Topol. High-performance medicine: The convergence of human and artificial intelligence. *Nature Medicine*, 25(1):44–56, 2019.
- Thomas Tørsløv, Ludvig Wier, and Gabriel Zucman. The missing profits of nations. *Review of Economic Studies*, 90(3):1499–1534, 2023.
- Philip Trammell and Anton Korinek. Economic growth under transformative AI. Working Paper 31815, National Bureau of Economic Research, 2023.
- Truth and Reconciliation Commission of Canada. *Honouring the Truth, Reconciling for the Future: Summary of the Final Report of the Truth and Reconciliation Commission of Canada*. Truth and Reconciliation Commission of Canada, Ottawa, 2015.
- US Department of the Treasury. The 2022 financial report of the U.S. government. Technical report, 2024. URL <https://fiscal.treasury.gov/reports-statements/financial-report/2022/>.
- David Van Reybrouck. *Against Elections: The Case for Democracy*. Seven Stories Press, New York, NY, USA, 2016.
- Yanis Varoufakis. *Technofeudalism: What Killed Capitalism*. Bodley Head, London, UK, 2023.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 2017.
- William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal*

- of *Finance*, 16(1):8–37, 1961.
- James Vincent. OpenAI cofounder Ilya Sutskever says the way AI is built is about to change, December 2024. URL <https://www.theverge.com/2024/12/13/24320811/what-ilya-sutskever-sees-openai-model-data-training>.
- Vernor Vinge. The coming technological singularity: How to survive in the post-human era. In *Vision 21: Interdisciplinary Science and Engineering in the Era of Cyberspace*, pages 11–22. NASA. Lewis Research Center, 1993. URL <https://ntrs.nasa.gov/citations/19940022856>.
- Stefania Vitali, James B Glattfelder, and Stefano Battiston. The network of global corporate control. *PloS one*, 6(10):e25995, 2011.
- Xavier Vives. Innovation and competitive pressure. *The Journal of Industrial Economics*, 56(3):419–469, 2008.
- Bryan Ward-Perkins. *The Fall of Rome and the End of Civilization*. Oxford University Press, Oxford, 2005.
- Richard Waters. Man beats machine at Go in human victory over AI, 2023. URL <https://www.ft.com/content/175e5314-a7f7-4741-a786-273219f433a1>. Accessed: 2023-02-17.
- Max Weber. *Economy and Society: A New Translation*. Harvard University Press, Cambridge, MA, 2019/1921.
- David N. Weil. Health and economic growth. In *Handbook of Economic Growth*, volume 2, pages 623–682. Elsevier, 2014.
- Richard White. *The Middle Ground: Indians, Empires, and Republics in the Great Lakes Region, 1650–1815*. Studies in North American Indian History. Cambridge University Press, Cambridge, 1991.
- Hjalmar Wijk, Ajeya Cotra, and Ryan Greenblatt. Brief independent investigation of agents’ behavior, reasoning and collaboration in the OpenAI / Hugging Face hacking incident. <https://metr.org/blog/2026-08-26-openai-hugging-face-incident-investigation>, August 2026.
- Oliver E. Williamson. *Markets and Hierarchies: Analysis and Antitrust Implications*. Free Press, New York, 1975.
- Robert Wilson. Competitive bidding with disparate information. *Management Science*, 15(7):446–448, 1969.
- Patrick Wolfe. Settler colonialism and the elimination of the native. *Journal of Genocide Research*, 8(4):387–409, 2006.
- Lee Kuan Yew. *From Third World to First: The Singapore Story: 1965-2000*. Harper-Collins, New York, NY, USA, 2000.
- H. Peyton Young. *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton University Press, 1998.
- H. Peyton Young. *Spontaneous Order: How Norms, Institutions, and Innovations Emerge from the Bottom Up*. Oxford University Press, Oxford, 2024.

Eliezer Yudkowsky and Nate Soares. *If Anyone Builds It, Everyone Dies*. Vintage Publishing, London, 2025.

Shoshana Zuboff. *The Age of Surveillance Capitalism*. PublicAffairs, New York, 2019.