

## INTRODUCTION

This Online Supplementary Appendix accompanies History's Most Revolutionary Innovation and provides expanded treatments of several arguments and empirical exercises that are necessarily compressed in the main text. The book's core thesis is institutional and political-economic: the Artificial Intelligence (AI) Revolution is best understood as a general-purpose-technology (GPT) shock whose diffusion, productivity effects, and distributional consequences depend on a prior scaffolding of standards, complementary investments, and governance.

The appendix deepens that account in four areas—history and mechanism, empirical corroboration, technical foundations, and geopolitical fragmentation—while keeping the book itself readable and tightly paced.

Section 1 places the AI Revolution in historical context by tracing the First, Second, and Third Industrial Revolutions through the lens of GPT diffusion.

Section 2 provides a deeper empirical and conceptual treatment of the innovation commons of the late Third Industrial Revolution, using patent citations and network analysis.

Section 3 offers a technical primer and historical narrative of modern AI, including deep learning, transformer architectures, and scaling dynamics.

Section 4 situates AI within a broader geopolitical and institutional environment, comparing U.S., Chinese, and European approaches to industrial policy and governance.

## ONLINE SUPPLEMENTARY APPENDIX TO HISTORY'S MOST REVOLUTIONARY INNOVATION, SECTION 1

To write college-level essays, populate spreadsheets organized around complex mathematical formulas, put together professional-caliber slideshows, and craft original, emotionally poignant poetry, generative AI exploits natural language processing (NLP), a field of AI dedicated to enabling computers to understand, interpret, and generate text and speech in a meaningful, context-aware way. The core technique that powers modern NLP is language modeling: a predictive task like autocomplete on your phone where an AI model learns from vast amounts of text to calculate what word will come next in any given sequence. By becoming exceptionally good at this single task, the model develops an emergent understanding of the complex rules of language—from grammar and syntax to meaning and context—allowing it to grasp the statistical relationship between words with great accuracy (see Jurafsky and Martin 2021).

But the deep neural networks and attention mechanisms that drive these linguistic feats are not limited to text (Vaswani et al. 2017). Variations of these same deep learning architectures power perception (computer vision), pattern recognition in complex data, and increasingly planning and control systems that help machines select sequences of actions to achieve goals (see Russell and Norvig 2009). In other words: AI is not one invention—it is a bundle of complementary capabilities that can be embedded across products, processes, and business models.

Previous industrial revolutions witnessed the confluence of several overlapping technologies too. The First Industrial Revolution relied on the synergy between the steam engine, advanced metallurgy, and mechanized textile production. The Second was built on the interplay of electricity, the internal combustion engine, and synthetic chemicals. The Third was driven by the convergence of semiconductors, computing architecture, and telecommunications networks.

But industrial revolutions are not just about discrete inventions—or even inventions that overlap and influence one another—but the commercialization of invention. It is the process of translating technical possibility into economic viability, turning a scientific breakthrough into a standardized, affordable product that generates mass demand. Often, this requires reimagining the product entirely, shifting the value proposition from the hardware itself to the services and experiences it enables. The Fourth Industrial Revolution should be no different.

For example, nowhere is the convergence of AI capabilities—vision, planning, and processing—more visible than in the driverless car. As automobiles continue to shed their traditional controls and embrace autonomy, they may no longer be sold primarily as standalone vehicles. Instead, they become mobility services—platforms that compete not only on the underlying autonomy stack but on the experience layer: safety, routing, entertainment, productivity, and integration with the rest of a customer's digital life (Berk 2025).

Similarly, the economic disruption and reorganization related to AI extends far beyond its ability to master the vagaries of language. As Chapter 11 will formalize, the labor-market impact of this transition involves far more than simple labor displacement. Staying with driverless cars for the moment: it may include new complementary work: fleet operations and dynamic-pricing analysts, predictive maintenance for sensor-heavy vehicles, and even hospitality-driven user-

experience design (ibid). Consistent with what happened during previous industrial revolutions, these roles—along with the regulatory and privacy officers needed to steward the data—should emerge once the technology is deployed widely.

### **What this Section of the Appendix Aims to Do**

This section of the appendix places the AI revolution in historical context by tracing the three earlier industrial revolutions and extracting the mechanisms they share.<sup>1</sup> The First Industrial Revolution revolved around steam power and textile automation. The second was about electricity, the internal combustion engine, and chemicals. The Third Industrial Revolution was centered on computers, semiconductors, and the internet. The boundary between the third and fourth revolutions is straddled by the mobile phone and the pre-generative AI digital, geolocational tailored economy that evolved around digital platforms; this is a transitional period I cover in several of the book's chapters.

This section's overarching message is that industrial revolutions are not inevitable. Even a superior technology can stall when switching costs are high, complementary components are missing, or firms and workers have already invested in the skills and routines of the old system. It is often rational for organizations to keep upgrading incumbent technologies—marginal improvements can offer higher short-term returns than rebuilding around a new, unreliable platform. Workers behave similarly: if most employers still run the old system, it can be rational to train for yesterday's tools rather than tomorrow's (Chari and Hopenhayn 1991).

History is full of such "stickiness." As Nuvolari (2009) documents, older technologies often fight back. Windmills and wooden sailing ships saw significant advancements in design that kept them competitive long after the arrival of steam. Waterpower remained a dominant energy source for decades in regions with abundant rivers. Most telling for our purposes, factories clung to the "group drive system"—complex networks of shafts and rubber belts powered by a central steam engine—long after more efficient electric motors became available. They did so because the new technology required a complete and costly architectural redesign of the shop floor (David 1990).

Older technologies also endure because new ones are often bedeviled by a chaotic array of incompatible designs. For example, the early days of railroads were marked by a lack of standardization in track gauges, leading to significant inefficiencies in transportation (Puffert 2002). Investors and speculators laid tracks with arbitrary gauges according to their individual preferences, necessitating the construction of railway carriages specifically designed for each unique track configuration (ibid). These carriages could only operate on their designated tracks and connect with their own carriage type, creating isolated segments of improved transportation (ibid). While these railways offered speeds and convenience far superior to horse-drawn

---

<sup>1</sup> While this section of the appendix mines the historical record for qualitative patterns, Chapter 11 of the book formalizes them into a quantitative model. It translates the historical lag described here—the "Standardization Phase"—into a variable I call "Friction." By modeling the gap between a technology's "Potential Yield" and its "Realized Yield," Chapter 11 demonstrates how market structure and policy choices determine whether we achieve a "High-Diffusion" equilibrium or get stuck in a "High-Friction" trap.

transport within their limited domains, this advantage evaporated when passengers needed to transfer between different rail lines (ibid). The fragmentation was so severe that, after accounting for connection delays between incompatible systems, journeys from New York to Washington remained nearly as time-consuming as before (ibid).

Therefore, industrial revolutions only happen when General Purpose Technologies (GPTs) such as steam engines and railroads become standardized and widely diffused. It is only then that these technologies will exhibit broad applicability, live up to their potential to undergo continuous improvements, and register transformative impacts (Bresnahan and Trajtenberg 1995). In turn, once a dominant design emerges and investment concentrates, performance often improves very quickly as scale economies, learning-by-doing, and cumulative R&D reinforce one another (ibid). Only then do follow-on innovations and commercial applications explode.

Because they are centered on GPTs, industrial revolutions are about the commercialization and maturation of a diverse group of technologies that contribute to the transformation of key industries and eventually the entire economy, leading to explosive productivity growth. They therefore require the creation and integration of new intermediate goods. Take railroads: an elaborate symphony of high-pressure steam engines, precision machine parts, and improved metallurgy for rail casting evolved simultaneously (Szostak 2014). To bring these components into profitable industrial relationships with each other, new supply chains had to blossom.

Moreover, while groundbreaking inventions like the steam engine, electricity, and semiconductors represent obvious quantum leaps in technical achievement, the adoption of GPTs is rarely linear: it's often characterized by slow initial uptake during a protracted gestation period. First, only the best-resourced organizations harness the new technology's potential. Second, it is only once new technologies become standardized that plug and play genericness, scale economies, and affordability help them spread across the economy—but only if firms, both big and small, also find ways to accommodate these technologies into their organizational DNA. This requires adjustments, both big and small, and even cultural changes. For example, for steel to become the bones and sinews of industrialized economies, mills had to first remake themselves around new metallurgical processes that required a cadre of trained chemists and mechanics to adapt new techniques to local conditions.

As this section will demonstrate, the combination of standardization and firm level adaptation is what allows GPTs to usher in transformative economic changes. Standardized steel components enabled advancements in railway and bridge engineering. Electric power advanced from short experimental demonstrations to powering entire factories and cities. Microprocessors went from handling simple logic circuits to running personal computers, data centers, and global networks. But these processes were themselves riddled with potential dead ends and nonstarters; it took the helping hand of the state to address several problems standing in the way of modularity.

Therefore, besides outlining the processes behind the disruptive innovation associated with industrial revolutions, this section also documents the pivotal role governments played during these transformations. It explores how patents fostered the invention and commercialization of key computer hardware technologies during the Third Industrial Revolution. Additionally, I will show how during each of the previous industrial revolutions, governments bankrolled and

coordinated basic research, promoted technological standardization, helped diffuse new GPTs, and furnished a host of public goods that built out supporting infrastructures and human capital.

The remainder of this section unpacks these dynamics. I begin by defining the engine of these transformations—General Purpose Technology (GPT) governed by an “S-curve” logic that dictates its diffusion. Next, I survey the historical record of the First, Second, and Third Industrial Revolutions to illustrate how this logic played out in practice. Then, I discuss the macroeconomic consequences of these shifts, specifically the “Productivity Paradox” and the recurring debate over automation and employment. I move on to detailing the policy toolkit governments have historically deployed to resolve the market failures that threaten these transitions, distinguishing between their roles as funders of basic science, enforcers of property rights, and coordinators of standardization. Finally, I explore how governments tend to grapple with distributional issues during industrial revolutions while still facilitating standardization.

## OVERVIEW OF THREE PRECEDING INDUSTRIAL REVOLUTIONS

As we walk through the first three industrial revolutions, keep three guiding questions in mind:

1. **Core Technology:** What was the core GPT, and what complementary technologies had to coevolve around it?
2. **Standardization Friction:** What specific elements—interfaces, components, protocols, or measurement systems—had to standardize before diffusion could accelerate?
3. **Complementary Investments:** What investments in infrastructure, skills, organization, and law were required to turn the invention into broad, enduring productivity gains?

As several chapters in the book show, technological co-evolution, interface standardization, and complementary investment are the same levers that will govern whether AI becomes a true GPT that transforms the economy and society, rather than a frontier demo that struggles to scale.

### First Industrial Revolution

The First Industrial Revolution was about solving a series of seemingly mundane problems: how to drain mines, move goods cheaply, and reliably convert energy into mechanical work inside factories. Steam engines proved to be the versatile X factor that helped solve all three. Early engines pumped water out of mines; improved engines efficiently turned that steam power into rotary motion that could drive factory machinery and, later, locomotives. Importantly, this created a virtuous cycle: steam engines drained the mines, making coal cheaper; cheaper coal, in turn, made steam power affordable for every other industry (Allen 2009). By freeing production from the geography of waterwheels and animal power, steam became the era’s first truly general-purpose source of industrial energy (Goldstone 2002).

In many ways, however, the demand for deeper and more efficient mining was downstream from blast furnaces, another innovation that took commercial flight during the industrial revolution.<sup>2</sup> Powered by a form of purified, processed coal called coke, they served as the technological

---

<sup>2</sup> This discussion builds on Menaldo (2016) and Menaldo (2021).

backbone that enabled the mass production of iron. These towering structures represented a quantum leap in metallurgical capability, allowing for unprecedented scale and speed in iron production that dramatically reduced costs while increasing output, making iron affordable and abundant enough to support rapid industrialization, as it was an essential material for machinery, railroads, bridges, and buildings.

However, producing iron called on first extracting both iron ore and coal through either open cast or underground mining methods. Once extracted, the ore underwent preprocessing—it was crushed, washed, and transported to blast furnaces for smelting. In the furnaces, the ore was combined with limestone and coke and subjected to powerful blasts of hot air that melted the ore. Their own exhaust gases were recirculated into the bottom of the furnace, heating the coke, limestone, and iron materials. When the iron reached its molten state, it was tapped from the furnace’s bottom and poured into molds called “pigs,” where it solidified into pig iron.

Early coal and iron mines faced a persistent challenge, however: they flooded constantly. As miners dug deeper shafts in search of richer veins of coal and ore, water seeped in, often rendering mines unusable. The invention of the steam engine was a direct response to this problem. Thomas Newcomen’s atmospheric engine, introduced in the early 18<sup>th</sup> century, became the first practical machine for pumping water out of mines. Although it burned considerable fuel and operated relatively slowly, it offered an unprecedented solution: it freed mines from the constraints of shallow drilling and expanded coal extraction. Over time, James Watt refined Newcomen’s design, dramatically improving the engine’s efficiency.<sup>3</sup>

The introduction of coal-powered steam engines enabled miners to bore significantly deeper shafts than previously possible. In turn, this unlocked access to vastly larger stocks of, not only coal, but also essential minerals like bauxite, copper, zinc, nickel, and iron—metals that served as key components in industrial plants, machinery, and finished products. As extraction capabilities improved, yields of coal and minerals increased dramatically across Europe, causing prices to plummet and further accelerating the pace of industrialization. Beyond its direct resource benefits, this mining revolution generated a deep stock of specialized knowledge around coal extraction and smelting processes.

These skills, particularly those centered on precision boring and calibration techniques, produced valuable technological spillovers that advanced machine design and numerous industrial applications, creating a virtuous cycle of innovation that propelled industrial development across multiple sectors (Jacob 2014).

The Boulton and Watt steam engine didn’t simply make deeper mining feasible; it also laid the groundwork for widespread applications of this power source to industrial processes. As the

---

<sup>3</sup> Watt significantly improved engine efficiency by ensuring the cylinder remained consistently hot while implementing a separate condenser chamber. After performing its work, the spent steam was drawn into this cooled vessel to condense, avoiding the need to repeatedly heat and cool the main cylinder with each stroke. As a result, his engines were far more efficient and consumed substantially less coal than their predecessors, making steam power more economical and practical for widespread industrial application (see Bottomley 2014).

textile industry boomed, mill owners increasingly turned to steam engines to replace waterpower. While waterwheels helped mechanize the textile industry and contributed significantly to the First Industrial Revolution by enabling the mass production of stronger cotton yarn that could withstand the tension required for warp threads—a critical advancement that previous spinning technologies had failed to achieve—they required fast-flowing rivers or streams.<sup>4</sup> Conversely, steam engines enabled a stable, centrally powered factory system that ran spinning mules, power looms, and other mechanized equipment, as they could be installed in urban centers or areas without large waterways, granting manufacturers more freedom in choosing mill locations (Goldstone 2002). Consequently, the textile sector flourished, driving down the cost of cloth and fueling a surge in consumer demand—as well as stimulating further demand for coal, iron, and mechanical parts that fed back into the emerging industrial economy (Hills 1970).

The next pressing challenge was finding a way to transfer large quantities of coal from remote mines to coastal piers for transport by ship or barge. Enter the railroad: initially a set of wooden or iron tracks over which animals or simple engines pulled carts of coal, it soon evolved into locomotive-powered lines (Wolmar 2009). Steam engines revolutionized freight transportation, drastically reducing the time and cost of hauling heavy resources. Just as with the steam’s other applications, railroads quickly found uses well beyond coal transportation: goods of all kinds could be shipped from city to city; letters and parcels moved more rapidly than ever before; and, crucially, people traveled in large numbers far beyond their hometowns. These rail links spurred the growth of new industrial centers, reshaped city life, and broadened personal horizons—entire families relocated for work opportunities that hadn’t existed a few decades earlier (see *ibid*).

Beyond mines, mills, and railroads, steam technology also began to solve other significant challenges. Steam-powered ships, for instance, overcame the limitations of wind and current, dramatically shrinking travel time across oceans and along rivers. Likewise, early steam-driven machines appeared in agriculture, such as portable traction engines used to power threshing machines. Although these developments were gradual and scattered at first, they steadily laid the foundation for more profound economic and social transformations (Mokyr 1990).

The industrial revolution eventually spread beyond England, into the European Continent, and the US, where several technologies were brought over from England or developed in parallel to English and sometimes French innovations (Haber et al. 2022: 37). Those included “[t]he invention of jigs and milling machines for cutting metal to precise tolerances, such that parts made from them would fit into any assembly of the same type” (*ibid*: 38).

---

<sup>4</sup> Richard Arkwright’s spinning machine—which became known as the water frame—represented a significant advancement in textile manufacturing. His 1769 patent utilized rollers to produce strong cotton yarn suitable for warp (lengthwise threads), surpassing James Hargreaves’s spinning jenny, which typically created weaker thread appropriate only for weft, i.e., filling yarn (see Bottomley 2014). While early prototypes were small, mature industrial versions could process 96 spindles simultaneously, dramatically increasing production capacity. These machines were widely installed in mills throughout Derbyshire and Lancashire, where the use of waterwheels to drive them gave rise to the name “water frame” (Hills 1970).

Textile manufacturing, which primarily blossomed in New England, created the “Silicon Valleys” of early American industrialization. Centered on inventions like self-acting mules (mechanized cotton spinning) and power looms (weaving), it had several spillovers. Consider Lowell, Massachusetts, which was named after Francis Cabot Lowell in honor of his numerous contributions to improving power looms and his pioneering efforts to developing the so-called integrated factory system.<sup>5</sup> During the early 1800s, it became a world-renowned hub for textile manufacturing. By the mid to late 19<sup>th</sup> Century, it had transitioned to producing machines such as steam engines and locomotives and machine tools such as lathes, planers, and milling machines (see Menaldo and Wittstock 2025).

As recounted in Chapter 1 of the book, the consequences of the first industrial revolution were profound.

Consider just one aspect here: as transportation and communication improved in the wake of the proliferation of railroads, steamships, and the telegraph, fragmented local markets were stitched together into integrated economies. While in the US a unified national market quickly emerged (Chandler 1977), eventually this happened globally too. As O’Rourke and Williamson (1999) document, prices between Europe and the Americas converged for both commodities—including wheat, cotton, and iron—and manufactured goods by around the middle of the 19<sup>th</sup> century. Financial markets became more interconnected as well, reducing price differences in securities and facilitating international investment flows. Finally, mass migration from European countries, China, and Japan to the Americas induced convergence in real wages (ibid).<sup>6</sup>

---

<sup>5</sup> Textile manufacturing also benefited from the invention of the cotton gin, which along with other innovations such as the steel plough, barbed wire, grain harvesters, refrigeration systems, and improved fertilizers, helped fuel an agricultural revolution during the 19<sup>th</sup> century. As railroads enabled dramatic improvements in transportation, they opened vast new territories across the United States for cultivation: farmers exploited these technologies to introduce new grains, fruits, vegetables, and livestock breeds better suited to different regions and market needs. The resulting mechanization significantly reduced labor requirements, allowing fewer farmers to manage larger operations more efficiently. These combined advancements generated massive increases in productivity per acre, fundamentally altering the agricultural landscape and creating food surpluses that supported growing urban populations while releasing labor for industrial production (Olmstead and Rhode 2008). The rise of Chicago, Illinois as a major commodity trading hub was enabled by these developments (Cronon 1991).

<sup>6</sup> However, the interconnectedness fostered by liberalized trade, global migration, and international finance faced a sharp reversal in the early 20<sup>th</sup> Century, culminating in a period of pronounced economic nationalism and escalating great-power conflict. This shift became pronounced during the interwar period (1919-1939), as the economic disruptions of World War I and the subsequent Great Depression fueled protectionist sentiments globally. America’s Smoot-Hawley Tariff Act (1930) sharply raised tariffs on thousands of imported goods in a bid to protect domestic industries. This triggered retaliatory tariffs from other nations, contributing to a steep decline in international trade. Concurrently, the classic gold standard, which had facilitated stable exchange rates and international financial flows during the 19<sup>th</sup> Century, fractured and ultimately collapsed in the 1930s, hampering international commerce and investment. Deglobalization

## The Second Industrial Revolution

The Second Industrial Revolution—running from approximately 1870 to the 1930s—harnessed electric power and the internal combustion engine’s motor force. Transportation evolved from rail to automobiles, fabrics changed from cotton to synthetic materials like rayon and nylon, and relatively cheap consumer goods such as canned food, radios, telephones, household appliances, pharmaceuticals, and photographic film were mass produced for the first time (McCloskey 2016; Gordon 2016). To manufacture these new products, equally seminal manufacturing techniques harnessed rubber, glass, petrochemicals, standardized machinery for molding and shaping, electric motors and equipment, turbines, aluminum, and prestressed concrete (Smil 2005).

The spread of relatively affordable energy underwrote the full electrification of factories with moving assembly lines and individuated (unit-driven) workspaces outfitted with machines plugged into electric sockets (David 1990). Between 1909 and 1929, the US experienced a sixfold increase in electricity use for manufacturing and residential power, along with a similar increase in horsepower per worker (Devine 1983). Cheap and reliable power enabled factories to abandon centralized line shafts with cumbersome pulleys and belts powered by centralized steam power. They transformed their layouts and workflows, making manufacturing more flexible and efficient while reducing capital outlays (David 1990).

Factories also implemented sophisticated time, space, and motion innovations that fundamentally restructured work environments, including the standardization of work hours through clocking in and out, the spatial reorganization of workspaces that maximized efficiency while minimizing machines’ physical footprints, and the development of assembly lines that broke complex manufacturing into simple, repeatable tasks (Thompson 1967; Biggs 1996). These changes were complemented by significant improvements in layout and workflow design, with factory spaces specialized to accommodate each distinct step in the production process (Taylor 2011).

This comprehensive restructuring of industrial space created unprecedented efficiencies in manufacturing operations, dramatically increasing productivity while standardizing output quality. This “Fordist” mass production of affordable manufactured goods fabricated with interchangeable parts on moving assembly lines leveraged cheap and reliable electricity to power individual machines and conveyor belts, enabling a continuous and highly efficient production process that surpassed the limitations of earlier line shaft-powered systems. Hyper specialized workers were able to better coordinate, allowing them to churn out large volumes of diverse consumer goods like appliances, electronics, furniture, and processed foods, as well as industrial equipment and even construction materials (Devine 1983; Hounshell 1984).

But for this to happen, the Second Industrial Revolution also had to witness changes to the industrial organization of industrial firms. As Chandler (1990) famously explained, increased vertical integration under multi-divisional corporate structures allowed companies to better control various stages of innovation and production, from R&D and design to manufacturing. This helped them achieve not only economies of scale but also of scope.

---

exacerbated the global economic downturn that contributed to the rise of fascism and other destabilizing events behind World War II (Kindleberger 1973).

For example, conglomerates like General Electric (GE) created an integrated ecosystem of more affordable products around electrical power.<sup>7</sup> It manufactured everything from the dynamos that generated electricity to the transmission wires that distributed it to the lightbulbs that consumed it. GE also developed comprehensive transportation networks, connecting neighborhoods via streetcars and linking cities with locomotives. Finally, it fabricated ovens and toasters for kitchens, radios and televisions for living rooms, curling irons and toothbrushes for bathrooms, and washers and dryers for laundry rooms. Meanwhile, a corporate-wide R&D program emerged within GE to help it develop these innovations.

Similarly, large companies beyond GE—including Eastman Kodak, B.F. Goodrich, Dow, DuPont, Westinghouse, RCA, US Steel, Unocal, and Goodyear—were able to embark on longer production runs by establishing a storage and distribution infrastructure that allowed them to hoard large inventories of raw materials and intermediate components. To sustain production speed, they developed systems to monitor the flow of these inputs and pioneered batch and continuous-process manufacturing. This allowed these companies to exploit standardized production techniques across disparate categories that included airplane engines, plastics, cannons, medical equipment, and oil-field drill bits (Chandler 1977).

Vertically integrated multidivisional firms like these also took direct control of suppliers, R&D, critical inputs, and distribution channels, allowing them to reach national, if not global, markets (Chandler 1977). They exploited railroads, telegraphs, steamships, and improved postal services to coordinate the flow of goods over larger territories, connecting numerous producers to reach a bigger pool of consumers (*ibid*).

This revolution in industrial organization extended to the retail sector as well. Mass retailers like A&P and Sears leveraged railroads and telegraphs to match wholesalers' purchasing power. By establishing internal purchasing organizations, investing in bulk storage and distribution networks, and buying directly from manufacturers, they eliminated the "middleman" and generated higher sales volumes—at lower margins—with increased stock turnover (Tedlow 1990; Chandler 1977). In turn, this strengthened their ability to negotiate lower prices with manufacturers, further reducing per-unit costs. They also exploited their distribution networks to reach consumers directly across the US, particularly outside the Eastern seaboard. And mass retailers' demands for standard products in bulk influenced how manufacturers designed assembly lines, reinforcing mass production (Chandler 1977).

As outlined in Chapter 1 of the book, the economic implications of these developments were as equally profound as those associated with the First Industrial Revolution. Overall, jobs required higher skills: a high school education prepared workers to master more sophisticated manufacturing techniques and college-educated workers became managers, engineers, chemists, and accountants (Goldin and Katz 2008). Indeed, the number of scientists and research engineers working in industrial labs also expanded dramatically (Lamoreaux and Sokoloff 1999). Unprecedented industrial growth coupled with a much more productive workforce culminated in an unprecedented explosion in living standards (McCloskey 2016; Gordon 2016).

---

<sup>7</sup> This paragraph closely draws on Gryta and Mann (2018).

## The Third Industrial Revolution

The Third Industrial Revolution key GPT was the centralized processing computer.<sup>8</sup> Its development rested on key technological breakthroughs, including the transistor, integrated circuits, microprocessors, personal computers, and user-friendly software. While integrated circuits first allowed for the rapid storage and processing of vast amounts of information, as well as the significant miniaturization of increasingly affordable devices, this revolution really began in 1971, when Intel invented the microprocessor, essentially a programmable computer within a chip. Rapid improvements in processing power that proceeded exponentially (more on this below) allowed personal computers to become smaller, cheaper, more powerful, and more versatile. The internet, digital networks, and portable wireless devices emerged downstream of this invention, fundamentally reshaping communication, commerce, work, media, and politics.

The transistor was the first key breakthrough. Made from semiconducting material, either Germanium or silicon, these miniature devices both switch and amplify electric currents.<sup>9</sup> In computers, they function as rapid on/off switches that represent the binary code essential to programming and the logic circuits of computers (see Isaacson 2014: 141).<sup>10</sup>

While Bell Labs invented the device, companies like Texas Instruments commercialized it—a topic I take up again in Chapter 5 of the book—crucially shifting from Germanium to silicon in 1954 to meet the high-temperature requirements of military hardware (Misa 1985). The US military and NASA immediately became the primary buyers, purchasing transistors in bulk to guide nuclear warheads and power the Apollo program. While this “guaranteed demand” drove down costs (O’Mara 2019), it also exposed a critical bottleneck: the “Tyranny of Numbers” (Riordan and Hoddeson 1997). Complex systems like ballistic missiles required thousands of

---

<sup>8</sup> A computer is fundamentally “a machine that [can] perform any logical task on any set of symbols” (Isaacson 2014: 108). This involves several key features: the ability to translate symbolic mathematical language into computer code; read-write memory for both storing program instructions and changing them in real time; and the capacity to use data and switch between instructions through variable-address program language.

<sup>9</sup> Broadly, transistors fall into two main categories: bipolar junction transistors (BJTs) and field-effect transistors (FETs). BJTs, such as the n-p-n junction transistor, regulate current by using a small base current to control a larger flow from an emitter to a collector. In contrast, Metal-Oxide-Semiconductor FETs (MOSFETs) utilize a voltage applied to an insulated gate to create an electrical field that controls conductivity. While they differ in operation, both allow a low-powered input to control a higher-powered output, enabling the amplification and switching functions foundational to modern computing (Horowitz and Hill 2015: 71–148; for the historical development of junction and field-effect transistors, see Riordan and Hoddeson 1997: 146–176).

<sup>10</sup> They operate as solid-state devices employing crystals to conduct current and feature at least three terminals, with a critical control terminal modulating current flow between the other two. When biased in a linear or active region—operating in a range where output is proportional to input—transistors amplify input signals, such as in radio receivers. Conversely, driving them between “cutoff” (fully off) and a low-resistance “on” state enables their function as binary switches (Horowitz and Hill 2015: 74–84; Sze and Ng 2007: 293–373).

transistors, each needing to be hand-wired to diodes, resistors, and capacitors. Workers, often women, assembled these tiny components using tweezers, where a single bad solder joint among thousands could render a missile useless. This precluded the emergence of economies of scale, rendering the devices unreliable and expensive (*ibid*).

The time was therefore ripe for an innovation that could enable multiple transistors and other components to be manufactured on a single semiconductor substrate, drastically reducing size and cost while increasing reliability. As I will detail further below, the US military, particularly interested in making components smaller, lighter, and more reliable, led a miniaturization and simplification drive (see O'Mara 2019). The integrated circuit demonstrated to engineers the immense potential of semiconductor technology for creating complex and compact electronic systems, well beyond transistors that were hitherto soldered together with cumbersome wires (Reid 2001).<sup>11</sup>

Monolithic integrated circuits revolutionized computing by simultaneously building and connecting transistors directly on silicon substrates (*ibid*). The reliable mass production of these devices was enabled by the groundbreaking planar manufacturing process, which was invented at Fairchild Semiconductor by Jean Hoerni and then refined by Jay Last (see Isaacson 2014: 184–185; Riordan and Hoddeson 1997: 257–262). Instead of hand-wiring components one by one, manufacturers could now chemically “print” intricate circuit patterns onto silicon wafers, allowing thousands of transistors to be fabricated simultaneously in a single batch process.<sup>12</sup>

The planar process’s ultimate triumph arrived in 1971 at the Intel Corporation, where the first commercially available single-chip microprocessor, the Intel 4004, was invented.<sup>13</sup> This represented a programmable computer within a computer. Unlike previous integrated circuits designed for specific tasks, the microprocessor allowed engineers to program all logic operations on a single, generic, and reprogrammable chip. Functioning as a central processing unit (CPU), it integrated instruction processing, memory access, and control logic onto one piece of silicon.

---

<sup>11</sup> The concept of a “solid block” circuit was first proposed in 1952 by Geoffrey Dummer of Britain's Royal Radar Establishment, but it was realized independently in 1958–59 by Jack Kilby at Texas Instruments and Robert Noyce at Fairchild Semiconductor. Kilby’s initial “solid circuit” used germanium and required some external wire connections, while Noyce’s “monolithic” version utilized silicon with printed aluminum interconnections on an oxide surface—the method that enabled modern mass production (Riordan and Hoddeson 1997).

<sup>12</sup> The planar process employs photomasks—essentially film negatives—to chemically print intricate circuit patterns on silicon wafers. It involves repeatedly layering thin semiconductor materials using chemicals, gases, and light. Through a series of precise exposures, conductive regions (formed by introducing dopants) and insulators (often silicon oxide) are created on the substrate, allowing layers to be selectively etched away or added to form complete functional units (see Riordan and Hoddeson 1997).

<sup>13</sup> Following Robert Noyce’s departure from Fairchild Semiconductor, he and Gordon Moore co-founded Intel in 1968; Andy Grove joined shortly thereafter. A team led by Federico Faggin—including Marcian (Ted) Hoff, Stanley Mazor, and Masatoshi Shima—were behind the microprocessor’s development (see Isaacson 2014: 177–181; for the definitive first-person account, see Faggin et al. 1996).

This paved the way for the bulk manufacturing of standardized computational power, effectively democratizing access to digital logic across countless applications (Ceruzzi 2003).

Initially, established firms undertook the commercialization of these technologies. Like the Second Industrial Revolution, the early digital era was dominated by large, vertically integrated companies, including AT&T, Texas Instruments, IBM, Xerox, and Hewlett Packard (Macher and Mowery 2004). These goliaths ran corporate research facilities—such as Bell Labs and Xerox PARC—that were generously financed with retained earnings; their labs focused on both practical engineering and basic science, introducing critical technologies like the Unix operating system, the C programming language, the computer mouse, and the laser printer (Gertner 2012; Hiltzik 1999).

However, the Third Industrial Revolution soon birthed a new institutional model: the venture-backed startup. The rise of Venture Capital allowed fledgling companies to capitalize on the appetite for digital technologies. This ecosystem was seeded by the "Fairchild Eight"—key personnel who departed Fairchild Semiconductor to establish new firms, including Intel, AMD, and venture firms (VC) like Sequoia Capital and Kleiner Perkins (O'Mara 2019). While each enterprise specialized in a particular segment, they remained interconnected through a shared pool of talent and professional networks (Saxenian 1994). Venture firms provided not just capital but "adult supervision," imposing managerial discipline and connecting startups to supply chains (see Mallaby 2022).<sup>14</sup>

It was in this ecosystem that Apple Computer rose on the scene. Partially financed by venture capital, it was the first company to bring a fully assembled personal computer to the mass market as a consumer appliance—eliminating the technical barrier of previous "kit" computers (Ceruzzi 2003).<sup>15</sup> The Apple II, released in 1977, offered color graphics and a friendly user interface. Then, in 1979, the release of the first spreadsheet software, VisiCalc, transformed the device into a serious productivity tool. Priced at roughly \$1,300—approximately \$6,900 in 2025 dollars—the Apple II moved computing from the hobbyist workbench to small businesses, households, and schools (O'Mara 2019).

As I outline in Chapters 4-6 of the book, the computer industry's structure eventually shifted from vertical silos to horizontal layers as the advent of the "Wintel" stack (Windows OS + Intel chips) standardized the interface between hardware and software. In the context of the constant exponential improvement in microprocessor performance known as Moore's Law (more on this below), which meant ever-falling costs for ever-improving hardware, Microsoft pioneered a platform with high R&D costs but nearly zero marginal costs of production, generating powerful supply-side economies of scale that attracted developers and users alike (Shapiro and Varian

---

<sup>14</sup> American Research & Development Corporation (1946), founded by Georges Doriot, was the first modern VC firm. A generation later, firms like Sequoia Capital and Kleiner Perkins were formed by the "Fairchild" diaspora to professionalize this model (see O'Mara 2019).

<sup>15</sup> Mike Markkula provided the initial angel investment; Don Valentine of Sequoia Capital invested the following year (Mallaby 2022: 83–86). Venrock participated in a subsequent 1978 funding round (Moritz 2009: 157–162).

1999).<sup>16</sup> Importantly, a PC appeared on nearly every office desk as the proliferation of the Windows standard further reduced the cost of computing and increased the utility of business software, fundamentally altering the daily tasks of the average worker (Autor, et al. 2003; Gordon 2016: 445–460). In Chapter 6 of the book, I explore the political economy of the various technologies that fed into this phenomenon, focusing particular attention on semiconductors.

Accordingly, just as the Second Industrial Revolution put a premium on a high school diploma, this new era biased the labor market in favor of college-educated workers who could operate complex digital tools (Goldin and Katz 2008). While computers boosted productivity, they also widened the wage gap between the highly skilled and the rest of the workforce (Acemoglu 2002; Goldin and Katz 2008); this militated against the wage compression of the mid-20<sup>th</sup> century, a topic I revisit in Chapter 11 of the book.

By the 1990s, the digital foundation paved by computers culminated in the commercialization of the Internet and, eventually, the smartphone and digital platform revolutions. While the PC commoditized computing power, the World Wide Web commoditized information distribution. Dedicated internet browsers and search engines drastically reduced the costs of finding, moving, displaying, and processing information (Bakos 1997). In doing so, these technologies effectively shattered the transaction-cost barriers of the analog age—collapsing the search, distribution, and verification frictions that had long served as bottlenecks to market entry and innovation.

By linking billions of discrete processors into a single, global network (see Greenstein 2015), the CDP architects moved human interaction into a digital environment where the marginal cost of data exchange approached zero. As observed by Bakos (1997) at the dawn of this shift, these diminishing search costs did more than just help buyers find sellers; they fundamentally reorganized the economy by enabling the "match" between diverse users that fuels multi-sided platforms. In Chapter 6, I document the industrial organization of the smartphone stack; in Chapter 8, I analyze the birth and consolidation of digital platforms, exploring how this wholesale disintermediation of content revolutionized the distribution of information and provided the high-volume training sets required for the AI revolution.

## INDUSTRIAL REVOLUTIONS ARE ABOUT GPTs

At first glance, these three transformations appear strikingly dissimilar. The First Industrial Revolution was powered by coal and steam, centered in Britain, and gave rise to the factory system. The Second harnessed electricity and petroleum, diffused across the US and Western Europe, and produced the vertically integrated, multidivisional corporation. The Third ran on silicon and software, clustered in places like Silicon Valley, and favored venture-backed startups competing within horizontal, modular industry structures. Each revolution also reshuffled the labor market differently: the first drew workers from farms to factories; the second rewarded the high school diploma; the third biased demand toward college-educated workers fluent in digital

---

<sup>16</sup> Microsoft later entered the advertising market and developed a robust cloud infrastructure (Azure), but its initial dominance was built on the zero-marginal-cost economics of packaged software.

tools. Even the signature consumer products—cotton textiles, automobiles, personal computers—seem to belong to entirely separate technological universes.

What ties these historic industrial revolutions together, however—and, as I explore later, what they share with the fourth one—is that they revolve around GPTs. Characterized by exponential technical progress and rapid product improvement cycles, as well as dizzying market expansion, GPTs are different than ordinary technologies like the zipper or the ballpoint pen. They eventually pervade the entire economy, spawn a myriad of complementary innovations and, eventually, after a pronounced delay, ratchet up productivity (see Bresnahan and Trajtenberg 1995). Outside of their economic effects, they also have profound impacts on society and culture (Perez 2002).

Realizing the commercial potential of GPTs requires significant upfront investments with long gestation periods, however (Perez 2002; Aghion and Howitt 1998). GPTs provide significant cost reductions and efficiency gains only over a considerable length of time, and after a relatively long period of protracted, non-linear diffusion. As Comin et al. (2006) have documented, adoption patterns for these technologies reveal characteristic S-curves: adoption accelerates sharply after crossing a critical infrastructure threshold, eventually tapering off as the technology approaches an upper plateau, a process I formalize in Chapter 11 of the book.<sup>17</sup>

GPTs aren't readily usable for several reasons identified in the literature.<sup>18</sup> First, the introductory vintage must undergo quality improvements, which require sustained R&D. Second, infrastructure must be built out and capital investments in machinery and tools must occur, as well as in knowledge and the training of workers to use the technology. Third, the technology must become standardized so it can be both produced and deployed at larger scales. Fourth, the development of complementary intermediate goods and components must occur, as organizations tend to switch from incumbent systems only after a critical level of supporting technologies makes the investment economically worthwhile. Fifth, to fully leverage a GPTs' potential, organizations must adapt their structures and processes—including hiring and training, and even their corporate culture.

## **Industrial Revolutions Require a Standardization Phase**

---

<sup>17</sup> In terms of a comparison of international adoption patterns, it is thus unsurprising that GPTs such as electricity, automobiles, and computers reached full penetration in their point of origin, the US, and that other countries that were also advanced adopters (e.g., Germany) show near-parallel growth lines once they pass key inflection points, gradually narrowing the gap over time. Even countries that started far behind eventually exhibit substantial growth rates, reflecting cross-border technology transfer and rising living standards. Moreover, different technologies spread at different rates. Newer innovations like the internet and smart phones diffused more rapidly in countries that were late adopters than older technologies. This implies leapfrogging effects where countries that adopted a GPT later sometimes used the newest generation of hardware or organizational methods to accelerate adoption (see Comin et al. 2006).

<sup>18</sup> See generally David 1990; Bresnahan and Trajtenberg 1995; Perez 2002.

Technological standardization marks a transition from a more “fluid” stage of technological development to a more predictable one. While the fluid phase is about introducing sweeping, foundational changes, innovation in the standardized stage becomes incremental: focused on refining small aspects of the established design, not reimagining core concepts. This allows some approaches to design and functionality to achieve dominance over others and for performance criteria to become more clearly specified (Utterback and Abernathy 1975; Clark 1985).

Only after new technologies become standardized and interoperability between versions increases that generic interfaces and best practices evolve, enabling far broader adoption that help transform the economy. Notable examples of technological standardization across each of the industrial revolutions embody this phenomenon.

Let’s start with the first one. The self-acting mule transformed textile production when its standardized designs allowed smaller mills to adopt automated spinning without employing specialized mechanics for custom machinery (Allen 2009). Similarly, the Linotype machine revolutionized printing once its standardized keyboard layout and casting mechanisms allowed newspapers and print shops to adopt mechanized typesetting without extensive technical expertise (Huss 2005). Bottlemaking underwent a similar transformation—Owens’ semiautomatic machines evolved from complex, custom installations requiring specialist operators into standardized units that smaller glass manufacturers could readily implement (Miller and McNichol 2012).

The Second Industrial Revolution is also instructive, especially when it comes to steelmaking. This era saw a surge in demand for steel, an iron alloy offering superior malleability and strength compared to cast iron.

Standardization made a big difference. Steelmakers eventually refined their diagnostic and calibration techniques through systematic trial and error in ways that standardized the approach. They honed techniques that furnished them with precise control over temperatures, movements, and timing sequences throughout the process—whether in pouring the molten pig iron, manipulating the converter’s tilt, controlling the air flow, or shaping and cooling the steel. The Bessemer converters gradually matured as standardized furnace designs, consistent chemical practices, and improved refractory linings emerged (see Menaldo 2021).

By the mid-20th century, international steel standards for chemical composition and tensile strength streamlined advanced techniques like basic oxygen steelmaking while ensuring consistent quality across applications—from railway construction to shipbuilding, machine manufacturing, and bridge building (Gilbert 2012; Nuvolari 2019). Plants clinging to older furnaces or partial improvements risked being outpaced in productivity and cost per ton.

In construction, uniform profiles for structural components like I-beams and H-beams provided engineers with predictable material properties, simplifying both the design of buildings and bridges and their fabrication. Over time, these developments underwrote a massive expansion in infrastructure investment (see Ollivier et al. 2014).

A similar process governed the innovations associated with the energy sources that powered the first and second industrial revolutions. Once boiler and engine designs became standardized and reliable, shipping lines could retrofit existing vessels to use steam without maintaining specialized engineering teams at each port (Ville 2004). As they crossed previously unimaginable horsepower thresholds, this precipitated a rapid decline in transportation costs. More efficient steam engines allowed shippers to more cheaply move massive cargos across continents and oceans, which incentivized the development of refrigerated shipping. Together, these technologies gave rise to trade in perishable goods at a global scale and international food brokerage. Moreover, they also facilitated international passenger travel and tourism. Similar standardization efforts created a truly integrated railway network (Gross 2016). The gradual move toward standardized gauges allowed rolling stock to move across lines and reduced the costly breaks-in-gauge that had previously forced passengers and freight to transfer between incompatible systems (Puffert 2002; Gross 2016).

Likewise, as I outlined in Chapter 1 of the book, electrification only achieved widespread adoption in manufacturing after standardized motors, uniform current specifications, and plug-and-play power transmission systems eliminated the need for each factory to essentially operate as its own power plant with unique mechanical configurations (Devine 1983). This itself depended on the standardization of electrical power delivery when the AC standard usurped the DC standard (David and Bunn 1988).

### *Computers and Standardization*

Perhaps the most instructive example of standardization processes that underpinned an industrial revolution is the evolution of computing and networking during the Third Industrial Revolution. Early computers were expensive, room-sized machines that operated in isolation, often requiring unique, proprietary code. For the digital revolution to scale, the industry had to converge on dominant architectures that allowed software and hardware to decouple.

As I document in Chapter 5, dominant semiconductor firms like Intel and Texas Instruments eventually drove this convergence. By establishing industry standards—most notably the x86 instruction set architecture for microprocessors—they fostered a collaborative ecosystem where software written for one machine would work on millions of others. This benefited both supply chain partners (suppliers of silicon wafers and specialized chemicals) and downstream competitors, creating a positive feedback loop that locked in the standard and lowered the cost of innovation for everyone else (see Gawer and Cusumano 2002).

The emergence of the personal computer (PC) in the late 1970s and early 1980s accelerated this shift. Standardized hardware architectures, such as the IBM PC standard, and the development of user-friendly operating systems like MS-DOS and later Windows, transformed computers from arcane tools into accessible devices for individuals and small businesses. Standardization fueled a surge in software development and broadened the PC's user base, creating a virtuous cycle of innovation and adoption.

The PC revolution enabled by these developments laid the groundwork for the next major technology that defined the Third Industrial Revolution: the internet. Its transformation from an

inscrutable network used primarily by academics and researchers that required significant technical expertise to navigate to the widespread use of a user-friendly product depended on a standardized foundation for data transmission across disparate networks. Specifically, TCP/IP enabled computer communication, HTML created a common data language, HTTP provided a platform to share this language, URLs established addressing conventions, and the World Wide Web offered a browsing portal.

Once these overlapping standards were in place, the technology's widespread adoption ensued. The interoperability and ease of use created by these common protocols and languages lowered barriers for developers and users alike, allowing a diverse range of applications and services to flourish and attract a massive user base (Greenstein 2015). The introduction of user-friendly web browsers with similar features like Mosaic and Netscape Navigator further democratized access, spurring the growth of e-commerce and online services.

The rise of the internet also created a rapidly increasing demand by organizations for a powerful and reliable server infrastructure to host websites, email services, databases, and other online applications (ibid). By extension, this catalyzed the development of on-site servers typically housed in dedicated server rooms or data centers within organizations that ran on operating systems like Unix, Linux, or Windows Server (West and Dedrick 2003).<sup>19</sup> Greater compatibility and interoperability between different hardware vendors was made possible by standardized server architectures (Baldwin and Clark 2000). This, in turn, facilitated the development of standardized server management tools, software, and networking technologies essential for connecting servers to the internet and to each other, including faster Ethernet standards, improved routing protocols, and sophisticated network devices like switches and routers (Mohindroo 2023).

The next stage of this evolution, and the twilight of the Third Industrial Revolution, saw the rise of cloud computing. Providers like Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP) began to offer Infrastructure as a Service (IaaS), Platforms as a Service (PaaS), and Software as a Service (SaaS).<sup>20</sup> Once again, standardization was critical: Cloud

---

<sup>19</sup> This development was itself contingent upon advancements in key hardware components: more powerful CPUs, increased RAM (Random Access Memory) capacity, larger and faster hard drives (later transitioning to solid-state drives or SSDs), and improved cooling systems to manage the increased heat generated by these high-performance components (Hennessy and Patterson 2019).

<sup>20</sup> IaaS gives users access to fundamental computing resources like virtual machines, storage, and networking, allowing them to manage the operating systems, middleware, and applications themselves. Examples of IaaS offerings include AWS EC2, Azure Virtual Machines, and Google Compute Engine. PaaS provides a platform for developing, deploying, and managing applications without the need to manage the underlying infrastructure, allowing users to focus primarily on coding and application development. AWS Elastic Beanstalk, Azure App Service, and Google App Engine are examples of PaaS. Finally, SaaS provides ready-to-use applications accessible over the internet through a web browser or other client application, completely abstracting away the need for users to manage any infrastructure or software. Examples of SaaS

computing relies heavily on virtualization technologies (like hypervisors) and established internet protocols (TCP/IP, HTTP, etc.) to deliver these services over the internet (Ledin 2020).<sup>21</sup> And, similar to the other stages of the Third Industrial Revolution outlined above, this shift further democratized access to powerful computing resources, allowing businesses and individuals to leverage scalable infrastructure—including virtual machines, storage, databases, and networking—without the need for large upfront investments in physical hardware, data center space, and ongoing maintenance (Armbrust et al. 2010).

In Chapters 5 and 6 of the book, I explore the standardization of Third Industrial Revolution technologies in greater detail and focus particular attention on the smartphone and its telecommunications infrastructure.

Having traced standardization's historical trajectory, we can now examine why it happens in the first place—or fails to. This matters greatly because the various examples I introduced above imply their obverse: a lack of process or product standardization can introduce a range of difficulties that slow the diffusion of GPTs. The absence of standards delays the development of complementary and compatible products. Alternatively, it may make it difficult to transfer knowledge and know-how about how to deploy and make best use of the new technology.

### *Political Economy of Standardization*

Basic microeconomics and canonical insights from industrial organization help explain differences in the timing of technological adoption. Some organizations may be more precocious than others if their industry is more competitive, the capital costs of adoption are relatively low, and their corporate culture embraces risk taking (see Mansfield 1961). The lesson is that firm-level characteristics can explain which organizations adopt first, but not whether a technology achieves widespread diffusion.

However, technological standardization and adoption dynamics are best understood through the lens of strategic interaction and coordination games, specifically the “Stag Hunt.” When all parties share identical interests—say, choosing a common technical standard—the socially optimal outcome (everyone “hunts stag”) dominates. However, the very expectation of what everyone else will do can pin players into a suboptimal outcome. In a Stag Hunt, there are two pure-strategy Nash equilibria: hunting stag together (high payoff for everyone, but disastrous for the lone hunter), and hunting hare alone (lower payoff, but safe even if your partner defects).

---

include Salesforce, Google Workspace, and Microsoft 365. On these general points see Mell and Grance (2011).

<sup>21</sup> Virtualization, enabled by hypervisors like Xen, KVM, and VMware vSphere, allows multiple virtual machines (VMs)—software emulations of physical computers—to run concurrently on a single physical server, maximizing hardware utilization and resource flexibility. The key internet protocols that make cloud computing possible are TCP/IP for reliable data transmission, HTTP/HTTPS for web-based services and APIs, and DNS for translating domain names into IP addresses, enabling clients to locate cloud servers (see Hwang et al. 2011).

Although the “hunt-stag” equilibrium is Pareto-superior, it carries risk: if you break ranks and go after stag while everyone else opts for hare, you get nothing.<sup>22</sup>

This is where focal points help. These are a solution to a coordination dilemma that actors tend to choose by default because it seems natural, special, or relevant (Schelling 1960). In the case of the Stag Hunt game, a focal point—such as a government-endorsed standard, a dominant firm's public commitment, or an industry consortium's declaration—can shift expectations so that all players anticipate others will “hunt stag,” thereby making it individually rational to coordinate on the Pareto-superior equilibrium.

However, it may also be the case that as path dependence deepens a collective expectation, the Pareto-inferior equilibrium may become locked in—once “hare” (a fragmented or inferior standard) is entrenched, coordinating a switch feels prohibitively risky. This explains why markets often stay stuck in inferior norms until an external focal pivot—such as a government mandate or a credible institutional commitment—reshuffles expectations toward the Pareto-superior outcome.

Similarly, in the context of early-stage technology diffusion, the market often faces a coordination game characterized by multiple potential equilibria—for example, the industry could standardize on *Design A* or *Design B*, but it remains trapped in a suboptimal state of fragmentation where neither alternative gains enough traction. In the absence of a dominant design, rational firms often delay investment to avoid the risk of adopting a “stranded” technology—a classic deadlock where the market splits between incompatible standards or stagnates entirely (Farrell and Saloner 1986). Consider once again the “fragmentation” of early 19<sup>th</sup> century rail gauges discussed earlier: without a coordinating body, regional lines remained trapped in a “hare” equilibrium of incompatible tracks, slowing down the railroad’s diffusion.

While I will explore the role of the state in breaking these kinds of technological impasses later in the section of this appendix, after that I will also discuss the logic of coordination games with distributional conflicts of interest (operationalized as a “Battle of the Sexes”) and how fights between interested parties over the spoils of getting on the same page to create common standards impact technological interoperability and continued innovation.

### **Standardization: Necessary, but Not Sufficient**

Besides standardization, new GPTs also require complementary investments to reach their full potential—everything from updated management structures to upgraded supply chains (David 1990; Bresnahan and Trajtenberg 1995). While competitive pressures drive firms to continuously refine and expand technology applications (Cohen et al. 2019), to integrate standardized technologies into their workflows and cultures effectively they must also rewire their

---

<sup>22</sup> In economics, an outcome is “Pareto superior” to another if at least one party is better off and no one is worse off. In the Stag Hunt scenario, the “Stag” equilibrium is Pareto superior to the “Hare” equilibrium because all participants earn a higher payoff by coordinating than they would by acting safely alone.

infrastructure and train their human capital. These investments can take years to implement effectively (Brynjolfsson and Hitt 2000). I now explore several examples of this phenomenon.

### *Steel*

The evolution of the modern steel industry provides a compelling example of the synergy between standardization and organizational adaptation.<sup>23</sup> Large-scale steel production came to rely on heavy machinery and assembly-line techniques to transform iron ore into finished products. However, revolutionizing the industry required more than just machinery; it called upon mills to train skilled smelters, codify workflows, and invest in massive new infrastructures.

As I outlined above, while early steelmaking involved byzantine purification processes, the Bessemer process perfected in the 1860s allowed for mass production by blowing compressed air through molten pig iron to burn off impurities.<sup>24</sup> While groundbreaking, this process was notoriously difficult to control, however. Early adopters struggled with inconsistent steel quality due to difficulties controlling oxygen content and impurities. Mastering the technique required significant expertise in temperature control, timing, and mechanical precision, leading many traditional producers to stick to older, slower methods like puddling or crucible techniques.

Diffusion required organizational learning. Steelmakers eventually refined their diagnostic techniques through trial and error, moving from ad-hoc experimentation to standardized furnace designs and chemical practices. Foundries shifted to complex two-stage melting processes—pouring molten iron directly into preheated converters—which required a total reorganization of the plant floor. Importantly, this "learning by doing" was often assisted by roaming technicians and consultants deployed by Henry Bessemer himself. Beyond teaching the core metallurgy, these experts helped clients install hydraulic systems, calibrate machinery, and modify coke-making ovens. Plants invested not only in physical infrastructure but in extensive employee training and cultural shifts to accommodate the relentless pace of the new batch processing.

### *Computers*

Computers are another quintessential example of how standardization requires organizational adaptation for a technology to consolidate and spread. Consider that the memory chips and silicon microprocessors that power personal computers were introduced in 1971. Yet, 20 years later, computer equipment still accounted for a negligible share of the nation's capital stock. Different sectors adopted computers at different times and for different reasons. When firms did pull the trigger, however, they shared one thing in common: they had invested in complementary changes to make the best use of these machines (David 1990).

Take retail as an example. In 1991, Walmart embedded a new software system—Retail Link—into its logistical operations, granting suppliers real-time access to sales and inventory data

---

<sup>23</sup> This narrative about the steel industry draws heavily on Menaldo (2021).

<sup>24</sup> Henry Bessemer's basic steelmaking patent was granted in 1856 to his firm Henry Bessemer and Company. Over the years, he also patented numerous improvements upon his original process. In the United States, William Kelley obtained a patent for a similar technique in 1857.

(Stalk, Evans, and Shulman 1992; Fishman 2006). This investment fundamentally transformed the superstore chain's supply chain from a more reactive, order-based system to a proactive, data-driven network where information flowed freely between the retailer and its suppliers, leading to significant improvements in efficiency, cost reduction, and responsiveness to customer demand. For it to work, however, Walmart had to make complementary changes to its hiring and training protocols, business processes, physical infrastructure, and software (Brynjolfsson and Hitt 2000).

Other large retailers soon followed suit, onboarding similar data-sharing platforms—but only after undergoing complementary changes of the sort Walmart undertook. This meant hiring specialists in demand forecasting and supply chain analytics to optimize inventory levels, alongside network engineers to maintain the technical backbone (Bresnahan et al. 2002). It also required investing in physical capital, such as robust data warehouses and automated reordering systems. Crucially, firms had to shift to frequent-delivery inventory management enabled by bar code scanning, a process that required not just new software, but a fundamental retraining of store managers and buyers to collaborate directly with suppliers on production planning (Holmes 2001; Abernathy et al. 1999).

### **GPTs Experience Exponential Improvements**

During previous industrial revolutions, once standardization took hold, the GPTs associated with these events experienced exponential—or at least quite rapid—gains across key performance metrics. These improvements were coupled with precipitous falls in quality-adjusted prices. Increased affordability further democratized access to these technologies, stimulating novel applications and catalyzing entirely new industries (Jovanovic and Rousseau 2005).

#### *The Paradigmatic Example: Microprocessors*

Famously, the Third Industrial Revolution was powered by Moore's Law, the observation—first articulated by Gordon Moore, Intel's cofounder—that the number of components (and later, transistors) on leading-edge integrated circuits rises at an approximately exponential pace, historically on the order of a doubling every year or two (Moore 1965; Burg and Ausubel 2021). In turn, this phenomenon drove remarkable improvements in the miniaturization, efficiency, and affordability of the computers these chips powered.<sup>25</sup> In 1971, Intel released the 4004, the first commercially available microprocessor, with just 2,300 transistors on a single chip. Although initially designed for calculators, the concept of a “computer on a chip” soon expanded to word processing, basic communications, and other general-purpose functions (Faggin et al. 1996). By 1978, Intel's 8086 processor had climbed to 29,000 transistors, laying the foundation for the IBM PC era. Exponential improvements continued during each of the subsequent decades. Table S.1 summarizes the major commercial accomplishments that occurred between 1971 and 2020.

### **Table S.1 Key Milestones in Processor Advancement**

---

<sup>25</sup> Market analysts track this progress in several ways, including: (1) transistors per microprocessor, (2) cost per transistor, (3) processor performance—often expressed in MIPS or FLOPS, (4) power consumption per transistor, (5) process node (nm) and die size, (6) memory capacity, (7) clock speed (GHz), and (8) economic output per transistor.

Year	Processor	Transistor Count	Significance
1971	Intel 4004	2,300	The world's first commercially available microprocessor
1974	Intel 8080	4,500	An 8-bit microprocessor that became the basis for many early computers
1978	Intel 8086	29,000	The first 16-bit microprocessor; established the x86 architecture
1982	Intel 80286	134,000	Introduced protected mode; used in the IBM PC/AT
1985	Intel 80386	275,000	The first 32-bit microprocessor; enabled multitasking
1989	Intel 80486	1,200,000	The first x86 chip to break the 1 million transistor barrier
1993	Intel Pentium	3,100,000	Introduced superscalar architecture
1997	Intel Pentium II	7,500,000	Improved multimedia performance (MMX)
1999	Intel Pentium III	9,500,000	Added SSE instructions for Internet/media
2000	Intel Pentium 4	42,000,000	NetBurst microarchitecture aimed at high clock speeds
2006	Intel Core 2 Duo	291,000,000	Shift to multi-core processing for efficiency
2008	Intel Core i7	731,000,000	Nehalem architecture; integrated memory controller
2011	Intel Core i7 (Sandy Bridge)	1,160,000,000	Integration of graphics on the same die
2017	AMD EPYC 7000	19,200,000,000	High-performance server chip utilizing multi-chip module design
2020	Apple M1	16,000,000,000	System-on-Chip (SoC) demonstrating the shift to high-efficiency ARM architecture

*Notes: Data for Intel processors (1971–2011) are official company figures. Data for 2017 and 2020 represent confirmed counts for AMD and Apple processors, reflecting the industry shift where Intel stopped publishing official counts for consumer desktop chips.*

*Sources: Intel Corp. (2011); Rupp (2018); Apple (2020); AMD (2017).*

Meanwhile, Figure S1.1 charts processor performance per inflation-adjusted dollar, revealing Moore’s Law. On the vertical axis, performance (measured in MIPS) is plotted on a logarithmic scale, while the horizontal axis spans 50 years of semiconductor progress. The data reveals a steep upward trajectory from the early 1970s to 2020. This means that every inflation-adjusted dollar spent on a microprocessor at any of the points in time depicted in this graph bought orders of magnitude more computing power and a relentless decline in computing costs.


While the conventional wisdom is that Moore’s Law proceeded inexorably, what actually powered transistor counts’ steep upward trajectory were deliberate and sustained efforts to improve both chip design and manufacturing methods (Burg and Ausubel 2021).<sup>26</sup>

Semiconductor firms determined to keep shrinking die sizes and pushing the limits of circuit design to increase transistor density deployed generous R&D money and engaged in constant experimentation (Galetovic 2021). Intel was the undisputed industry leader behind these efforts—a topic I elucidate in Chapter 5 of the book.

### **Figure S1.1 Charting Moore’s Law Over Time**

---

<sup>26</sup> Engineers were able to devise creative solutions to the formidable technical challenges posed by heat dissipation and electron leakage at ever-smaller scales, ensuring the viability of continuous transistor miniaturization. Firms poured resources into R&D to invent or refine new process technologies—examples include strained silicon, where it is physically stretched at the atomic level to let electrons move more freely; high- $\kappa$  dielectrics, which reduce power loss and permit thinner insulating layers; and FinFET transistor architectures, three-dimensional designs featuring thin “fins” for better control over electron flow. Working together, these breakthroughs steadily improved efficiency and speed while preventing the chip from hitting physical limits (see generally Hennessy and Patterson 2019).



Notes: MIPS (Million Instructions Per Second) is a measure of the execution speed of a computer processor. It indicates how many millions of instructions a processor can execute in one second. The data is for key processor models released between 1971 and 2020, selected based on their significance and availability of performance data. In cases where data for specific years was not available, values were interpolated based on trends and known data points. Processor list prices and performance figures are compiled by the author from manufacturer documentation and contemporaneous trade-press reporting. The costs were adjusted for inflation using the Consumer Price Index (CPI) values. The MIPS per inflation-adjusted dollar was calculated by dividing the processor performance (MIPS) by the adjusted cost in 2020 USD. The values were then transformed into their natural logarithms.

Sources: Intel Corp. (2011); Roser and Ritchie (2013); Moore (1965); Burg and Ausubel (2021); BLS (2025a); and author calculations.

As computing and communications capabilities spread in the wake of the plummeting price for ever-improving chips, consumers and businesses benefited greatly. In the second half of the 1990s, semiconductor-driven advances and improved processor technology drove sharp declines in constant-quality price indexes for information technology—a mechanism identified by Jorgenson (2001) as central to the late-1990s growth resurgence. Practically speaking, fully configured early personal computers such as the 1981 IBM PC cost roughly \$10,000 in 2025 dollars, yet delivered only modest computing resources (see IBM 2025). By the late 1990s, however, much more powerful and accessible PCs were available at about 1/10<sup>th</sup> the quality-adjusted price (McCarthy 2001).

During the early 21<sup>st</sup> century, progress continued apace. Continuous gains in transistor density allowed processors like the Intel Core series to pack billions of transistors, delivering massive increases in computational power and energy efficiency within ever-smaller chassis. This phenomenon enabled the development of lightweight, portable machines that retailed for \$300 to \$500 (in 2025 dollars) and accelerated the diffusion of interconnected personal computers into

businesses and households (Martin 2021), setting the stage for near-universal market penetration. By 2021, 95% of US households reported owning a computer (Mejía 2024). The world would never be the same.

Just as the steady cadence of Moore’s Law governed the Third Industrial Revolution, the Fourth is governed by its own set of “scaling laws.” As I demonstrate in several of the book’s chapters, during the early 2020s Artificial Intelligence models mirror this distinctive pattern of exponential performance gains coupled with plummeting costs. Specifically, I show that AI models were growing like gangbusters across key metrics, including training compute and parameter counts.

## INDUSTRIAL REVOLUTIONS’ MACROEFFECTS

The long run economic consequences of industrial revolutions are profound, as they are characterized by compounding growth in both the productive capacity of the economy and improvements in average living standards (McCloskey 2016). This is almost entirely due to increases made on the intensive, rather than extensive, margin (McCloskey 2016; Mokyr 2016). For example, consider that US Total Factor Productivity (TFP) grew at an unprecedented 1.89% annually between 1920 and 1970, which Gordon (2016) attributes to the widespread diffusion of GPTs such as electricity.<sup>27</sup> Consumers benefited twice over, enjoying new products and plummeting prices (Chandler 1990; Gordon 2016).

To visualize the power of this compound growth, we can hold quality constant and look at the cost of production. If annual TFP growth clocks in at roughly 1.9% over a 50-year period, the efficiency of the economy improves by a factor of 2.5. This implies that by the end of the period, a factory could produce the exact same widget using less than 40% of the inputs—labor, energy, and capital—required at the start. A famous precedent for this is the Ford Model T: between 1909 and 1924, the car itself barely changed, yet Ford’s relentless process innovations slashed the nominal price from \$950 to \$290. In real terms, the input intensity required to put a Model T on the road fell almost 70%, attesting to the “free lunch” of productivity growth (Hounshell 1984).

Yet, many economists have noted that the GPTs associated with industrial revolutions suffer from a “productivity paradox”: upon introduction, productivity growth lags the outsized expectations about their transformative impact; it may even fall pronouncedly (Solow 1988; Brynjolfsson et al. 2021). In the language I develop in Chapter 11, there is a persistent gap between a technology’s “Potential Yield” (its theoretical capacity to transform production) and its “Realized Yield” (the actual output gains observed in the economy).

I outlined some of the reasons for this phenomenon above. In the early stages, firms must undergo substantial reorganization before realizing the full benefits of innovation; as they learn,

---

<sup>27</sup> It’s little wonder then that the Second Industrial Revolution witnessed the ascendance of American economic supremacy. Indeed, as early as the 1890s, the U.S. had become the world’s largest economy and one of its most prosperous nations, surpassing the United Kingdom in technological prowess, industrial intensity, and overall prosperity (Nelson and Wright 1992). Consider that, by 1929, the American share of global motor vehicle exports exceeded 70%, signaling the country’s unparalleled industrial might (ibid: 1945).

train, and retool to adopt new GPTs, firms divert resources and time away from their everyday activities. While resources are redeployed to get the technology up and running and widely disseminated, learning by doing means inefficiencies pile up during the short term. As a result, measured labor productivity across an economy can stagnate or even dip while companies figure out how best to integrate new technologies. However, once the necessary adjustments and complementary innovations take hold and new technologies are widely deployed, productivity begins to climb—often sharply. During the transition process, productivity growth moves from the new technologies themselves to their use in society.

For example, during the 1950s, productivity growth moved from being dominated by manufacturing to being dominated by distribution (Field 2003). While the spread of new infrastructure, such as the interstate highway system, accounted for part of this, the biggest factor was containerization, which allowed the older infrastructure (ships and railways) to utilize the newer technologies and infrastructure (trucks and highways) in a way that made them both far more productive (Brooks et al. 2021).

The same was true during the Third Industrial Revolution. Central processing units (programmable microprocessors) were widely adopted throughout the US economy in novel ways. As I reviewed above, companies like Walmart pioneered innovations in supply chain management and warehousing—including the widespread implementation of barcodes, computer networks, and data storage—which subsequently spread to competitors like Target. These technological advances enabled retailers to forecast sales accurately, reduce inventories, minimize waste, cut costs, adopt standardized product labeling, and create interoperable production and sales information systems that could automatically share standardized data throughout global supply chains and between retail locations (see Bonacich and Hardie 2006: 108). The economy wide knock-on effect was faster productivity growth (Gordon 2016).

This dynamic partly explains why US labor productivity followed a J-curve trajectory.<sup>28</sup> It averaged only about 1.4% from 1970 to 1994, despite the emergence of computers and early internet services (Gordon 2016). During this period, firms were investing in IT hardware and software, but the organizational changes needed to leverage these technologies across the economy took another decade or more to unfold. However, once these organizational changes and complementary investments began to take effect, and businesses learned how to effectively leverage these technologies, US labor productivity experienced a significant upswing. It grew about 2.3% per year from 1994–2004 (Gordon 2016), demonstrating the delayed, powerful impact of the IT revolution.

In Chapter 11 of the book, I formalize these recurring historical patterns—slow diffusion, the friction of standardization, and the gap between potential and realized yield—into a unified economic framework. That model will demonstrate mathematically why the “Productivity Paradox” is not a bug, but a feature of the transition to a new GPT and foretells a similar trajectory for AI.

## **Automation and Jobs**

---

<sup>28</sup> The productivity numbers can be found in Jorgenson et al. 2008 and Gordon 2016.

In terms of employment, GPT waves have caused short-term disruptions requiring worker reskilling, but have ultimately created more jobs than they destroyed. When automation reduces production costs, the resulting scale economies and quality improvements can boost final-product demand, expanding both industry size and labor demand—though often accompanied by transformed job requirements and wage structures. Thus, automation has typically led to higher employment, albeit with significant changes in workers’ required skills and productivity expectations (see Autor 2015).

Consider the textile industry’s transformation during the Industrial Revolution. Before automation, textile production was highly labor-intensive, with spinners and weavers working from homes or small workshops using hand-operated equipment. When innovations like the Spinning Jenny (1760s), Water Frame (1770s), and Power Loom (1780s-1790s) arrived, they dramatically increased production efficiency. While weavers initially revolted, fearing job losses, the reality proved more complex. As cloth production costs plummeted, demand for textiles—especially cotton fabrics—soared. British textile mills expanded capacity and ended up hiring more workers overall, though for different tasks: preparing machines, handling raw materials, finishing and packing cloth, and maintaining equipment (Bessen 2015).

Indeed, while fewer workers performed basic weaving, more took on supervisory or technical roles managing multiple looms or conducting repairs that called on higher levels of human capital. As Bessen (2015) argues in the case of New England’s “Rhode Island System,” the process of developing expertise in textile production involved several distinct types of knowledge acquisition and skill development. Workers acquired specific procedural knowledge, such as learning the precise techniques for tying knots essential to the weaving process. They simultaneously developed physical proficiency through practicing manual operations to discover the most efficient movements, such as quickly replacing an empty shuttle without disrupting the weaving rhythm.

Perhaps most importantly, workers gained technical knowledge through continuous experimentation, learning critical adjustments like how to properly regulate tension on the warp threads to minimize breakage and mastering the complex coordination required to operate multiple looms simultaneously. This multilayered learning process combined explicit instruction, physical practice, and experiential problem-solving to create the comprehensive expertise needed for efficient textile production. Those who could operate and maintain machinery in this manner earned higher wages than simple laborers. The automotive industry offers another compelling example across two waves of automation.

In the first wave, Ford’s introduction of the assembly line in 1913 standardized production tasks and slashed costs, making cars affordable to a mass market. As demand exploded, automakers hired more workers to run assembly lines, operate stamping presses, and staff paint shops. The automaker’s famous \$5 daily wage—more than double the prevailing rate—was neither charity nor a mere recruitment tactic. It was a rational economic strategy grounded in the massive productivity surplus the assembly line created. Because the new continuous-flow process was intense and fragile, annual worker turnover had skyrocketed to nearly 370%. By sharing a portion of the technology-driven windfall with workers, Ford slashed these turnover costs, combated absenteeism, and secured the strict discipline required for mass production (Raff and

Summers 1987). While traditional mechanics were replaced by specialized line workers, new roles emerged for technicians, foremen, and industrial engineers who organized production, maintained equipment, and improved processes (Hounshell 1984).

Later, when robotic arms arrived in auto plants in the 1980s for welding, painting, and lifting, predictions of massive job losses again proved premature. As output and quality standards rose, carmakers needed more design engineers, robotics technicians, quality managers, and supply-chain coordinators. While robots took over repetitive tasks, humans shifted to roles requiring dexterity, judgment, or problem-solving. Workers programming robots or maintaining automated systems often earned above-average manufacturing wages (Graetz and Michaels 2018).

Famously, it is a widely repeated fact that while bank tellers faced automation pressure as ATMs substituted for some routine transactions, more bank teller jobs were ultimately created than destroyed: yes, the number of tellers per branch fell but banks expanded branches and shifted teller duties toward customer-facing service tasks rather than simply eliminating the occupation (Bessen 2015; Autor 2015; Klein 2019).<sup>29</sup> ATMs made it economical for banks to operate more branches and expand services like business advising and mortgage consulting. While each branch needed fewer traditional tellers, the growing branch network and broader service offerings created new roles in relationship banking, financial product sales, IT, security, and compliance. Many of these positions commanded higher wages given their enhanced skill requirements.

And even though from the 2000s onward optical character recognition and robotic process automation decimated data-entry roles, rendering the manual keying of information obsolete (Ramirez 2024), this destruction paved the way for creation: as data-entry positions withered, demand surged for the data analysts and engineers needed to govern these complex information systems (Jaiswal 2025).

Similarly, while desktop publishing software in the 1980s and 90s decimated the ranks of traditional typesetters and paste-up artists, it birthed the entirely new (and far larger) professions of graphic design and web development, demonstrating that automation often lowers the barrier to entry for creative work rather than eliminating it entirely (Bessen 2015).

These cases demonstrate how automation can trigger virtuous cycles: lower production costs expand markets, creating demand for new types of skilled labor that complement the technology. Job content evolves—often toward more technical and service-oriented roles—and overall employment can grow as industries scale up. While certain tasks may be automated away, expanding markets and evolving job roles can maintain or increase labor demand, particularly for workers who develop the knowledge and flexibility to work alongside new technologies. In Chapter 11 of the book I formalize this insight.

However, these virtuous cycles were neither automatic nor guaranteed. While the long-run view shows equilibrium and growth, the short-run reality of these industrial revolutions was often

---

<sup>29</sup> Pace the discipline's penchant for trotting out this old horse of an example (of an old occupation thriving after creative destruction), BLS (2025o) projects a 13% decline in teller employment over 2024–2034 as digital banking increasingly substitutes for physical branches.

characterized by chaotic friction. As the next section details, the transition to a new GPT is frequently beset by severe market failures—from coordination problems to underinvestment in basic science—that private actors cannot solve on their own. It is here that the state enters the stage, not merely as a regulator, but as an indispensable midwife to the new economy.


#### MARKET FAILURES MUST BE SOLVED FOR GPTs TO REACH THEIR POTENTIAL

While industrial revolutions motored by the emergence and diffusion of GPTs represent transformative periods in economic history that dramatically reshape production processes, business models, and social structures, they do not materialize automatically through market forces alone. Significant market failures at various stages of innovation mean that GPTs are underprovided or fail to diffuse. This necessitates strategic government intervention to help ignite and sustain industrial revolutions.

To help understand why, consider Romer's (1986) endogenous growth model, which is predicated on the notion that non-rival ideas generate increasing returns at the aggregate level. Unlike standard economic thinking, where diminishing returns typically curb long-term growth (e.g., Solow 1956), knowledge spillovers enable the productivity benefits of R&D to multiply across the entire economy.

Empirical studies lend support to this perspective. Researchers have found that the social return to R&D often exceeds the private return (Griliches 1992; Hall 1996). In other words, when a firm invests in creating new ideas or technologies, much of the benefit accrues not just to that firm, but to other firms and individuals as well. Indeed, Nordhaus (2004) estimates that innovators can capture only about 2.2% of the total surplus from innovation, with the remaining benefits accruing to society at large. Faced with such a meager slice of the pie, rational private actors will consistently underinvest in the basic research required to launch a GPT, leaving trillion-dollar bills on the sidewalk simply because they cannot privatize the pickup.

#### **Figure S1.2 The Underprovision of Ideas as a Market Failure and the Government's Role**



*Notes: Supply curves are drawn as linear and parallel for expositional clarity; the qualitative results hold under more general functional forms.  $Q_1$  = market equilibrium;  $Q_2$  = socially optimal equilibrium. The framework follows the standard Pigouvian treatment of positive externalities (Pigou 1920).*

Following Pigou (1920), Figure S1.2 explicates the logic of how this divergence between private and social returns leads to underinvestment in innovation—and how government intervention solves this market failure. The solid downward-sloping line represents the aggregate Demand for Ideas, which reflects the value consumers and society place on each additional innovation.

The graph displays two different supply curves. The higher, dashed line represents Private Marginal Cost—the full, unsubsidized cost that private firms must bear to generate new ideas. The lower, dash-dot line represents the Social Marginal Cost. This lower curve reflects the "effective" cost to society when the positive externalities of innovation are accounted for. The vertical distance between these two curves represents the per-unit positive externality—the spillover benefit that each innovation generates for the broader economy but that private firms cannot capture. Left to its own devices, the market settles at  $Q_1$ , where firms produce fewer ideas at a higher price ( $P_1$ ).

The hatched triangle represents deadweight welfare loss—the economic inefficiency that occurs when the market produces fewer ideas than socially optimal. It captures the lost value of the trades that didn't happen: at the high price ( $P_1$ ), downstream users are priced out of the market, and facing high private costs, innovators are unwilling to supply the good.

Two features of the graph determine the size of the deadweight loss.<sup>30</sup> First, holding the externality wedge constant, the elasticity of demand governs how much quantity responds to the price distortion. If demand for ideas is highly elastic—meaning downstream users are price-sensitive—then the quantity shortfall ( $Q_2 - Q_1$ ) is large, and the welfare loss from underinvestment is correspondingly substantial. If demand is inelastic, the quantity response is muted, and less social value is foregone. Second, holding demand elasticity constant, a larger vertical gap between Private and Social Marginal Cost amplifies the deadweight loss more than proportionally.<sup>31</sup>

The empirical literature suggests that both conditions hold for innovation.

On the externality side, estimates consistently find that the social return to R&D far exceeds the private return. Jones and Williams (1998) calculate that the social rate of return is at least 30 percent—implying that optimal R&D investment is two to four times actual investment. Jones and Summers (2020) estimate gross social returns to total U.S. R&D of approximately 67 percent, while Bloom et al. (2013) find social returns to private R&D of roughly 55 percent. These estimates imply a substantial wedge between what firms can capture privately and what their innovations generate for the broader economy.

Returning to demand elasticity, the demonstrated responsiveness of R&D activity to tax credits and subsidies suggests that innovation investment is indeed price-sensitive (Hall and Van Reenen 2000), reinforcing the case that the foregone welfare from underinvestment is substantial. This suggests that with the right government interventions—the details of which I will explain below—the market moves toward the socially optimal equilibrium at  $Q_2$ . By subsidizing R&D, the state effectively shifts the supply curve down to the Social Marginal Cost line. This allows more ideas to be produced ( $Q_2$ ) at a lower effective price ( $P_2$ ), thereby eliminating the deadweight loss and capturing the full value of the innovation for society.

---

<sup>30</sup> A deadweight loss is the loss in economic efficiency that occurs when mutually beneficial trades that would otherwise occur fail to take place. It represents potential value lost to both seller and buyers and thus society at large. For example, imagine you are willing to sell a concert ticket for any price above \$50, and a friend is willing to pay up to \$100 for it. A trade at any price between \$50 and \$100 would make you both better off, increasing economic efficiency in the process and make society better off too. However, if the only way to make the transaction is through a broker who charges a \$60 fee, no trade will occur. The potential value that both of you would have gained is the deadweight loss—in this case due transaction costs.

<sup>31</sup> This is because a wider externality wedge not only increases the height of the triangle but also expands its base: larger spillovers imply that the socially optimal quantity ( $Q_2$ ) lies further from the market equilibrium ( $Q_1$ ). The welfare loss thus scales roughly with the square of the externality wedge.

Fortunately, there are four critical areas where government action can overcome the market failures detailed above. These are: 1) undertaking or subsidizing the basic science research that inspires new inventions; 2) establishing intellectual property rights (IPRs) to incentivize the development of those ideas; 3) supporting the commercialization of the resulting technologies; and 4) accelerating mass adoption by investing in standardization, education, and infrastructure.

### **Subsidizing Basic Science Research**

If, as Nordhaus (2004) estimates, innovators capture only a meager slice of the value from applied R&D, the problem is exponentially worse for basic science. As Nelson (1959) argues, “basic research” is essentially the production of pure information. Since this information possesses the characteristics of a public good—it is non-rival (one person’s use doesn’t diminish another’s) and non-excludable (it is difficult to prevent others from utilizing it)—it defies standard market mechanics even more than commercial innovation does. As Arrow (1962) famously demonstrated, private entities systematically underinvest here because they cannot fully appropriate the benefits. Consequently, the marginal social benefit of basic research exceeds the marginal private benefit to the firm undertaking it (see also Stiglitz 1999).

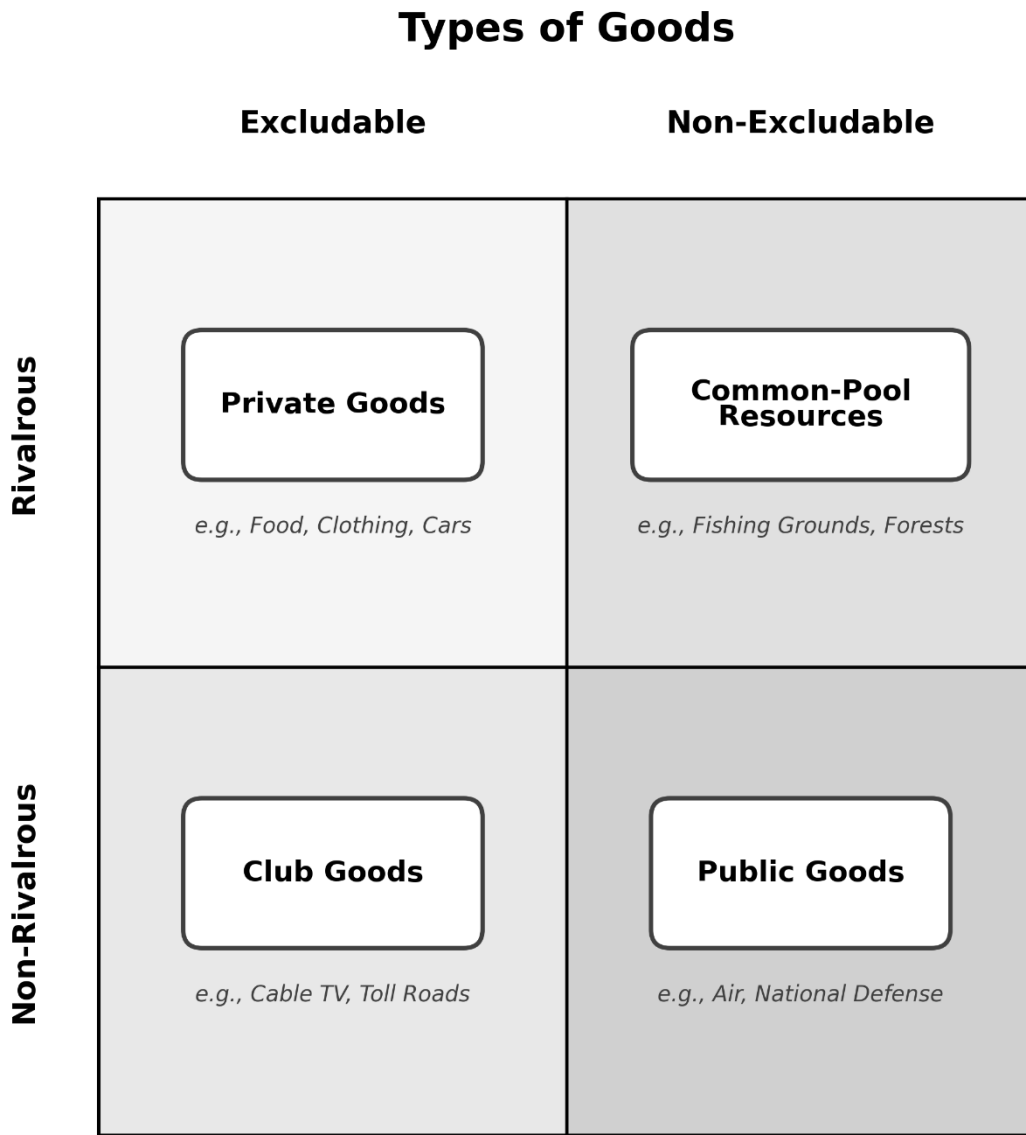
To understand why this is the case, Figure S1.3 compares the different types of goods, including the key distinctions between private and public ones. This 2x2 matrix categorizes economic goods based on two fundamental characteristics: excludability and rivalry.

Let’s start at the top. In the top-left cell, where both excludability and rivalry exist, we find private goods such as food, clothing, cars, and personal electronics. These function well in markets with private property rights since owners can prevent others from using them and one person’s consumption prevents simultaneous use by others. The top-right cell contains common-pool resources like fishing grounds, irrigation water, and forest resources, which are rivalrous, but difficult to exclude people from using. This makes them susceptible to overuse and suffer the “tragedy of the commons”—collective resource degradation.

Now I proceed to the bottom. The bottom-left cell represents club goods such as private parks, swimming pools, and toll roads, which can be controlled through access restrictions while remaining non-rivalrous up to a point of congestion. This allows for their efficient provision through membership fees or other access mechanisms. The bottom-right cell encompasses public goods, the primary focus of this discussion. Besides basic scientific research, this cell also includes national defense, street lighting, and clean air, all of which are neither excludable nor rivalrous—once provided, anyone can benefit without reducing availability for others.

A free-rider problem therefore bedevils public goods such as basic scientific research, and thus the marketplace underprovides it. This under-provision reflects a classic Prisoner’s Dilemma: even if all firms would be better off collectively funding basic research, the dominant strategy for any individual firm is to defect. If others contribute, one benefits most by free riding on their efforts; if others don’t, there’s no reason to waste one’s own resources. Rational self-interest thus drives everyone toward defection, even though universal cooperation would yield better outcomes for all. Therefore, without a binding mechanism to enforce cooperation, the group ends up at a suboptimal equilibrium where the research is never funded.

**Figure S1.3 Different Types of Goods**



*Notes: Excludability refers to the ability to prevent others from consuming a good. Rivalry refers to whether one person's consumption diminishes the quantity available for others. Private goods exhibit both properties; public goods exhibit neither. Knowledge and ideas—the outputs of basic research—are classic public goods: once produced, they can be consumed by additional users at zero marginal cost (non-rivalrous), and it is difficult to prevent others from accessing them (non-excludable). This combination creates the appropriability problem central to the economics of innovation (Nelson 1959; Arrow 1962).*

As I discuss in Chapters 5, 6 and 10 of the book, a “hold-up” problem with similar DNA occurs in the domain of technological commercialization: an investor refuses to make a relationship-specific investment if she fears that, once the investment is sunk, her counterparty can

renegotiate or expropriate the surplus (Klein et al. 1978; Williamson 1985). In both cases, cooperation unravels not because actors lack goodwill, but because the underlying payoff structure makes “defection” the safe, self-interested choice.

Because they can use coercion and inducements to incentivize cooperation, such as compelling citizens to pay taxes, governments play crucial roles in funding public research institutions and subsidizing universities that engage in it. Moreover, basic research suffers from highly uncertain returns—most of it cannot be commercialized (see Mazzucato 2013)—and requires relatively long investor time horizons, further strengthening the case for state intervention. This is more than a theoretical abstraction. Mokyr (2002) demonstrates how government support for institutions of scientific knowledge creation and diffusion has been vital to sustaining technological progress across modern history.

### *How These Lessons Apply to Industrial Revolutions*

The British government played a foundational role in fostering the Industrial Enlightenment, an era that valorized science and technical advances in the run up to the First Industrial Revolution (Mokyr 2002). It established and supported organizations that facilitated the free exchange of scientific knowledge and ideas. A prime example was the Royal Society, founded with a royal charter in 1660, which created an institutional framework for scientific inquiry that could inform practical problem-solving (Mokyr 2009). Similarly, Count Rumford and Joseph Banks co-founded the Royal Institution in 1799 to produce and disseminate scientific knowledge (Mokyr 2002; 2009).

This pattern of state support was scaled up dramatically in the 20<sup>th</sup> Century. The US Federal Government helped fund much of the basic research that underwrote the Third Industrial Revolution. Starting in World War II, the US military recognized computers’ immense potential to assist in codebreaking efforts and calculate ballistic missile trajectories (O’Mara 2019). To achieve these goals, and others like it, the government directed substantial public funding to several leading universities and research labs, including the University of Pennsylvania when it was developing the Electronic Numerical Integrator and Computer (ENIAC), one of the first computers capable of general-purpose electronic computing.<sup>32</sup>

During the Cold War, Washington, D.C., financed basic and applied scientific research in a multifaceted manner. The Defense Advanced Research Projects Agency (DARPA), housed in the Pentagon, primarily funded mission-driven research at universities and corporate research parks. The National Science Foundation (NSF) was created in 1950 to directly support more basic research at U.S. academic institutions that did not fall directly under the more mission-driven R&D conducted by other federal agencies (see Kleinman and Solovey 1995). The NSF issued generous grants to scientists working in major research universities. The Atomic Energy Commission, which later became the Department of Energy (DoE), financed and coordinated

---

<sup>32</sup> While the computational power of the ENIAC would later prove valuable for complex calculations related to nuclear weapons development, including the hydrogen bomb pioneered by Edward Teller in the post-war era at Los Alamos National Laboratory, its initial design and funding were specifically tied to wartime needs like improving artillery accuracy (O’Mara 2019).

R&D around nuclear power. The Office of Naval Research, as well as other arms of the Department of Defense (DoD), assumed a major role in funding and supporting academic research in the physical sciences. The National Institutes of Health (NIH) oversaw most health-related research, including biomedical research conducted in universities (Hurt 2015).

The federal government also mobilized substantial funds to build scientific installations for government-funded research. Prominent examples of publicly subsidized organizations that engaged in basic and applied scientific research were the Stanford Research Institute, the Lincoln Laboratory, a military funded research center affiliated with MIT, and the Rand Corporation in Santa Monica, California (see O'Mara 2019). These efforts significantly accelerated the development of key technologies that powered the Third Industrial Revolution, including semiconductors, the Global Positioning System (GPS), and the internet (see Mazzucato 2013).

Take semiconductors first. The transistor was enabled by significant basic research into how to manipulate the electrical conductivity of semiconductors rooted in solid-state physics and materials science that was strongly supported by the federal government. This included the creation of new knowledge around the band theory of solids, a fundamental concept in quantum mechanics that describes the allowed energy levels of electrons in solid materials (Riordan and Hoddeson 1997).

Fundamental research also seeded advances in integrated circuits and microprocessors; the continuous miniaturization of transistors and chip performance improvements hinged on breakthroughs in materials science, including the development of high-purity silicon, and the discovery of new semiconductor materials. The Office of Naval Research, and later the NSF, supported extensive investigations into the electronic properties of semiconductors like Germanium and silicon and advanced the science of growing pure single crystals of these materials with controlled impurities (doping) (Lécuyer 2006). Electrical engineering breakthroughs funded by the NSF and DARPA led to innovations in circuit design, allowing for more efficient and powerful processors (NRC 1999).

Advancements in photolithography, the process used to etch intricate patterns onto silicon wafers crucial for reducing transistor sizes and increasing chip density, were also funded by these agencies (Lécuyer 2006).

To help these myriad advances along, the Semiconductor Research Corporation (SRC), a non-profit consortium of leading semiconductor companies, government agencies, and academic institutions, was established in 1982.<sup>33</sup> It advanced cutting-edge, mid- to long-term research that advanced extreme ultraviolet (EUV) lithography, critical for manufacturing the most advanced microchips with nanoscale features (see Chapter 6 of the book). SRC also catalyzed breakthroughs in novel semiconductor materials beyond traditional silicon, such as silicon Germanium and gallium nitride, which offered improved performance for specific applications like high-speed transistors and power electronics. Furthermore, it contributed to advancements in both interconnect technologies—the microscopic wiring that links billions of transistors together

---

<sup>33</sup> For a detailed history of the SRC's research initiatives and cooperative structure, see Rea et al. (1997).

to form a functioning circuit—that improved the speed and efficiency of data transfer within complex chips and sophisticated modeling and simulation tools used to design and optimize these chips (Cavin et al. 2012).

The US government also heavily supported the development of GPS, which relies on a network of satellites and ground stations to identify both stationary and moving objects' locations. A host of federal agencies supported fundamental research in satellite technology, signal processing, and advanced mathematics, all of which enabled the technology to triangulate between multiple satellites to detect an object's position (Parkinson and Spilker 1996).<sup>34</sup>

Meanwhile, the US government bankrolled foundational work in computer science and network protocols that were crucial for the internet's development. The internet's precursor, ARPANET, was directly funded by DARPA; it financed groundbreaking research in packet switching theory—how to break down data into smaller packets, route them independently, and reassemble them at the destination—which paved the way for the efficient transmission of data across networks (Abbate 1999).

### **IPRs also Play a Critical Role in Fostering Industrial Revolutions**

Beyond undertaking or subsidizing pure research pursued for the sake of knowledge, a government may also assign and enforce IPRs to address the systematic under-provision of practical new ideas and inventions I outlined above. Figure S1.3 clarifies the logic behind how this solves the market failure by mapping different types of goods onto the excludability/rivalry framework.

The top row establishes the baseline for physical goods. In the top-left cell, where both excludability and rivalry exist, we find Private Goods such as Homes. These function well in markets because owners can prevent others from using them, and one person's occupancy physically prevents simultaneous use by others. The top-right cell represents Common-Pool Resources like Common Grazing Areas, which are rival (one person's animals consume grass others cannot use) but difficult to exclude others from accessing. As noted earlier, this often leads to resource depletion or the "tragedy of the commons."

The bottom row reveals the unique dynamics of the knowledge economy. The bottom-right cell represents Public Goods, specifically Laws of Science. These are neither excludable (difficult to prevent access once published) nor rival (one person's use doesn't diminish another's ability to use it). Because markets underprovide these goods, government funding is required.

The bottom-left cell is where IPRs enter the picture. It contains patents, copyrights, trademarks, and trade secrets. By legal design, these are akin to Club Goods: they are excludable through legal mechanisms (the patent right), but remain inherently non-rival (multiple people could use the invention simultaneously without diminishing it).

---

<sup>34</sup> DARPA funded early research in satellite technology and signal processing; NASA contributed to satellite technology research; and the NSF supported fundamental research into mathematics underpinning the system's logistics (Mazzucato 2013).

**Figure S1.3. How IPRs Make Ideas into Quasi-Private Goods**

		<b>EXCLUDABLE</b>	
		<b>YES</b>	<b>NO</b>
<b>RIVAL</b>	<b>YES</b>	<b>HOME</b>	<b>COMMON GRAZING AREA</b>
	<b>NO</b>	<b>PATENT</b>	<b>LAWS OF SCIENCE</b>

*Notes: Rivalry refers to whether a good is "used up" by consumption; Excludability refers to whether access can be restricted. Markets function naturally in the Top-Left (Private Goods). Market failures occur in the Top-Right (overuse/depletion) and Bottom-Right (underproduction/free-riding). Intellectual Property Rights (Bottom-Left) are a government intervention designed to correct the Bottom-Right failure by creating a proprietary asset out of non-rival knowledge.*

By artificially shifting an invention from the "Public Good" quadrant to the "Club Good" quadrant, the government creates a proprietary asset that incentivizes private investment. By

endowing ideas with excludability in a ring-fencing approach that establishes clear boundaries around a unique combination of technical claims that potentially bring to life new products and services, patents in particular incentivize innovation. Akin to ownership over physical assets established by a deed to a house or business or piece of equipment, they allow capital holders to secure a return on investments in time and resources that carry opportunity costs.<sup>35</sup>

Most importantly, well-assigned and enforced IPRs underwrite the commercialization of innovation. Like any property right, IPRs represent a bundle of privileges that not only spell out how access to an asset is restricted, but also outline rights to use, divide, and transform it, as well as the right to earn income from the asset and contract with others to license control of it or sell it. In turn, these privileges can underpin “a market for ideas” where divisible and tradeable rights can be exchanged, leased, or transformed through arms-length contracting (Haber 2016).

In this vein, Spulber (2015) argues that patents enhance market efficiency through exclusion, transferability, disclosure, certification, standardization, and divisibility. These attributes reduce transaction costs and promote competition in the invention marketplace. Patents create a “market for innovative control,” where IP owners have rights to determine how their inventions are developed and used, helping allocate inventions to their highest-value applications. Patents also facilitate financial separation between inventions and inventors, enabling inventors to secure funding through various channels. They function as intangible real assets that appear on corporate balance sheets, affecting firm value both through direct earnings from licensing or implementation and through growth opportunities. Inventors can even post patents as collateral to raise money to commercialize their ideas.

Patents foster specialization and help new supply chains materialize. Both their ability to exclude and their public disclosure incentivizes different parties—investors, entrepreneurs, assemblers, distributors, retailers, and marketers—to seek out original inventors to help bring new products to market (Kieff 2006). The bottom line is that an inventor can specialize in their comparative advantage, coming up with new ideas, and other market players can specialize in what they do cheapest—with the fewest opportunity costs to their time and resources—whether that’s assembling or marketing or distributing the ideas’ applications (Spulber 2015). Moreover, patent agents, bankers, lawyers, and patent assertion entities help reduce the transaction costs of contracting incurred by these parties when they help create the supply chains that commercialize new innovations (Kieff 2006; Spulber 2015).

Patents also make it easier for the different parties that embody an innovation’s supply chain to gain access to essential knowhow that inventors acquired by “learning by doing” and “trial and error.”<sup>36</sup> To put the invention into practice, they cannot rely solely on information available in the patent itself, nor can they exclusively lean on their knowledge of basic science, exposure to technical literature and exhibitions, and previous work experience. However, patent licensing

---

<sup>35</sup> To be sure, a robust system of IPRs does not operate in a vacuum. It relies on a broader institutional framework that includes impartial courts, the rule of law, and deep, liquid capital markets—the “background conditions” that allow corporations to form, raise capital, and enforce contracts (see Haber 2016).

<sup>36</sup> This and the subsequent two paragraphs draw strongly on Menaldo (2021).

contracts outline how the tacit knowledge critical to translating ideas into goods and services will be conveyed from licensors to licensees, including how a licensee will gain access to plans, drawings, blueprints, machinery, services, and onsite tutorials and training.

These contracts can also help licensees adapt processes and products to accommodate differences in raw materials and other inputs; and help them address logistical challenges. The latter includes initial implementation efforts, ongoing maintenance issues, and continued technological enhancements. Moreover, as licensees themselves develop improvements to the technology upon adopting and adjusting it, the knowledge transfer process can become bidirectional: licensees may patent these enhancements and they themselves lease these to the original inventors.

### *Patents are not Monopolies*

Finally, I hasten to emphasize that patents' role commercializing ideas stands in contrast to a simpler understanding of how they promote innovation. That is, that they are government-imposed barriers to entry that allow firms to recover their fixed, sunk costs by setting prices above their marginal costs for a delimited period, usually between 15 and 20 years (Posner 2005). In this view, while patents help solve the free rider problem inherent to the production and exchange of ideas because they punish the type of copying and duplication that allows rivals to avoid incurring fixed R&D and learning by doing costs, they create deadweight loss because they allow inventors to ration supply and thus increase prices.

However, this assumes market power that rarely exists in practice. If the cross-elasticity of demand between a patented product and its potential substitutes is high, the firm cannot maintain significant Lerner markups (Tirole 1988; Haber 2016).<sup>37</sup> With the notable exception of pharmaceuticals—where precise chemical recipes lack functional substitutes (Levin et al. 1987)—there is scant evidence that patents confer genuine monopoly power. Moreover, because the patent system mandates immediate disclosure, it enables rivals to build new supply chains and “invent around” the core claims (Cohen et al. 2000; Furman et al. 2021). Thus, theoretical “static efficiency” losses are bounded by competition, while the “dynamic efficiency” gained by incentivizing commercialization remains the primary engine of growth (Nordhaus 1969).

### *Some Evidence for these Ideas*

---

<sup>37</sup> I provide the formal derivation of the Lerner monopoly condition in Chapter 11 of the book, but the intuition is essential here. The Lerner Index, defined as  $L = (P - MC)/P$ , measures a firm's market power—its ability to set prices ( $P$ ) above marginal cost ( $MC$ ). Standard price theory dictates that this index is inversely proportional to the price elasticity of demand faced by the firm ( $L = 1/\varepsilon$ ). Thus, if a patent exists but rivals can “invent around” it to create substitutes, the elasticity of demand ( $\varepsilon$ ) rises. As  $\varepsilon \rightarrow \infty$ , the firm's market power ( $L$ ) approaches zero, and the optimal markup ( $\mu^*$ ) collapses toward the competitive price, regardless of the legal patent right.

While the British patent system underwent important reforms in 1852 that considerably strengthened it, its earlier incarnation played a critical role in fostering innovation.<sup>38</sup> Exploiting archival evidence to challenge the view that patents were either ineffective or peripheral to the First Industrial Revolution, Bottomley (2014) convincingly shows that they incentivized the invention and commercialization of signal technologies such as Boulton and Watt’s separate condenser steam engine. Bottomley documents how, by the 1780s, a more “modern” understanding of patents emerged—rather than a royal privilege that conferred economic rents, they evolved into a social contract between inventors and the public, whereby the former could exclude others (securing profits and licensing options) in exchange for the full disclosure of their discoveries, thus promoting learning and follow-on invention.<sup>39</sup>

The US Patent Act of 1793 improved upon the modern British patent system by simplifying patent filing and IP enforcement and reducing patenting costs by 95% (Khan 2005). This ushered in the era of “the great inventor” during America’s early industrialization. As the transaction costs of purchasing or leasing patents plummeted, US firms eschewed in-house R&D and instead “outsourced innovation” to serial independent inventors.<sup>40</sup> In turn, this created vibrant, albeit geographically segmented, technology markets in which patent agents served as crucial intermediaries linking firms in search of new ideas to commercialize and original inventors (Lamoreaux and Sokoloff 1999). Similarly, IPRs were also important during the second and third industrial revolutions.

Indeed, they played a pivotal role in fostering the commercialization of the ENIAC. In 1947, J. Presper Eckert and John Mauchly, the ENIAC’s developers, filed what became the ENIAC patent, ultimately granted by the USPTO as US Patent Number 3,120,606 (Eckert and Mauchly 1964)”. While the legal strength of this specific patent would later be contested, due in part to how broad it was, as it covered everything from electronic memory to arithmetic units, the very

---

<sup>38</sup> The Patent Law Amendment Act of 1852 reformed the British patent system by simplifying the application process, reducing fees, and making it easier for inventors to protect their inventions.

<sup>39</sup> To be sure, during this era technological innovations were frequently shared and diffused between inventors, entrepreneurs, producers, and distributors through the vibrant exchange of knowledge and knowhow and industrial fairs (Allen 1983; Nuvolari 2004). This “culture of collaboration” complemented a decentralized production system marked by significant outsourcing and reliance on cottage industries. In the book’s later chapters, I explore how a culture of collaboration emerged during the Third Industrial Revolution on the back of strong IPRs.

<sup>40</sup> Prominent examples include Thomas Edison, who held over 1,000 patents across telegraphy, electric lighting, and sound recording; Alexander Graham Bell (telephone); Eli Whitney (cotton gin); Samuel Morse (telegraph); and Cyrus McCormick (mechanical reaper). However, as Lamoreaux and Sokoloff (1999) emphasize, the system also supported hundreds of lesser-known serial inventors who made their livelihoods by developing and selling patented technologies to manufacturing firms.

act of filing it marked the definitive transition of the computer from a wartime government project to a private commercial asset.<sup>41</sup>

This move to secure proprietary rights was the fulcrum upon which the computer industry turned. As Stern (1981) documents, Eckert and Mauchly's ability to attract early capital from the American Totalisator Company and secure contracts with the Census Bureau hinged entirely on their claim to exclusive ownership of the technology. The patent application effectively "assetized" their tacit engineering knowledge, transforming it into a transferable property right that could be valued on a balance sheet and sold to investors. Without the shield of potential patent protection, the high-risk capital required to scale production from a university lab to a commercial factory would likely never have materialized (Norberg 2005).

Importantly, Eckert and Mauchly left the University of Pennsylvania's Moore School of Electrical Engineering after a dispute over patent rights. The university's IP policy would have required them to assign all their inventions to the institution, which conflicted with their commercial ambitions (Lukoff 1979). Armed with the patents they had amassed during their work on the ENIAC, in December 1947 the duo essentially established the first startup of the computer era, the Eckert-Mauchly Computer Corporation (see Ceruzzi 2003).<sup>42</sup>

The company's formidable portfolio included concrete technical components alongside their broader system claims. In 1947, Eckert and Mauchly had also applied for a patent on their mercury acoustic delay-line electronic memory system, later granted as US Patent Number 2,629,827 (Eckert and Mauchly 1953). Basically, by introducing a "regenerative recirculation loop" that kept data pulses alive by constantly cycling them through the mercury and refreshing them electronically, the patent asserted rights to the machine's version of random-access memory (RAM). And unlike the contested ENIAC patent, this invention provided a stable IP foundation and served as the technical backbone for the company's first major commercial device: the groundbreaking UNIVAC, their Universal Automatic Computer (Weik 1961). In turn, it became the first device to gain widespread acceptance as a reliable computer memory system—a fundamental building block of modern computing that was widely adopted in early commercial architectures (Williams 1997).

IPRs continued to play an essential role during the evolution of semiconductors. While in 1947 the transistor was invented by John Bardeen and Walter Brattain at AT&T's Bell Laboratories, it was in securing US Patent Number 2,524,035 for the technology in 1950 (see Bardeen and Brattain 1950) that this scientific breakthrough became a defining economic asset. The patent established a sweeping legal claim over the very concept of solid-state amplification. By securing exclusive rights to the fundamental principle of using semiconductor materials to control electric current, this document effectively served as the "birth certificate" of the modern

---

<sup>41</sup> In 1973, a U.S. District Court declared the patent invalid after it found that prior work—particularly that of John Vincent Atanasoff—had anticipated key ENIAC patent claims (see Mollenhoff 1988).

<sup>42</sup> The Eckert-Mauchly Computer Corporation sold itself to Remington Rand in 1950. Through a series of corporate mergers and acquisitions, the rights to the ENIAC patents eventually passed to Sperry Rand (formed by the merger of Sperry Corporation and Remington Rand in 1955).

electronics industry, creating a singular choke point through which all subsequent semiconductor development had to pass.

Rather than maintaining exclusive control over this technology, however, AT&T shared it widely through a strategic licensing program launched in 1952. As Jack Morton, a leading Bell Labs engineer, explained: “We realized that if this thing was as big as we thought, we couldn’t keep it to ourselves and we couldn’t make all the technical contributions. It was in our interest to spread it around. If you cast your bread on the water, sometimes it comes back angel food cake” (Tilton 1971: 75-6). Moreover, internal memos revealed that Bell Labs engineers understood that “by involving engineers around the world in the evolution of the device—making it better, cheaper, more reliable—the hope was that everyone would profit from the advances, especially the Bell System” (Nagler et al. 2022: 839).

In that spirit, AT&T’s approach to licensing went beyond merely leasing their patent. The company actively facilitated technology transfer by organizing educational symposia.

The first event, held in 1951, specifically targeted military users and applications and attracted nearly 300 guests for a five-day summit (Bell Laboratories 1951; Nagler et al. 2022). The symposium’s proceedings were published as *The Transistor*—a volume that came to be widely known as “Ma Bell’s Cookbook” or “the semiconductor bible” because it codified the tacit knowledge required to build the devices (Nagler et al. 2022).

In 1952, Bell Labs organized a more extensive nine-day Transistor Technology event for over 100 representatives from 40 companies that had paid the \$25,000 fee to license its transistor patent (ibid). The symposium included not just theoretical presentations, but hands-on training and even a visit to Western Electric’s transistor manufacturing plant in Allentown, Pennsylvania. AT&T’s willingness to share their manufacturing know-how, which complemented their inclusive patent licensing strategy, precipitated the rapid diffusion of transistor technology (ibid).

Meanwhile, Texas Instruments exemplifies how licensing enabled new market entrants to become industry leaders in the semiconductor space. While the company emerged in 1951 from Geophysical Service Incorporated, a firm that manufactured equipment used in seismic exploration and defense electronics, merely a year later it was among the first companies to acquire a patent license to produce Germanium transistors from Western Electric, the manufacturing arm of AT&T (Texas Instruments 2021).

Texas Instruments later spearheaded the mass production of affordable memory chips and licensed its technology widely, including to Japanese companies during the early 1980s, when such chips were integral components of consumer electronic products like the Sony Walkman. Starting in the mid-1980s, Japanese (and one Korean) semiconductor firms began paying Texas Instruments substantial licensing fees after the company asserted key Dynamic Random Access Memory (DRAM) related semiconductor memory patents in U.S. litigation and settlements (Kinukawa 2006).<sup>43</sup>

---

<sup>43</sup> Incidentally, US Patent Number 3,138,743, the Kilby patent, formally titled “Miniaturized Electronic Circuits,” was filed with the USPTO in 1959 and awarded to Texas Instruments in

Patenting and patent licensing were also critical to Silicon Valley's most storied semiconductor firm, Intel. Even before the microprocessor era, the company had accumulated several patents over core memory technologies, including DRAM, and integrated circuit designs. During the 1970s, Intel forged cross-licensing agreements with AMD, IBM, National, Texas Instruments, Mostek, Siemens, NEC, and several other semiconductor firms covering memory chip designs and manufacturing processes.

Moreover, cross-licensing became an industry-wide pattern (Grindley and Teece 1997; Hall and Ziedonis 2001). For example, when MOS Technology developed the 6501 and 6502 microprocessors to compete with Motorola's 6800, the latter sued the former. The dispute was settled with MOS withdrawing the 6501 and the parties agreeing to a cross-license of microprocessor patents (Electronics 1976). This pact helped ensure that multiple firms could produce compatible parts for large customers (Griffin 1988).

Patent licensing arrangements helped standardize and diffuse key microprocessor technologies. When IBM was developing the original PC, released in 1981, it insisted on having at least two suppliers for its microprocessor to ensure supply chain stability. This requirement effectively compelled Intel to license its technology to AMD as a "second-source" guarantee, enabling AMD to legally produce and sell chips that were x86 pin-compatible. This had far-reaching consequences; it helped cement the x86 architecture as the dominant industry standard for PCs spanning into the 1990s and early 21<sup>st</sup> century across several generations of chips—the Intel 8086, 8088, 80286, 80386, and 80486 (Grindley and Teece 1997).

The consequences of this development were far reaching—and among them was a fundamental transformation of the computer industry's industrial organization. By transforming proprietary inventions into industry-wide standards, firms like Texas Instruments and Intel effectively positioned themselves at the center of a vast, invisible web of technological interdependence. As I show in Chapter 5 of the book, these “superstar” firms evolved into the central hubs of an innovation commons, dictating the pace and direction of the Third Industrial Revolution. And I show in Chapters 3 and 6, IP, in the form of Standard Essential Patents, were pivotal to the creation and consolidation of the smartphone supply chain that helped bring about the Fourth Industrial Revolution, as they made interoperability and sustained digital innovations possible.

However, as Branscomb and Auerswald (2002) document, even when IPRs are well-defined and enforced, private capital markets frequently fail to fund the early-stage, high-risk technologies that underpin industrial revolutions. I now turn to explaining why this is the case and how the government can help address this problem.

---

1964 (see Kilby 1964). While this invention was not officially granted by the Japanese patent office until 1989, Japanese companies that licensed it began paying Texas Instruments royalties decades before that. Robert Noyce also filed a patent for his version of the integrated circuit in 1959 for Fairchild Semiconductor; it was granted by the USPTO in 1961 as US Patent Number 2,981,877 (see Noyce 1961). After a legal battle over “priority,” the two firms recognized each other as co-inventors and cross-licensed the technology in 1966 (see Isaacson 2014: 173–180).

## **Governments Must Also Solve Other Market Failures**

Of course, the relative scarcity of new ideas given an outsized demand for them is nested within a larger problem—the chronic under-provision of public goods. Because private actors cannot easily exclude others from using basic infrastructure or the general pool of skilled labor, they tend to underinvest in both (see the discussion of Figure S1.2). Historically, the U.S. government has stepped in to fill this gap, providing the physical and human capital foundations upon which private innovation can build upon.

For example, its efforts have extended into providing supportive infrastructure. As David and Wright (1997) document, the federal government bolstered the mining sector by subsidizing the training of geologists and commissioning USGS maps that identified mineral deposits. Furthermore, during the 19<sup>th</sup> Century, the federal government allocated 5% of its spending to capital investments in harbors, roads, and canals. Following a wave of state bankruptcies in the 1830s, states began issuing bonds to finance infrastructure, catalyzing a precipitous decline in transportation costs (see Taylor 1962). As a result, industrialization took off precipitously at a massive scale (ibid; Chandler 1977). And perhaps the quintessential modern example is Eisenhower’s interstate highway system, where a similar phenomenon was unleashed.<sup>44</sup>

Yet another classic example of the state solving a coordination failure to “crowd in” private investment is the Tennessee Valley Authority (TVA). In the early 1930s, private utilities viewed rural electrification as financially irrational due to high fixed costs and low immediate returns. By underwriting the initial capital outlay for hydroelectric dams and transmission lines, the federal government didn’t just light up farms; it lowered the input costs for heavy industry. As Kline and Moretti (2014) document, this public investment ignited a sustained boom in private employment that outlasted the subsidies themselves. Rather than crowding out private capital, the provision of this public good created the necessary conditions for it to flourish.

In the modern era, the US government was instrumental in bankrolling the infrastructure that underpins the digital economy (Mazzucato 2013). ARPANET (Advanced Research Projects Agency Network), created in the late 1960s, was the first operational packet-switching network (O’Mara 2019). Building on its success, the NSF established NSFNET in the mid-1980s to link supercomputing centers, accelerating scientific collaboration (ibid). To further buttress this, the federal government supported the expansion of the internet backbone through initiatives like the High-Performance Computing and Communication Act (US Congress 1991) and various rural broadband programs.<sup>45</sup> In Chapter 8 of the book I discuss other initiatives like these.

---

<sup>44</sup> Fernald (1999) demonstrates that the construction of the system was a unique, unrepeatable shock that significantly raised the productivity of vehicle-intensive industries, effectively explaining the robust productivity growth of the 1950s and 60s.

<sup>45</sup> Under the American Recovery and Reinvestment Act of 2009, the Broadband Technology Opportunities Program provided approximately \$4.7 billion in grants to expand broadband access and adoption in unserved and underserved areas (US Congress 2009). This program supported the deployment of infrastructure, enhancement of public computer centers, and promotion of sustainable broadband adoption projects. The Broadband Initiatives Program,

Public investments in education have been comparable in scale and impact.<sup>46</sup> The Morrill Land-Grant Acts established technical universities that helped train the skilled land surveyors, miners, and engineers who unearthed and refined the minerals and energy that powered American industrialization (David and Wright 1997).<sup>47</sup> Later, the G.I. Bill expanded higher education after World War II, providing the skilled human capital that helped make the Third Industrial Revolution possible. And, in 1965, President Lyndon B. Johnson significantly expanded access to higher education by signing the Higher Education Act, which established federal aid programs for colleges and universities (US Congress 1965). Following this legislative milestone, states across the country invested in building new community college campuses and, between 1970 and 2016, enrollment in American higher education institutions more than doubled, growing from about 9 million to about 20 million students (Goldin and Katz 2008).

This allowed universities to become talent incubators; Stanford University's industrial park, for instance, nurtured hundreds of companies, from Varian to Facebook (Isaacson 2014: 156; see also O'Mara 2019). In general, the Third Industrial Revolution matured across innovation clusters—from Silicon Valley to Route 128—anchored by well-funded research universities (Menaldo and Wittstock 2025). The upshot was that both basic and applied research flourished; as Chapter 9 of the book documents, this helped nourish the AI revolution too.

### *The Government's Role in Financing Innovation*

Hall and Lerner (2010) identify why there is often a relatively large wedge between the rate of return required by an entrepreneur investing her own funds and the one sought by external investors. First, a quite extreme level of uncertainty often pervades R&D investments, especially those with relatively long maturities. Second, asymmetric information and adverse selection distort the market: while inventors possess more accurate knowledge about their projects' viability and value than potential investors, those who actively seek outside funding may perversely possess the least marketable ideas.<sup>48</sup> Consequently, investors may demand a higher

---

another component of the 2009 stimulus package, allocated about \$2.5 billion in grants and loans specifically for rural broadband infrastructure projects (see *ibid*). Separately, the FCC's Connect America Fund provided financial support to service providers to expand broadband services to rural and high-cost areas where the market alone may not make those investments worthwhile (FCC 2011).

<sup>46</sup> As Goldin and Katz (2008) argue, America's educational system evolved in numerous, small, and fiscally independent school districts that competed for students. It was built on several key principles: public funding that provided free education for all; non-sectarian public schools that maintained separation between Church and State in both financing and control; gender neutrality in access; and an open and forgiving system designed for mass education.

<sup>47</sup> This built on a vibrant universal public primary and secondary education system; by 1850, over 56% of school-age children (ages 5-19) were enrolled in educational institutions, and approximately 90% of white adults had achieved literacy (Goldin and Katz 2008).

<sup>48</sup> Moreover, Arrow's "Information Paradox" exacerbates this friction. Because ideas are easily stolen once shared, "[f]irms are reluctant to reveal their innovative ideas to the marketplace and the fact that there could be a substantial cost to revealing information to their competitors

risk premium for R&D investments compared to conventional assets, as distinguishing promising R&D projects from unpromising ones presents greater challenges than evaluating short-term or lower-risk opportunities (ibid: 614).

Finally, moral hazard concerns associated with the divergence between the interests of shareholders and managers can depress investment in R&D. While managers may divert scarce resources toward activities that provide them with personal benefits (such as corporate empire-building or workplace amenities), they may also exercise excessive caution and avoid investing in risky, yet valuable R&D projects to avoid incurring blame if these were to go awry (ibid: 615). Governments have historically addressed the credit market imperfections associated with these dilemmas through various mechanisms and, in the process, underpinned important GPTs. This includes early-stage funding for commercialization efforts when private markets go missing.<sup>49</sup> Various historical examples are particularly interesting.

Take the Board of Longitude, established by the British Parliament in 1714. It offered substantial financial incentives—up to £20,000 (equivalent to millions today)—for solving the critical navigation problem of determining longitude at sea (MacLeod 1988). This approach effectively financed John Harrison’s groundbreaking chronometer development in the face of private investor skepticism about its viability; over the course of decades, Harrison received over £23,000 in government payments that sustained his work, ultimately setting the stage for safe, scalable ocean navigation during the First Industrial Revolution (Sobel 1995).

Similarly, during the mid-to-late 18<sup>th</sup> century, the Society of Arts (later Royal Society of Arts) used government-supported prize competitions to stimulate innovations in agriculture, manufacturing, and chemistry when traditional financing was unavailable (Mokyr 2009). The result was the rapid diffusion of "useful knowledge" and practical micro-inventions—such as improvements in textile machinery, new agricultural tools like the scythe, and safer industrial chemical processes—that fueled the productivity gains of the First Industrial Revolution.

In terms of a more recent example, in targeting the notorious “valley of death” between initial concept and commercial viability, the US Small Business Innovation Research (SBIR) program helped finance several startups that went on to play a pivotal role in the digital revolution. Those include Qualcomm, which received SBIR funding in the 1980s to develop its digital wireless technology before going public, and Symantec, which utilized early government grants to advance software capabilities that proved integral to natural language processing (Lerner 1999)—an AI capability I examine in detail in Chapter 9.

### *The Government’s Role in Promoting Standardization*

---

reduces the quality of the signal they can make about a potential project” (Hall and Lerner 2010: 615; see Arrow 1962).

<sup>49</sup> Other mechanisms I will not elaborate upon here include several fiscal instruments such as R&D tax credits, accelerated depreciation allowances—which incentivize firms to upgrade their capital stock by allowing them to deduct the cost of new assets from their taxable income more quickly than those assets actually wear out—and relatively low capital gains tax rates (see Hall and Lerner 2010; O’Mara 2019).

Similarly, market mechanisms alone also often fail to achieve optimal technological standardization due to the coordination problems I outlined earlier and competing private interests (Farrell and Saloner 1985). Governments can play critical roles in overcoming these standardization related market failures.

First, they can act as a coordinator or regulator, explicitly mandating a unified standard or convening standard-setting bodies to harmonize conflicting technical specifications. Consider the government-led standardization of railroad gauges in the 19<sup>th</sup> Century and electric power systems in the early 20<sup>th</sup> Century (Shapiro and Varian 1999).

Or take the internet's underlying architecture. In the 1980s, the world faced a "protocol war" between the US-developed Transmission Control Protocol/Internet Protocol (TCP/IP) and the international OSI standard. The US federal government resolved this coordination failure by mandating the adoption of TCP/IP across its military and research networks, effectively establishing a dominant design that the private sector subsequently adopted (Abbate 1999). Later, its role in establishing organizations like ICANN ensured the stable and unified management of critical internet resources (O'Mara 2019).<sup>50</sup>

Second, by providing massive, guaranteed demand for a specific specification as a "customer of first resort," the government can make a specific design the salient choice, coordinating private expectations and tipping the market toward a single, dominant equilibrium. To secure these contracts, rivals are compelled to converge on the government's specifications, forcing the interoperability and reliability that the private market was too fragmented to produce on its own (see Mazzucato 2013).

Consider that the US military and NASA were the primary buyers of early transistors, purchasing them in bulk—and often at premium prices that commercial buyers balked at (Misa 1985). For example, Texas Instruments secured major defense work for its first generation of integrated circuits; this included winning the bid for the Minuteman II missile guidance system, the first major commercial deployment of its chips (Texas Instruments 2021). Earnings from these government contracts were then reinvested in R&D to commercialize digital devices such as pocket calculators (Reid 2001).

Finally, there is the idea of the government as a nudger. Government agencies like the US's National Institute of Standards and Technology (NIST) facilitate standardization processes that enable compatibility, interoperability, and widespread adoption of new technologies. Rather than impose these standards coercively on private parties, they serve as a neutral arbiter, convening industry stakeholders to develop voluntary consensus standards. This approach solves the

---

<sup>50</sup> TCP/IP (Transmission Control Protocol/Internet Protocol) is the fundamental suite of communications protocols used to interconnect network devices on the internet, defining how data is packetized, addressed, transmitted, and routed. ICANN (Internet Corporation for Assigned Names and Numbers) is the non-profit organization responsible for coordinating the internet's naming system (DNS) and IP address allocation, ensuring that every domain name maps to the correct unique address.

coordination problem I outlined above by establishing a clear technical “focal point” that firms can adopt with confidence, knowing their competitors are likely to do the same (Tassey 2000).

The rise of voluntary Standards-Setting Organizations (SSOs) perhaps embodies this role best. Over the 21<sup>st</sup> century, many governments moved from being direct architects to strategic arbiters, working behind the scenes to support a system where private firms collaborate to define technical standards (like 3G, 4G, and Wi-Fi) while relying on specific property rights rules—such as Fair, Reasonable, and Non-Discriminatory (FRAND) licensing terms—to balance incentives for innovation with the need for universal access. In Chapter 3 of the book, I show how the US federal government’s involvement in helping private actors develop wireless telecommunications standards in this way helped underpin the digital revolution.

### *Market Failures around Technological Commercialization*

Even after technological standardization occurs, significant barriers to the widespread adoption of GPTs often remain. These include adjustment costs, organizational inertia, and skill mismatches. Exacerbating these problems is the fact that firms tend to earn lower revenues and enjoy slimmer profit margins when making the initial investments associated with technological acquisition and installation (Brynjolfsson et al. 2021). Moreover, hiring specialists to facilitate this process is expensive. Indeed, consider that 50% or more of R&D spending consists of the wages and salaries of highly educated scientists and engineers, whose “efforts create an intangible asset, the firm’s knowledge base, from which profits in future years will be generated” (Hall and Lerner 2010: 612).

Yet, this is far from a unilateral dilemma, as the logic of coordination failure extends beyond technical protocols to these complementary investments as well. Firms often suffer from collective inaction, waiting indefinitely for other organizations to make the first move when it comes to training workers or building infrastructure (see Hoff 1999). For instance, a firm might hesitate to invest in specialized training if it fears that newly skilled employees will be poached by competitors who didn’t incur those training costs—a classic free-rider problem. Conversely, workers may hesitate to specialize in a new technology if they aren’t sure there will be enough employers to hire them. This mutual hesitation creates a trap where the necessary ecosystem of skills never materializes.

What’s more, the economic stakes associated with all these transaction costs are significant. As Brynjolfsson and Hitt (2000) demonstrate, for an organization to realize efficiency gains from new technologies, they must make complementary investments in organizational capital, business processes, and especially human capital. The longer it takes to install a new technology, the greater the short-run economic pain: productivity often temporarily declines as resources are diverted from production to learning and reorganization (Brynjolfsson et al. 2021). Consequently, workers may have to accept lower initial wages while acquiring necessary skills before they can fully translate the new technology into higher output. This is a phenomenon I revisit in Chapter 11 of the book.

The US federal government historically undertook several activities to break this deadlock.

During early American industrialization, it stimulated the diffusion of the telegraph by providing land grants that incentivized the expansion of lines across the country (Du Boff 1983). It also used the agricultural extension service to promote inventions such as the mechanical reaper by demonstrating the technology to farmers and facilitating the sharing of best practices (Huffman and Evenson 1993).

In the modern era, the government encouraged federal laboratories to actively collaborate with industry through CRADAs (the Technology Transfer Act and Cooperative Research and Development Agreements). These legal frameworks helped diffuse sophisticated innovations such as lasers; while early R&D occurred in federal labs, these policies made it easier for private companies to license foundational patents, leading to widespread applications in medicine and telecommunications (Bozeman 2000). Finally, the Manufacturing Extension Partnership (MEP) provided technical assistance to local manufacturers to expedite the transition from manual drafting to digital workflows using Computer-Aided Design (CAD) and Computer-Aided Manufacturing (CAM) software (Shapira 2001).

Another example of Washington’s ability to spread advanced technology during the Third Industrial Revolution, and in doing so mitigate the coordination problem outlined above, is SEMATECH (Semiconductor Manufacturing Technology), a non-profit consortium established in 1987 by the federal government and 14 leading U.S. semiconductor firms—including Intel, IBM, Texas Instruments, and Motorola. These otherwise rivals pooled resources and expertise to jointly invest in more advanced manufacturing processes for both logic and memory chips. By facilitating collaboration to standardize and improve products while reducing effort duplication, SEMATECH helped diffuse several innovations such as advanced lithography techniques, improvements in chip materials, and enhanced semiconductor manufacturing processes. This intervention helped the U.S. industry enhance yields and reliability, effectively reversing the market share losses to Japan that had characterized the previous decade (Irwin and Klenow 1996).<sup>51</sup>

## INDUSTRIAL REVOLUTIONS AND DISTRIBUTION

While this section of the appendix has primarily focused on how markets fail due to simple coordination or commitment problems, when promoting industrial revolutions or grappling with their consequences, governments must also address distributional conflict. In game theory, this can be conceptualized as a coordination game involving how to divide the social surplus created when social actors coordinate: when everyone benefits from working together but disagrees over

---

<sup>51</sup> This is not to say that the SEMATECH strategy was without its critics. As Rodgers (2025) notes, the consortium delayed next-generation equipment sales to nonmembers for a year to “slow Japanese competition,” a tactic that potentially stifled innovation among smaller US chip firms while offering questionable benefits to the larger incumbents.

how to divide the gains. Specifically, each side prefers a different focal outcome, and without a neutral arbiter, the stronger side can impose its favored rules, leaving the weaker either acquiescing or bargaining endlessly. This creates a risk that mutually beneficial cooperation will collapse entirely.

The classic example, the Battle of the Sexes, involves a couple who both want to spend the evening together but have different preferences: one prefers the opera, the other prefers a boxing match. While they both receive the highest payoff from coordinating on the same location (rather than going alone to their preferred event), they disagree on which location that should be. In political economy, this analogy describes industries that must coordinate on a single standard to unlock a market (the surplus) but fight over the specific rules that determine who captures the lion's share of that surplus.

Unlike the pure coordination problems outlined earlier in this section of the appendix, Battle-of-the-Sexes games yield equilibria favoring one side. Internalizing this logic, Social Conflict Theory emphasizes how power asymmetries shape institutions: actors who are relatively more powerful choose the rules that deliver policies that they benefit from the most. When actors' interests diverge, those with greater resources impose rules guaranteeing them a larger slice of the surplus, even at the expense of overall efficiency, because "they are in a position to accept more risk, because they have more withholding power, and because other people are less willing to sanction them if they violate a norm" (Acheson and Knight 2000: 213). These rules are designed to give them an advantage and reflect unbalanced political power, not the mutually perceived attractiveness of the most efficient solution to a collective action problem (Acemoglu et al. 2005).

When grappling with distributional conflicts, governments can play three crucial roles. First, by establishing binding revenue-sharing or pricing formulas—whether through regulated interconnection rates in telecommunications or catch-share schemes in fisheries—public authorities guarantee that each participant receives a transparent, predetermined portion of the surplus, so that even the weaker party has confidence in cooperation. Second, by convening jointly governed standards bodies or licensing pools, they create credible forums where all stakeholders have a seat at the table and decisions must meet supermajority or consensus thresholds, preventing any one actor from unilaterally rewriting the rules. Third, by embedding lock-in mechanisms—such as statutory veto points, fixed-term appointments on regulatory boards, or mandatory review periods—governments protect cooperative arrangements from ex post capture by better-resourced interests.

Under this approach, the state does not "pick winners" but rather designs processes and default outcomes that align divergent incentives, transform zero-sum struggles into positive-sum agreements, and allow innovation-driven coordination to flourish without perpetual renegotiation or outright domination by the politically powerful.

### **The Third Industrial Revolution as an Example**

In several of the chapters of the book I discuss how, by establishing binding rules that guaranteed each participant a predictable portion of a growing surplus, the US government repeatedly found ways to realign incentives, ensuring cooperation across several regulatory domains.

For example, in Chapters 5 and 6, I explore a sharp distributional conflict rending the smartphone supply chain that pitted upstream patent owners seeking to maximize royalties versus downstream implementers aiming to minimize licensing costs. This dynamic embodies a Battle of the Sexes style coordination dilemma, where all parties benefit from a uniform set of interoperability standards, but differ on what set of standards to adopt and how to share the producer surplus associated with a bigger and more lucrative market for products that can communicate with each other. To help address the conflict, the federal government stepped in and used a variety of tools to coax industry rivals to cooperate under the aegis of voluntary standard-setting bodies governed by supermajority formulas. In essence, it compelled industry adversaries to abide by transparent FRAND (Fair, Reasonable, and Non-Discriminatory) IP licensing rules that prescribed predetermined royalty formulas and gave its blessing to transparent dispute-resolution forums, both of which locked in a particular division of the producer surplus.

However, as we shall see ahead, the specific distribution of benefits in these Battle-of-the-Sexes arrangements was inherently unstable: during the CDP era, the federal government intervened to tilt the scales toward one faction, only to subsequently reverse course as the balance of political power shifted.

## CONCLUSION

The first three industrial revolutions share a common through line: groundbreaking GPTs such as steam-powered engines, electricity, and microprocessors gradually evolved from niche, highly criticized inventions into mature technologies with multiple applications. They reshaped production, commerce, and daily life. Once they permeated downstream industries, they also ushered in impressive booms in economic productivity that improved living standards and unlocked significant expansions of leisure time.

But, as this section of the appendix illustrates, industrial revolutions do not unfold automatically: even a compelling new technology can languish when faced with high switching costs, competing standards, entrenched skill sets, or meager private investment in R&D.

Industrial revolutions therefore hinge on (1) standardization, which ensures interoperability and economies of scale; (2) organizational adaptation, through which firms retool operations and retrain workforces; and (3) government support that solves multiple market failures across several domains. That includes providing strong patent protections for ideas and financing for early-stage innovation. It also includes basic research funding, infrastructure buildup, education and training, and help with technological standardization and diffusion efforts.

In many ways, the lessons from the past three industrial revolutions I outlined in this section of the appendix set the stage for the rest of the book. Capturing AI’s full potential will demand not only scientific progress and entrepreneurial ingenuity, but also the slow, unglamorous work of standardizing tools and interfaces, financing complementary investments, and rebuilding organizations so that frontier capability becomes routine in day-to-day production. These are topics I take up in Chapters 11 and 12 of the book.

Each industrial revolution also encountered a “productivity paradox”: measured output often sputters before a new technology’s benefits can be fully realized. Companies initially pour resources into learning, reorganizing, and complementary investments—which can dampen productivity metrics—yet eventually, once the reorganizations are complete, technologies diffuse broadly, bringing exponential cost reductions and long-delayed but dramatic gains in performance and affordability. Whether steam, electricity, or the microprocessor, once the core GPT is standardized and widely adopted, subsequent improvements tend to follow rapid or exponential curves, fueling fresh applications, surging demand, and new industries.

In Chapter 11 of the book, I will explore this phenomenon alongside other macroeconomic and socio-economic shifts, including wages, jobs, profits, and distribution. There, I formalize the key historical mechanism from this section of the appendix: the gap between what a frontier GPT can do and what the economy can realize before standards, complements, and organizations catch up. I will inform these modeling exercises with the historical evidence reviewed here: that, while automation can displace specific tasks, growing markets and new skill demands often increase overall employment. Thus, workers who adapt to the new technology and cultivate relevant skills frequently see improved wages and job opportunities as income inequality eventually narrows.

However, before we get there, several chapters of the book move from the general historical patterns of the past to the specific origins of the Artificial Intelligence revolution.

Chapter 3 begins this investigation by exploring the unique policy environment that birthed Silicon Valley. I explore how, between the late 1970s and 2000s, a broad “populist-statist” consensus was supplanted by a new paradigm. Earning strong bipartisan support, this new approach prioritized evidence-based regulation, established novel property rights over data and software, and aggressively reduced transaction costs. This specific institutional mix solved the market failures described above in a new way, eventually laying the legal and financial groundwork for the AI boom.

## ONLINE SUPPLEMENTARY APPENDIX TO HISTORY'S MOST REVOLUTIONARY INNOVATION, SECTION 2

In the popular imagination, the dominant firm is often cast as a hoarder—a dragon guarding a fortress of intellectual property designed to starve competitors and stifle the market. Yet, the history of the Third Industrial Revolution reveals a reality that defies this caricature. The era that birthed the personal computer and the mobile internet was not defined by walled gardens, but by a vibrant, invisible architecture of shared knowledge. Far from the stagnant “dual economy” that plagued Japan during the same period (see Katz 1998; McKinsey Global Institute 2000), the U.S. high-tech sector thrived on a peculiar form of “selfish benevolence”: a system where technological superpowers cemented their dominance not by hiding their inventions, but by broadcasting them, turning their proprietary breakthroughs into the industry’s common language.

To understand how superstar firms in high-tech sectors became their industries’ innovation leaders during the Third Industrial Revolution (1976-2006), it is helpful to first consider a parallel example from biotechnology. The Cohen-Boyer patent for recombinant DNA, the foundational technology of the entire biotechnology industry, was filed in 1974 by Stanford and the University of California, covering the fundamental technique for gene splicing—the ability to cut and paste DNA. The universities' strategy of offering broad, non-exclusive licenses spurred massive horizontal technology transfer, as hundreds of startups and established pharmaceutical firms licensed the patent to build the entire modern biotechnology industry. Rivals were compelled to license not by a standard, but by the sheer necessity of using this foundational technique to compete. This phenomenon, in turn, drove vertical technology transfer, creating a new market for specialized suppliers of enzymes, lab equipment, and other tools needed to perform the gene-splicing methods outlined in the patent.

This section of the appendix will show that innovation in high-tech industries during this era, with semiconductors leading the way, took on a similar networked character. However, unlike in biotechnology, where universities often took center stage, in this innovation commons firms collaborated and competed to create the foundational technologies of the digital world. Their collective efforts produced crucial precursors to the modern smartphone, from the first personal computers that connected to the internet to the later generation of wireless devices like mobile phones and personal digital assistants.

However, the term “collective” should not be confused with either “centralized” or “egalitarian.” On the one hand, there was no hierarchical, formal system that orchestrated technological creation and transfer from the top down. To be sure, while Chapter 6 of the book documents efforts of that ilk in terms of establishing telecommunications standards around cellular phones, the decentralized approach to innovation I explore here was about technologies that were much more foundational. On the other hand, only a few firms were clear technological standard-bearers; for example, companies such as Intel were responsible for introducing and standardizing technologies like the CPUs that power laptop computers, notwithstanding the fact that a critical mass of “follower” firms implemented them and sometimes helped “incrementally” improve them.

Strikingly, as I will show ahead, the pronounced inegalitarianism characterizing technological initiative and influence that marked the computer and internet age reflected a “win-win dynamic,” benefiting both leaders and followers. Both types of firms together sought to create new markets, which required shifting out the demand curves for embryonic digital electronic devices in a context where perhaps the biggest challenge they faced was to persuade consumers to acquire and learn how to use operating systems and myriad applications. In other words, the challenge for the industry as a whole was: to fundamentally increase a consumer's willingness to buy a device by making it dramatically more useful and desirable than its potential substitutes or earlier vintages, rather than just making the existing version “nominally” cheaper.<sup>52</sup> The challenge was to convince people that these new gadgets were worth both the money and the effort to learn.

Consider the evolution of the personal computer in the mid-1990s. Early PCs were complex and required technical knowledge, limiting their appeal. However, a new generation of machines powered by faster processors came bundled with game-changing software like the “Windows 95” operating system and a “web browser.” This combination transformed the PC from a difficult-to-use office machine into an easy-to-use gateway to the new world of the internet, e-mail, and multimedia. This massive surge in value is what shifted the entire demand curve, convincing millions of households that a computer was no longer a niche hobbyist's toy but an essential appliance.

Ultimately, in fast moving and novel technology spaces, technology leaders could not do this on their own. It was only when supply chain partners, such as suppliers and assemblers, and even bitter rivals, worked together that they could hope to create an interoperable—and eventually interconnected—ecosystem of devices, applications, and complementary markets. Returning to the example of PCs, the industry's success hinged on a coordinated, positive-sum, effort across the different layers of its technology stack. It was often Intel's relentless push for more powerful and affordable microprocessors, famously marketed through its “Intel Inside” campaign, that set the pace. Each new generation of its chips enabled more demanding software, creating the

---

<sup>52</sup> However, this is not to say that “quality adjusted” prices won't be lower. To understand how this may be the case, first consider Arrow's (1962) explanation of how a drastic process innovation that displaces a previous product vintage and significantly reduces the costs of the new product below the price of the old product; even if the new product is sold at the monopoly price (half of the quantity is produced vis-à-vis a competitive market), this price lies below the competitively determined price of the old product. Now, take drastic product innovation, which we can extrapolate this logic to: while the nominal price for an item may be higher after a new product that boasts new bells and whistles that consumers greatly value comes to market and displaces the older vintage, consumers may still derive greater consumer surplus than they did before its introduction—and even if they pay the monopoly price (and thus deadweight losses obtain) and even if the previous product vintage was offered in a context of competitive, marginal cost pricing. Specifically, the new monopoly price is less than the increase in the value of the product after the innovation: the difference in consumers' highest willingness to pay for new product and consumers' highest willingness to pay for old product (see Galetovic and Haber 2017).

necessary foundation for Microsoft's revolutionary Windows 95 operating system and the first web browsers.<sup>53</sup> In turn, PC manufacturers like Compaq and Dell competed to bundle these powerful hardware and software into increasingly affordable machines for the mass market, creating a much larger pie for the entire industry.

Indeed, this tenet will be even clearer in the Chapter 6 of the book, when I explain how designers and implementers complemented the decentralized approach I outline here with joint efforts to set and implement technology standards promoting interoperability through their membership in SSOs.

### **Using Patent Data to Lend Empirical Corroboration**

While the prolific cross-licensing of Intel's foundational x86 architecture—a practice I discuss in Chapter 2 of the book—was initially driven by IBM's demand for a "second source" manufacturer to ensure a stable supply of microprocessors for its first PC, it helped create a powerful "innovation commons" that accelerated the development of this revolutionary device far beyond what a single firm could have achieved. This requirement compelled Intel to license its designs to competitors, most notably AMD, which had the long-term effect of cementing the x86 architecture as the dominant industry standard, enabling an entire ecosystem of PC manufacturers—including Compaq, Dell, and Gateway—to build compatible machines around a common architectural foundation. These firms could reliably source chips and design systems knowing the underlying instruction set was standardized and would be supported by a robust supplier base.

This ecosystem was defined by substantial but complementary roles for its leaders and followers. Leaders like Intel defrayed the costs of industrywide public goods beyond merely disclosing their ideas in patents—they sponsored conferences and trade shows, contributed to the scientific literature, invested in developer tools, and maintained backward compatibility across chip generations. They did this not out of altruism but because they were uniquely positioned to monetize these activities and capture returns from the expanding PC market. Followers benefited as well, but their role was different: they helped leaders create new markets for standardized components and products, and they exploited the fact that these technologies were already proven and could be widely adopted at lower cost and risk.

While the initial licenses allowed rivals like AMD to produce compatible chips, and while Intel remained the overwhelming innovation leader in this industry, this did not lead to a static, one-

---

<sup>53</sup> Windows 95 was revolutionary because it replaced the intimidating, text-based "MS-DOS" command line with a user-friendly "graphical user interface" (GUI) that was accessible to the average consumer. It introduced now-standard elements like the "Start Menu" for launching programs and a "taskbar" for easily switching between open windows, which made multitasking intuitive. Technically, its most significant breakthrough was "Plug and Play," a new standard that allowed the operating system to automatically detect and configure new hardware (like printers or modems), eliminating a major source of frustration for users of earlier systems. It also had built-in support for networking, which greatly simplified the process of connecting to the early internet.

way flow of technology. The most famous example of this occurred in the early 2000s, when AMD introduced a crucial innovation of its own: a 64-bit extension to Intel's 32-bit architecture. This new x86-64 standard was so successful that Intel was compelled to license the improved standard back from AMD to incorporate into its own processors, despite being the architecture's original source. Scenarios like these where the licensee's innovation is adopted by the original licensor are a prime example of how patent cross-licensing can create a competitive ecosystem where a shared technology platform is advanced by multiple players, benefiting the entire industry (see Menaldo 2021).

To grasp how I can systematically use patent data to put flesh on these technology creation and diffusion stories, we now move beyond licensing to the citations that link patents together. When a new patent application cites the "prior art" it builds upon, these references are called backward citations—like an academic paper's bibliography. The more telling metric for our purposes, however, is a patent's forward citations. Just as a paper's true impact is measured by how many future scholars cite it, a patent's influence is measured by the number of subsequent inventions that reference it. This makes tracking forward citations instrumental in identifying technological leaders and mapping the flow of ideas (Henderson et al. 1998; Brantle and Fallah 2007).

This makes forward citations instrumental in both identifying technological leaders and mapping the flow of ideas (Hall et al. 2001). A patent with thousands of forward citations represents a “foundational” piece of technology that served as a critical building block for a new field of innovation (ibid: 14-15). This is true whether a succeeding patent is making an incremental improvement on an antecedent patent's technology or attempting to "invent around" it (Brantle and Fallah 2007).

Importantly, there is a sincere, rather than simply strategic, reason why inventors cite previous patents in their patent applications and granted patent documents. Not only does an applicant face a legal duty to disclose its knowledge of the “prior art,” but a professional patent examiner, an expert in the relevant technology area, holds the patent seeker accountable by conducting a thorough search for any existing technology that might be like the claimed invention and actively helping the applicant identify the most relevant patents to cite. For these reasons, forward citations should be reliable indicators of an invention's importance and impact.

By extension, the process around discovering and acknowledging prior art and the subsequent citations linking patents together represents a powerful proxy for the actual technological links between inventions. This should therefore help capture the spillover of ideas and technology from one firm to another (Jaffe et al. 1993; Caballero and Jaffe 1993; Jaffe et al. 2000).

Returning to the example of recombinant DNA that I outlined above, a patent network analysis of the biotechnology industry during this era would show the Cohen-Boyer patent as a singular, massive hub. Forward citations would radiate out to an expanding universe of firms in both the pharmaceutical industry and its supply chain, demonstrating how one groundbreaking invention can seed an entire “innovation commons.”

*Preparing the Necessary Data*

To understand the dataset that follows and the multiple variables I construct from it and analyze, a brief primer on the US patent system during the relevant time frame, 1976-2006, is helpful.

First, note that the key features of patents are their statutory lengths of protection; how broadly one interprets the exclusive right over the idea; what is disclosed in the patent and when this occurs (immediately vs. delay); the size of inventive step required to earn protection (drastic vs. incremental); and what constitutes patentable subject matter (Scotchmer 2004). On these dimensions, the US system can be categorized as favoring broad scope and incremental innovation during this period (see *ibid*). Along these lines, it was—and continues to be—defined by three fundamental statutory requirements. Inventions must be novel (“new”), prohibiting patents on anything “known or used by others” or “described in a printed publication” before the effective filing date. They must also be non-obvious, barring patents on inventions that would have been obvious at the time to a person having “ordinary skill in the art.” Finally, they must be useful: awarded for processes, products, machines, combinations of materials, and improvements upon previous patents (US Congress 1952).<sup>54</sup>

In terms of what is patentable, the primary instrument is the “utility patent,” which protects the functional aspects of an invention. Patents are also granted to inventors for plants and modifications of plants, which are classified by the USPTO differently (see USPTO 2001). Finally, this also includes, as discussed in Chapter 3 of the book, software patents without affiliated hardware.

The nuts and bolts are straightforward. In the US, a patent application must describe the invention in detail and, in doing so, outline a series of explicit claims that comprise the invention. It must also include visual diagrams that explain the claims. A professional examiner must then screen, evaluate, and approve the application and decide whether it should be granted or rejected. If granted, utility patents last for 20 years. Design patents are granted for ornamental features of products and last 14 years (USPTO 2001).<sup>55</sup>

To map these types of patent citation networks for the technologies that were instrumental in driving the Third Industrial Revolution, I leverage a remarkable patent database published by the National Bureau of Economic Research (NBER) called the “Utility Patent and Patent Citation Data File” (see Hall et al. 2001). This dataset identifies and observes patents and patents’ features granted by the USPTO between 1976 and 2006. The variables in the dataset include granted patents; the number of claims made by each patent; patents’ backward and forward citations; and the cited patents’ and citing patents’ degree of “generality.” Critically, the 2006 end date coincided with the eve of Apple’s introduction of the iPhone, a fact I will capitalize on ahead.

---

<sup>54</sup> The patent application process offers multiple pathways, including provisional patents that provide preliminary protection before full review, and non-provisional applications that undergo comprehensive examination. For time-sensitive innovations, expedited review options exist to accelerate the approval process. This tiered approach balances the need for thorough evaluation with the practical demands of rapidly evolving industries (USPTO 2001).

<sup>55</sup> I am citing the 2001 edition of the MPEP (i.e., USPTO 2001) because it was the USPTO’s governing manual during the bulk of my dataset’s timeframe.

Out of the 3,279,509 patents in the NBER database, almost 93 percent are utility patents granted by the USPTO to inventions deemed genuine, novel, useful, and not obvious. Plant patents are less than 1% of all patents. Design patents are 6.5% of total patents. The analyses that follow focus exclusively on utility patents, which protect how an invention works or is used, as opposed to design patents, which protect an item's ornamental appearance.

While the individual “patent” is the NBER dataset’s original unit of analysis, individual inventors almost always assign their patents to a particular company when they obtain patents at the USPTO—allowing me to aggregate each of these by publicly listed/traded firms.<sup>56</sup> To do so, I leverage the COMPUSTAT database published by Standard and Poor’s.<sup>57</sup> While it covers thousands of publicly traded companies worldwide, with information dating as far back as 1950, I relegate attention to all publicly traded firms that existed between 2000 and 2006—a period representing the mature phase of the dot-com era immediately preceding the launch of the iPhone in 2007—and match all utility patents awarded between 1976 and 2006 to those firms.<sup>58</sup> There are 63,266 patents with multiple assignees in the NBER dataset; about 2% of them. To address this issue, I matched these patents to each of the firms assigned to them as documented by the USPTO.<sup>59</sup>

All told, these efforts yielded a cross-sectional dataset of 25,480 firms holding 3,023,482 patents. In total, these patents received 28,292,441 forward citations, representing over 28 million unique links (citing-cited patent pairs) between inventions.

It is important to note that older, incumbent technological fields tend to cite more and are cited relatively less, while emerging fields at the technological frontier are cited more, but cite less outside of their area (Hall et al. 2001: 16). Indeed, consider that 50% of citations are made to patents at least 10 years older than the citing patent, 25% to patents 20 years older or more, and 5% of citations refer to patents that are at least 50 years older than the citing one (Hall et al. 2001: 17). Therefore, I weight the patent data by the amount of time the firm has been around, establishing a “floor” on the number of years a firm has been in existence by starting the count in

---

<sup>56</sup> Not all patents are assigned to public firms; some are assigned to individuals, universities, or private firms. I do not include any of these in the dataset.

<sup>57</sup> Throughout the rest of this chapter, I classify firms into their respective industries (e.g., “semiconductor,” “telecommunications,” “computer manufacturing”) based on their North American Industry Classification System (NAICS) code. For this analysis, each firm was assigned to its primary industry using its 6-digit NAICS code as reported in the COMPUSTAT database. This system provides the most specific, standardized definition for a firm's business activities (e.g., ‘334413’ for Semiconductor and Related Device Manufacturing).

<sup>58</sup> This means they were born before or after 2000, but not after 2006.

<sup>59</sup> I address the transfer of patents from one firm to another resulting from corporate mergers, acquisitions, and spinoffs. For example, consider the case of Lucent: it is observed as “owning” patents initially assigned to AG Communication Systems Corp., VMX Inc., Otel Communication Corp., Lucent Tech Remote Access Business Unit, Lucent Tech, Excel Switching Comp., Yurie Systems Inc., Ortel Inc., and International Network Services.

1950 for all firms older than 1950 that helps control for the fact that older firms have simply had more time to accumulate patents.

### **Preliminary Corroboration of Superstar Economics**

After normalizing the data in that manner, the mean number of (weighted) patents is 4.5, the median is .17, and the standard deviation is 28.10 patents.<sup>60</sup> These descriptive statistics about patenting are a preliminary signal that there is a reason to believe that a small handful of superstar firms may also represent technology leaders that set the pace of innovation. The vast gap between the average number of patents per firm and the median number is telling: a handful of patent-rich "superstar" firms pull the average sky-high, while the typical firm holds only a tiny number of patents. This extreme inequality potentially points to a "winner-take-most" dynamic where a few "innovation leaders" are responsible for a vastly disproportionate share of technological assets.

This "winner-take-most" pattern is even more pronounced when it comes to patents' forward citations. The vast 150-fold gap between the average (1,139.5) and the median (7) immediately signals an immense "innovation inequality."<sup>61</sup> This is confirmed by the distribution's huge standard deviation (11,147.59 forward citations per firm), its extreme positive skewness (28.47), and a kurtosis of over 1,100, which points to the predictable presence of massive outliers, or "fat tails." The most striking illustration of this dynamic is the fact that the top 20% of firms account for a staggering 98% of all forward citations.<sup>62</sup> In short, the data reveals that a few firms completely dominate in terms of technological influence, while most have very few citations to their patents or none whatsoever.

This is all well and good, but a larger question about this strategy for identifying superstar firms remains: does a firm's status as an innovation hub, as measured by patent citations, help predict its real-world economic performance? To explore this issue, I now turn to looking at whether there is a systematic relationship between Total Factor Productivity (TFP) and firms' patent forward citations.

Constructed by Bahar (2018), TFP is calculated as the residual from a Cobb-Douglass production function.<sup>63</sup> He estimates it using firm-level data on operating revenue, tangible fixed assets,

---

<sup>60</sup> In terms of the unweighted version, the average number of patents per firm is 124.4, the median is 3, and the standard deviation is 1,108.2.

<sup>61</sup> In terms of backward citations, the average firm cites 1,294.69 patents across its patent portfolio, with a median of 21 citations.

<sup>62</sup> Focusing attention on the average forward citations per patent, per firm (instead of total citations per firm across all its patents), does not really change the story. The mean is 5.39 citations per patent, and the median is 2. The standard deviation, however, remains rather high at 9.9, with a skewness statistic of 6.39, and a kurtosis statistic of 80.41.

<sup>63</sup> TFP is the residual from a Cobb-Douglass production function with constant returns to scale ( $Y = A \cdot K^\alpha \cdot L^\beta \cdot M^\gamma$ ), which is estimated via OLS. After calculating the elasticities, the TFP

employees, and cost of materials.<sup>64</sup> The TFP data is measured as the firm-level average between 2007 and 2015 for all firms in the COMPUSTAT dataset; any independent variable is measured between 1976 and 2006.

Before conducting this analysis in particular, I made some important adjustments to the dataset. Firms that exited it between 2001 and 2007 do not contribute to any of the regressions that follow.<sup>65</sup> Excluding them reduces the number of observations from 5,926 to 5,060.<sup>66</sup> I also drop economic sectors with only one firm since we are interested in focusing on data about sectoral citation leaders and followers throughout the analysis that follows and, by definition, a firm cannot either lead itself or follow itself. This therefore rules out including sectors characterized by singletons. This restriction accounts for losing an additional 1,356 observations.

The descriptive statistics alone are suggestive that we are on the right track. They reveal a distribution that is not normal, but is instead highly skewed, consistent with a Pareto-like "fat tail." While the median firm produces .21 outputs per unit of input, the mean is dragged to .36. This massive skew is driven by a handful of hyper-productive outliers, confirmed by a maximum value of 186.64 and a staggering standard deviation of 274. In short, this "fat tail" is the superstar effect in numerical form, providing a clear picture of the productivity divide that defines this era.

Moreover, it appears that "superstar" firms are, in fact, the very "technological leaders" that revealed themselves in the preceding analysis. First, consider that patent forward citation leaders' median TFP is .32 (outputs per unit of input), far exceeding the .20 level of the "citation followers": the firms that cite them. Second, the results of a bivariate OLS regression, estimated with standard errors clustered by economic sector, are even more stark: a firm's status as a patent forward citation leader is associated with a 100 percent higher level of TFP than its followers (p

---

residuals are estimated in the following manner:  $\ln(TFP) = \ln(Y) - \hat{\alpha} \ln(K) - \hat{\beta} \ln(L) - \hat{\gamma} \ln(M)$ .

<sup>64</sup> Following Gal (2013), missing values for material costs are imputed by computing the difference between operating revenue and value added. The firm-level data is then weighted using the OECD's Structural and Demographic Business Statistics (SBDS) database to correct for sampling bias in firm size and industry employment. Furthermore, all monetary variables are converted to constant 2010 U.S. dollars using industry-specific deflators, including the PPI for revenues and investment goods, a specialized input-output weighted PPI for materials, and the ECI for employee costs.

<sup>65</sup> I do this because I want to neutralize the potential that speculative financial bubbles are driving our results. Consider the NASDAQ stock exchange, where the lion's share of the biggest and most influential American high-tech public firms' shares traded during our observation window: many of those corporations witnessed an epic boom and implosion between 1998 and 2000 known as the .com bubble and bust. Many of the firms that gained notoriety during this period are infamous: their meteoric rise was followed by spectacular bankruptcies. They had untenable business models and turned out to be flashes in the pan rather than technologically innovative firms.

<sup>66</sup> This does not mean that the obverse is true, however: any firm that enters the dataset between 2000 and 2015, and has TFP data between 2006 and 2015, makes it into the regressions.

< .001), suggesting that technological influence and economic productivity are two sides of the same coin.<sup>67</sup>

This powerful bivariate correlation, however, could be spurious: A firm's status as a citation leader might simply be a byproduct of other sector-wide characteristics—such as a high density of patents, a proliferation of general-purpose technologies, or high levels of R&D spillovers—that independently boost the productivity of all firms in that industry. To isolate the true effect of technological leadership, I therefore introduce a battery of control variables to account for these potential confounders.<sup>68</sup> Even in a fully specified multivariate model, a firm's status as a citation leader remains a powerful predictor, correlating with a 67% higher level of total factor productivity ( $p < .001$ ). More decisively, introducing sector fixed-effects—which expunges all sector-level omitted variables and focuses exclusively on variation within each industry—yields a materially identical result.

The above analysis provides powerful economy-wide evidence that innovation leaders are also productivity superstars. The link between patent citation dominance and superior TFP appears quite robust.

Yet, a central argument of this section of the appendix is that, in the run up to the Fourth Industrial Revolution, innovation was not necessarily uniformly distributed across the economy but was instead driven by and concentrated in the high-tech sector. These firms were purportedly characterized by a pronounced "rich-get-richer" phenomenon: a small number of pacesetters became massively influential hubs, while the vast majority remained on the periphery with very few connections. Was it really the case that a few firms at the cutting edge of the era's frontier technology industries assumed the role of innovation hegemony and took a lead role in diffusing technology across their industries? The next logical step is to formally test the contention that a winner takes most dynamic characterized high-tech firms' innovation commons in an exceptional manner.

### **Cross-sectional Power Law Assessment**

---

<sup>67</sup> Standard errors are clustered by specific high-tech sectors (e.g., Semiconductors, Software Publishing, Computer Manufacturing) rather than a standard industrial classification. This clustering strategy is empirically justified by a Multiple Regression Quadratic Assignment Procedure (QAP) analysis, which confirms that the dyadic distance in TFP levels between any two firms is significantly and negatively correlated with co-membership in each of these high-tech categories ( $p < .05$  for all tested sectors). This finding suggests that firms within these purpose-built clusters share a distinct productivity profile, justifying their use for standard error correction. The main results are also robust to clustering at the 3-digit NAICS level, however.

<sup>68</sup> To isolate the effect of citation leadership from other industry-wide characteristics, the multivariate model controls for the following sector-level variables, measured for the 1976 to 2005 period: (1) the average number of patents; (2) the median number of claims per patent; (3) the median number of forward citations per patent; (4) the median number of backward citations per patent; (5) the generality of the citation leader's forward citations; (6) the generality of the citation leader's backward citations; (7) the generality of the citation followers' forward citations; and (8) the sectoral forward citation Gini coefficient.

As a first step in identifying and mapping this phenomenon, I expect the patent distribution for high-tech firms to be highly unequal.<sup>69</sup> Rather than clustering around a typical value, I anticipate a pattern best described by a Power Law distribution. Most simply, this describes a "rich-get-richer" phenomenon. Its most famous statistical feature is a characteristic "fat tail," which means that massive outliers are not just possible, but are an expected and predictable feature of the distribution. This results in a dynamic where a handful of technological leaders completely dominate the "innovation landscape."

Technically-speaking, I need to evaluate whether the distribution of total patents weighted by firm age for high tech firms follows a Power Law in terms of a constant multiplied by a power of the variable in question, which implies that firms with relatively higher shares of patents tend to have orders of magnitude greater patents than firms with fewer ones.<sup>70</sup> In other words, a relative change in one quantity should map onto a proportional relative change in the other quantity. In folk language, such a distribution subsumes the idea that the top 20 percent of the firms in the patent data distribution account for 80 percent (or more) of the total patents. The simplest case is a  $\frac{1}{f}$  function: if the values are sorted from most common to least common, the second most common frequency occurs half as often as the first; the third most common frequency occurs 1/3 as often as the first; and the  $n^{th}$  most common frequency occurs  $1/n$  as often as the first. I focus particular attention on ascertaining whether the right tail of the data's distribution obeys such a pattern, speculating that this is more likely than a power law corresponding to the variable's entire range of values.

More formally, a quantity  $x$  is governed by a power law if it is drawn from a probability distribution  $p(x) \propto x^{-\alpha}$ , with  $\propto$  denoting proportionality. This can also be written as  $p(x) = Cx^{-\alpha}$ , where  $\alpha > 0$  is a constant, shape parameter,  $C$  is a constant, scale parameter, and  $x \geq C > 0$  is a random variable. The right tail of the distribution is heavier as  $\alpha$  gets smaller. This distribution therefore has cumulative distribution function (c.d.f.) equal to  $F(x) = 1 - \left(\frac{C}{x}\right)^{-\alpha}$ ; its probability density function (*p. d. f.*) is  $f(x) = \alpha(C)^{\alpha}/x^{\alpha+1}$ .

Several things follow from these facts that I can exploit to diagnose and describe a power law. If we take logs of both the data's rank order (quantiles) and its frequency, we obtain  $\ln p(x) = -\alpha \ln x + C$ . This means I should observe a straight line on a doubly logarithmic plot, with  $-\alpha$  as the slope of that line. Moreover, I can calculate a gini coefficient to gauge the level of

---

<sup>69</sup> To qualify as "high-tech," a firm must be classified by the NAICS as either operating in wireless telecommunications, semiconductors, computer manufacturing, software publishing, data processing and hosting, computer system design services (Network Systems Integration Design Services), other computer services, and custom computer design services (Computer Program or Software Development).

<sup>70</sup> This is often referred to a Pareto 1 distribution or a Zipf distribution. For a review of how pervasive power law distributions are across several processes studied by economists, demographers, biologist, physicists, engineers, and other researchers, as well as how to identify them statistically, see Clauset et al. (2009).

inequality implied by calculating the following equation,  $1 - (2\alpha - 1)$ , which is derived from the Lorenz curve ordinates at each  $s = F(x_s)$  given by  $L_s = 1 - (1 - s)^{1-1/\alpha}$ , with 0 denoting a uniform distribution and 1 denoting that one firm has all of the patents. However, what is more likely is that the power law applies to values greater than some minimum,  $x_{min}$ , so that it is most likely the tail of the distribution that obeys a power law. By extension, only a portion of the slope of the line above a designated  $C$  that  $F$  is greater than zero (after taking logs) will be straight.

I therefore proceed in a few steps. First, I graph a double log rank frequency plot of the data and visually inspect it to establish whether it depicts a straight line or not. Second, I assume that the entire data's distribution obeys a power law and estimate the  $\alpha$  parameter via maximum likelihood, as well as the Gini coefficient. Third, I search for the distribution's  $x_{min}$  by varying  $C$  while calculating a goodness of fit statistic between the data that lies above the stated threshold and the power law distribution to formally assess the null hypothesis that the data does not follow a power law. I thus try to narrow down *where* in the distribution the data obeys a power law.

As a first step, to evaluate whether the patent data for high tech sectors follows a power law, I graph the log of the number of patents weighted by firm age and the log of the rank of firms (the firm that patents the most is assigned the highest rank). As outlined above, I expect a linear relationship. Figure S2.1, which depicts said graph, largely fits the expected pattern; it also appears that the line is straighter as it gets closer to the right tail of the distribution, which I expect given that the power law probably kicks in at relatively higher data values.

As a second step, I follow Jenkins and Van Kerm (2007) and derive the likelihood function for a sample of observations on patents for high tech firms specified as the product of the densities for each observation. This yields an  $\alpha$  coefficient of 1.04 (p-value < .001), which implies, as expected, a very steep power law relationship. I also calculate the Gini coefficient using this approach, which is .93.<sup>71</sup> Moreover, I note that the top twenty percent of high-tech firms have 95.4 percent of the total patents held by firms in these sectors.

Based on Figure S2.1, I surmise that power law dynamics seem to really kick in at relatively higher values of the number of patents. Therefore, we exploit the technique pioneered by Urzúa (2020) to both test the goodness of fit of the power law to the data and locate the appropriate value for  $x_{min}$ . I test for different thresholds of  $C$  at the patent data's 5<sup>th</sup>, 25<sup>th</sup>, 50<sup>th</sup>, 75<sup>th</sup>, 90<sup>th</sup>, 95<sup>th</sup>, and 99<sup>th</sup> percentiles and fail to reject the hypothesis that the data follows a power law once we reach the median of the data's distribution, which is 1.5 patents (p-value = .73). I note that the results are robust to measuring patents as a firm's total patents instead of a firm's patents

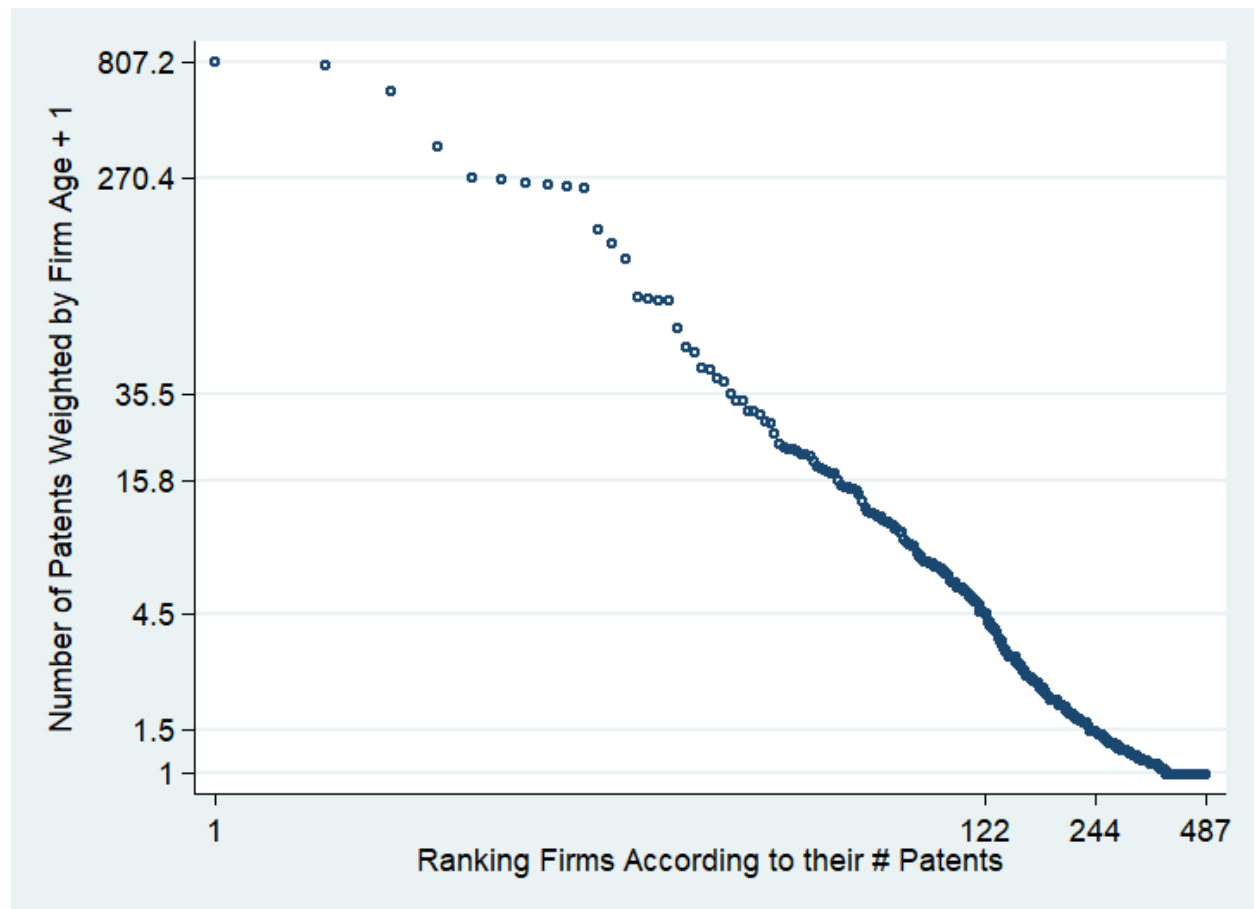
---

<sup>71</sup> The Gini coefficient measures statistical dispersion, with 0 representing perfect equality and 1 representing maximal inequality. The Gini coefficient for a Pareto (Power Law) distribution can be calculated directly from the scaling parameter alpha ( $\alpha$ ). The formula is  $1 / (2\alpha - 1)$ . Given the estimated  $\alpha$  of 1.04, the calculation is:  $1 / (2 \times 1.04 - 1) = 1 / (2.08 - 1) = 1 / 1.08 \approx 0.926$ , which rounds to the 0.93 reported.

weighted by its age. The only meaningful difference is that for the former the value of  $x_{min}$  now lies at the data's 75<sup>th</sup> percentile, 31 (p-value = .87).

As expected, things are different if I focus attention on the non-high-tech sample. Returning to the number of patents weighted by firm age, I find that at no values of  $C$ , including the variable's lowest observed value, do that data seem to follow a power law. Specifically, I reject the hypothesis that the data follows a power law for the following values of the data's distribution: the minimum, 1 (p-value < .001); the median, 1.15 (p-value < .001); the 75<sup>th</sup> percentile, 2.21 (p-value = .01); the 90<sup>th</sup> percentile, 6.09 (p-value < .001); the 95<sup>th</sup> percentile, 13.4 (p-value < .001); and the 99<sup>th</sup> percentile, 60.07 (p-value < .001).

**Figure S2.1: Power law fitted to patent count data (Rank-Size Plot)**



*Notes: both the y and x axes are logged. I add 1 to the number of patents weighted by firms because several firms are observed as having zero patents. On the y-axis, I label the lowest value (0 + 1), median, 75<sup>th</sup> percentile, 90<sup>th</sup> percentile, 95<sup>th</sup> percentile, 99<sup>th</sup> percentile, and highest value. On the x-axis, I label the lowest value, 25<sup>th</sup> percentile, median, and highest value.*

Taken together, these formal tests provide rigorous, statistical validation for the claim that during the Third Industrial Revolution the patent distribution for high-tech firms follows a Power Law, revealing a "winner-take-most" dynamic where a handful of innovation hegemony possessed

orders of magnitude more IP than their peers. Crucially, this is not a universal feature of the American economy during this period. The patent distribution for non-high-tech firms fails to conform to a power law at any threshold, indicating a fundamentally different and more equitable structure of innovation. This stark divergence provides the empirical justification for this section of the appendix's deep dive into some of the specific mechanisms of the high-tech innovation commons, to which we now turn.

### *The Patent Citation Gap in High-tech Sectors*

The same winner-take-most dynamic should be even more pronounced when examining patent forward citations. Among high-tech firms, I expect to find a vast citation gap between the handful of technological leaders and the broader periphery of followers. Because followers are more likely to adopt and build upon the leaders' standardized technologies, they will disproportionately cite the patents of the most influential firms. This dynamic implies that, unlike their non-high-tech counterparts, high-tech sectors will exhibit a highly unequal distribution of forward citations, a pattern again best described by a Power Law: Instead of citations clustering around a typical value, I expect to find a right-skewed distribution with a characteristic "fat tail," indicating a massive ratio between the most-cited firms and the least-cited.

I therefore apply the same rigorous protocol used in the previous section to see whether I can identify a Power Law in the forward citation data. First, I graph a log-log plot of the forward citation data to visually check for the characteristic straight-line pattern. Next, I estimate the overall inequality by calculating the Gini coefficient via maximum likelihood. Finally, I formally test the Power Law hypothesis by searching for the optimal  $x_{\min}$  threshold where the tail of the distribution best fits the theoretical model.

The results from these tests are decisive. The log-log plot in Figure S2.2 provides the initial visual evidence, revealing the expected linear relationship, particularly in the right tail of the distribution for firms with higher citation counts. The overall inequality is staggering: a maximum likelihood estimation yields a Gini coefficient of .96, a figure consistent with the finding that the top 20% of high-tech firms command 98.8% of all forward citations. Formal testing confirms that this extreme concentration is driven by a Power Law dynamic. The data conforms to a Pareto distribution above the 75<sup>th</sup> percentile (a threshold of 218 citations), where I fail to reject the null hypothesis (p-value = .40). This finding is robust to measurement: when I use average citations per patent, the results are identical.

As expected, a starkly different pattern emerges for the non-high-tech sample. The Power Law hypothesis is overwhelmingly rejected across the distribution; instead, the data for these firms appears to follow a log-normal distribution, a fact highlighted by a hangroot plot that I do not report here to conserve space. To quantify the magnitude of this structural difference, I next compare the level of inequality on a sector-wide basis using the Gini coefficient for forward citations.<sup>72</sup> A simple bivariate regression across 669 industries reveals a dramatic gap. The average Gini coefficient for high-tech sectors is .725, nearly double the .377 average for non-

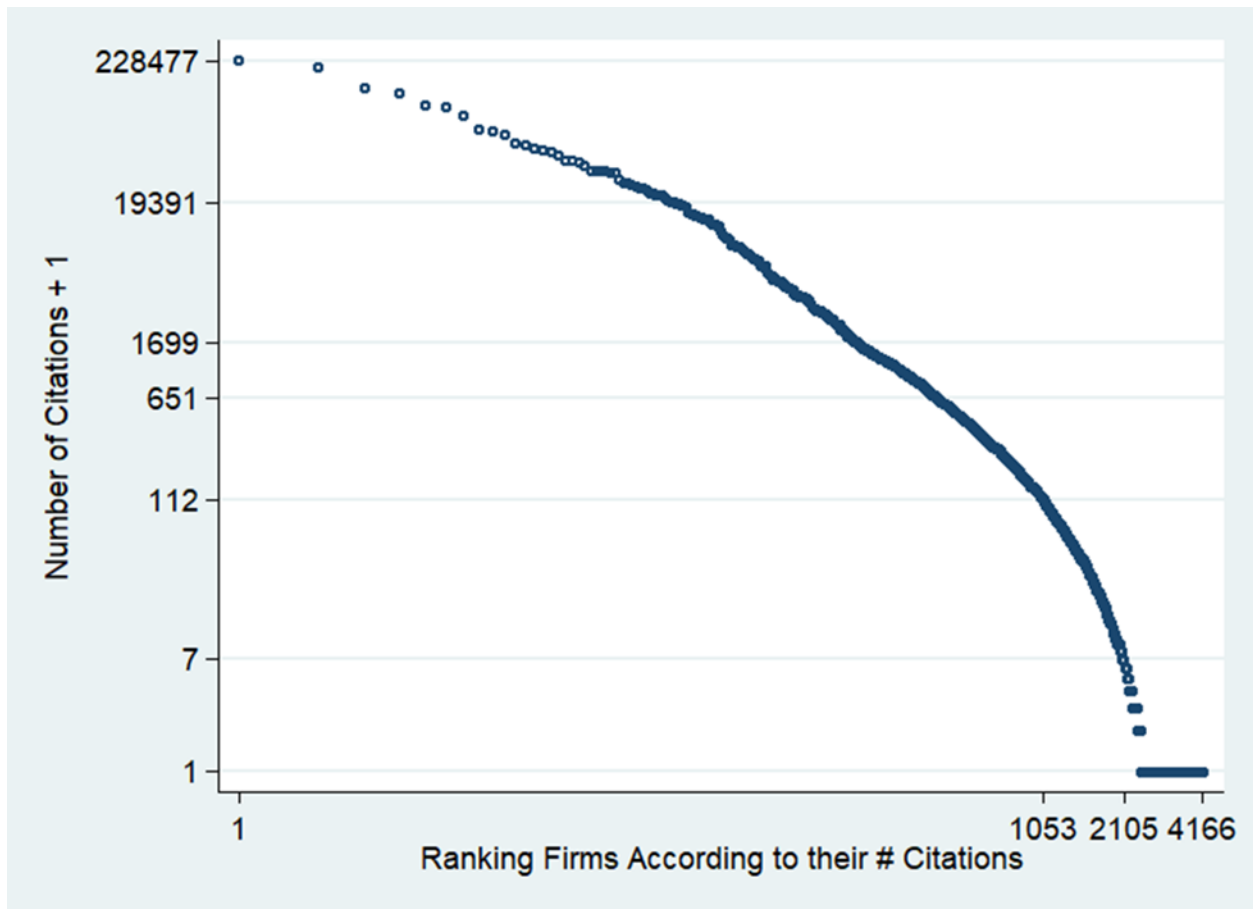
---

<sup>72</sup> That is, each firm in each sector is a data point that contributes to the calculation of this sector wide measure of inequality: economic sectors are the unit of analysis.

high-tech sectors (a difference of means with  $p < .001$ ). This finding provides powerful, cross-sectional evidence that the "winner-take-most" innovation dynamic is a unique and defining feature of the high-tech ecosystem.

The results from these tests are unequivocal, strongly corroborating the idea that, between 1976 and 2006, the innovation economy embodied by the high-tech sector obeyed Superstar Firm dynamics. Its patent citation patterns conform to a Power Law, creating a "winner-take-most" dynamic where a few influential firms attract orders of magnitude more citations than their peers. Crucially, this is not a universal feature of that economy, but a unique characteristic of the high-tech ecosystem. This empirical confirmation provides the justification for the next phase of our analysis: a deep dive into the specific network mechanisms that drove this exceptional concentration of technological influence.

**Figure S2.2: Power law fitted to patent forward citation data (Rank-Size Plot), High Tech**



*Notes: both the y and x axes are logged. I add 1 to the number of patent citations weighted by firms because several firms are observed as having zero patents. On the y-axis, I list the lowest value (0+1), 25th percentile, median, 75th percentile, 90th percentile, 95th percentile, 99th percentile, and highest value. On the x-axis, I list the 25th percentile, median, and highest value.*

## Patent Citation Network Datasets

I now seek to more clearly identify the distinct “innovation leaders” that cut their teeth developing and diffusing inventions central to the transmission of technology vital to their industry across high tech sectors. Broadly, I am looking to establish several patterns that satisfy the idea that distinct network leaders are the wellspring of their industry’s foundational technologies and act as the “hubs” in a technological exchange in which peripheral firms serve as “spokes” by acquiring, honing, and helping spread the formers’ critical innovations. To address these claims systematically, I constructed two network datasets that observe citation ties between high tech firms existing between 2000 and 2006.<sup>73</sup>

The first network dataset I built allows me to evaluate the transfer of information and technology *between* firms in the semiconductor industry during the Third Industrial Revolution, again spanning from 1976 and 2006. I call this the “horizontal network.” The first half is dedicated to mining the horizontal connections between semiconductor firms. This allows me to both operationalize and test hypotheses about the transfer of information and technology between firms in roughly the same segment of the supply chain; indeed, some inventions may spread between rivals competing for the same market share.

The second dataset is a “vertical network” designed to map the supply chain connections between the upstream semiconductor firms and the downstream companies that used their chips in the pre-2007 era. Unlike the first network, this one excludes horizontal ties between rivals and focuses exclusively on the flow of technology up and down the digital device “stack.” The downstream firms include the hardware manufacturers (producing PCs and early mobile devices like the Palm Pilot), wireless telecommunications companies, and software publishers that together created the pre-2007 digital ecosystem.

### *Network Analysis Basics*

At its core, a network consists of two simple things: “nodes” and “edges”. The nodes (or vertices) are the individual items being studied. In a professional social network like LinkedIn, for example, the nodes are the individual user profiles. The edges (or ties) are the connections that link these nodes together. On LinkedIn, this could be a “1st-degree connection,” the act of following an influencer, or membership in the same industry group.

To analyze these connections, researchers represent the network as a dataset, most commonly in the form of an “adjacency matrix”. You can think of this as a large spreadsheet where every

---

<sup>73</sup> I am not the first researcher to convert patent data into network datasets to explore technological leadership and technological diffusion. Jaffe and Trajtenberg (2002) use patent citation data to track the flow of knowledge between universities, corporations, and different countries, famously showing that these “knowledge spillovers” are often geographically concentrated. Brantle and Fallah (2007) seek to show that highly connected nodes, or hubs, are the most cited firms across all industries, not just semiconductors or the larger electronic device industry supply chain. Meanwhile, Kim and Song (2012) use patent lawsuit information to identify technological leaders in the smartphone industry, which differs from the citation-based approach I undertake here.

node—every person on LinkedIn—is listed as both a row and a column. A cell in the spreadsheet receives a "1" if a tie exists from the row person to the column person and a "0" if it does not. This simple grid of 0s and 1s provides a complete map of the network, which can then be used to perform statistical analysis at two levels.

Keeping with our running example, a “node-level” analysis could focus on a single influential person on LinkedIn to measure their importance. On the one hand, we could calculate their “degree centrality” by simply counting their connections. On the other hand, we could identify their “betweenness centrality” to see if they act as a crucial “bridge” connecting different professional groups.

A “network-wide” analysis, on the other hand, looks at the structure of the entire group. For example, it might show that senior biotech researchers in Boston form a “high-density”, tight-knit community where most individuals are already connected to one another. In contrast, the same analysis might show that retail store managers across the U.S. form a “low-density”, fragmented network where most managers are strangers with few professional connections between them.

Translating this into more formal language, a network with  $n$  nodes is represented by an  $n \times n$  adjacency matrix,  $Y$ , where  $Y_{ij} \neq 0$  if a relationship exists between nodes  $i$  and  $j$ . The nature of the relationship determines the matrix's properties. In an undirected network, where a relationship is always mutual, the matrix is symmetric ( $Y_{ij} = Y_{ji}$ ). Conversely, in a directed network, where a relationship can be one-way, the matrix is asymmetric, as a tie from node  $i$  to node  $j$  does not imply a tie from  $j$  back to  $i$ .<sup>74</sup> The fundamental unit of analysis in these networks is the dyad, or a pair of actors ( $i, j$ ), whose complete relationship is described by the configuration of its tie variables ( $y_{ij}, y_{ji}$ ).<sup>75</sup>

The relationship within each dyad can express one of three states: it can be mutual, where a tie flows in both directions between the two nodes; asymmetric, where a tie flows in only one direction; or null, where no tie exists at all.

The next level of social structure is the triad, which is a set of three nodes and the ties among them. A triad is the smallest unit in which we can analyze more complex social dynamics like balance, hierarchy, and transitivity (the principle that "a friend of my friend is also my friend"). The relationships within even this small group can be surprisingly complex; in a directed network there are 16 possible configurations of ties that can exist within a triad.

### *Expectations about our Particular Networks*

---

<sup>74</sup> An example is a social network like X (formerly Twitter), where you can "follow" an influencer (a tie from you to them), but they do not automatically have to follow you back.

<sup>75</sup> In terms of graph theory and notation, a network can be described as the graph  $G = (V, E)$ , where  $V = \{1, 2, 3, \dots, n\}$  is a set of vertices (nodes) and  $E \subseteq \{\langle i, j \rangle \mid i, j \in V\}$  is a set of edges (ties). Edges are pairs of vertices; for example,  $\{\langle 1, 5 \rangle, \langle 2, 5 \rangle\}$ .

Applying this framework to the patent citation networks, we should expect to observe a scale-free structure with clear network leaders where technology is widely distributed. This means uncovering a relatively large patent citation gap between technological leaders, which should be cited extensively, and peripheral firms, which should be cited minimally or not at all. The structural characteristics of such a network should manifest in several specific ways.

First, we expect relatively short average path lengths between nodes—that is, the mean number of steps along the shortest citation paths connecting all possible pairs of firms should be small. In practical terms, this means that knowledge can flow from one firm to another through only a few intermediate patent citations, enabling rapid technology diffusion across the industry. This efficiency in knowledge transmission is a hallmark of scale-free networks, where hub firms create shortcuts that dramatically reduce the degrees of separation between otherwise distant peripheral firms.

Second, the network should exhibit high clustering coefficients combined with low transitivity, a seemingly paradoxical but revealing combination. The clustering coefficient measures the probability that two firms citing the same third firm also cite each other, with values ranging from 0 to 1. High clustering indicates dense local neighborhoods where firms tend to cite common sources. However, transitivity—the overall probability that if firm *A* cites firm *B*, and *B* cites *C*, then *A* also cites *C*—should remain relatively low. This pattern suggests that while firms may cluster around common technological sources (creating high local clustering), they don't necessarily form closed triangles of mutual citation. Instead, knowledge flows through central hubs rather than circulating within tightly knit subgroups, preventing technological balkanization.

Third, the node degree distribution should exhibit extreme inequality, potentially following a power law dynamic. In a citation network, a node's in-degree is the number of times it is cited by other firms, while out-degree is the number of other firms it cites. What this means in our case is that most nodes should exhibit few connections and a small number of "hub" nodes should have many connections. This is because a network's degree distribution describes the probability distribution of node degrees (number of connections) across the entire network.

Indeed, I can search for the same Power Law distribution identified in the cross-sectional data here: When applied to a network, the probability  $P(k)$  that a randomly selected node has  $k$  connections obeys the relationship  $P(k) \propto k^{-\beta}$ , where  $\beta$  represents the scaling exponent. Unlike a normal distribution where connections would cluster around an average, this polynomial relationship dictates that a few 'hub' firms receive vastly more connections than others. This structure creates the characteristic "fat tail" we observed earlier, where a handful of technological leaders dominate the citation landscape while most firms occupy the periphery with minimal connections.

Fourth, network leaders should exhibit high centrality across multiple dimensions. Centrality encompasses several complementary measures of a node's importance: degree centrality counts direct connections, betweenness centrality measures how frequently a node lies on the shortest paths between other nodes, eigenvector centrality weights connections by the importance of

connected nodes, and closeness centrality calculates average distance to all other nodes. Technological leaders should score highly across these measures, confirming they not only maintain numerous direct connections to peripheral firms but also serve as critical intermediaries mediating knowledge flows between otherwise disconnected parts of the network. These firms function simultaneously as knowledge sources, transmission hubs, and bridges connecting diverse technological communities.

Fifth, the network should exhibit a relative absence of cliques—subsets of nodes where every member cites every other member, forming fully connected subgraphs. Rather than fragmenting into smaller, isolated neighborhoods or subnetworks, we should observe a single, dense, tightly connected village. In short, the horizontal network reveals strong interconnectedness that is uniform—tending towards universality—and an absence of cliques, which would indicate that technological knowledge doesn't get trapped in isolated communities but instead flows through the broader network via hub firms. This structure ensures that innovations can diffuse widely rather than remaining confined to technological silos.

Finally, firms that acquire technology from network leaders should seek access to foundational types of technology from them. Therefore, when peripheral firms' patents cite the latter's patents they should be citing relatively more “general” technology than what they possess in their own portfolios: a high level of generality for backward citations means that a patent cites other patents that have an impact across a wide range of fields. Conversely when firms cite peripheral firms' patents, they should not necessarily be seeking to acquire “general” technology from them but, rather, bespoke technological applications.

### **Horizontal Network: Semiconductor Firms**

The first dataset is a “horizontal network” that observes patent citations among 127 semiconductor firms.<sup>76</sup> It is a directed network, meaning a tie from Firm *A* to Firm *B* does not imply a tie from *B* to *A*. For this analysis, a tie is recorded as a binary relationship: a link exists if one firm cites any of the other firm's patents at least once, irrespective of the total number of citations. This allows for several possible relationships between any two firms, including a null tie (no citations in either direction), an asymmetric tie (one cites the other, but not vice versa), or a reciprocal tie (both cite each other), as well as self-citations.

#### *Visual Topology and Basic Network Features*

Figure S2.3, a global visualization of the horizontal network, reveals that the 127 semiconductor companies in the dataset are connected through 2,347 citation linkages in a complex manner. The network's most striking feature is its heterogeneous connectivity pattern, a core-periphery structure with three types of nodes. At the center there are dominant hubs: large, more lightly


---

<sup>76</sup> I excluded some “edge” cases from this dataset for principled reasons. First, I omitted subsidiaries of larger foreign corporations (e.g., Samsung Semiconductor Inc.). Second, I excluded firms with complex acquisition histories that predate the analysis period (e.g., NXP Semiconductors N.V., formerly VLSI). Finally, I omitted firms for which I could not find patent records in the NBER patent dataset (e.g., Aeroflex Holding Corp.).

shaded circles that represent highly cited firms including Texas Instruments, cited by 99 firms, and Intel, cited by 89. The middle tier shows up as medium-sized, moderately shaded nodes forming bridges between the core and periphery, which include firms such as Cirrus Logic (cited by 57 firms and ranked 10th) and Maxim Integrated Products (cited by 61 firms and ranked 8th). Finally, small, dark-colored nodes scatter along the edges—25 firms receive no citations at all, while 47 receive five or fewer citations. These minimally cited firms appear as the smallest nodes on the periphery, often connected to the network through only one or two edges leading to major hubs.

The network's visual structure—with its clear core-periphery organization, dominant hubs, and absence of isolated clusters—suggests specific patterns of knowledge diffusion. The thick edges converging on Texas Instruments, Intel, and Qualcomm indicate these firms generate foundational technologies that others build upon. The peripheral firms' thin, often singular connections to these hubs suggest they are technology adopters rather than generators, accessing standardized or general-purpose technologies through their citations. In short, the network structure implies that innovation flows from a concentrated core of technological leaders to a broader periphery of specialized firms, with middle-tier companies serving as bridges that help diffuse knowledge throughout the ecosystem.

### **Figure S2.3 Visualizing the Semiconductor Network**



*Notes: This network visualization displays the citation relationships among 127 semiconductor firms through 2,347 directed ties (links) based on patent citations from 1976 to 2006. Node size is proportional to the number of times a firm is cited by other semiconductor firms in the network (in-degree centrality). The largest nodes represent the most influential technology hubs: Texas Instruments (cited by 99 firms), Intel (cited by 89 firms), and Qualcomm (cited by 39 firms). The network exhibits a core-periphery structure with a density of 14.7%, meaning that only about 15% of all possible citation ties exist. Twenty-five firms receive no citations, while 47 firms receive five or fewer citations, appearing as the smallest nodes on the periphery. A tie is recorded as a binary relationship: a link exists if one firm cites any of another firm's patents at least once during the analysis period, irrespective of the total number of citations.*

*Sources: NBER Utility Patent and Patent Citation Data File (Hall et al. 2001); S&P Global Market Intelligence COMPUSTAT database (2025).*

The network's basic features complement this visual assessment. With only 2,347 existing ties out of 13,655 potential connections, the network's low density of 14.7% suggests a pattern of strategic, rather than random, knowledge acquisition. The visualization also reveals a key structural feature: while firms tend to form dense clusters by citing the same major hubs, creating "interconnected neighborhoods," these clusters do not appear to be isolated. The general absence of separate, tightly bound subgroups, or "cliques," suggests that knowledge flows relatively freely across the entire network via the central hubs, rather than being balkanized into technological silos.

Indeed, the semiconductor firms in the network both draw from and contribute to a broad technological base. The average number of forward citations for semiconductor firms' patents (from all citing firms, both inside and outside this network) is 5,341, with an average generality score of 0.47. Backward citations average 7,511, with an average generality score of 0.51. Notably, 93 firms cite themselves, indicating internal knowledge building and patent portfolio development. These self-citations appear as reflexive loops in the network structure, though they are less visually prominent than the inter-firm connections that shape the overall topology.

The firms in this network are deeply embedded in the broader technological ecosystem, both consuming and producing a vast amount of knowledge. On average, each firm's patent portfolio receives over 5,000 forward citations from later inventions across all industries, signaling its outward influence. At the same time, each firm makes over 7,000 backward citations to prior art, demonstrating its reliance on the existing stock of knowledge. Furthermore, the fact that 93 of the 127 firms cite their own prior patents indicates a significant level of internal, cumulative innovation.

### *Network Leaders*

While further below I formally defend the claim that Texas Instruments, Intel, and Qualcomm are the semiconductor network's clear leaders, here I familiarize readers with each of them in a stylized manner, paying particular focus to their core technologies and flagship inventions. In their own way, these firms played instrumental roles across different semiconductor segments during the analysis period. Texas Instruments did so for analog/embedded processing; Intel for microprocessors; and Qualcomm for wireless technologies and mobile communications. Putting aside these different specializations, they shared a similar playbook: aggressive risk-taking, massive R&D investments, and pioneering new models of industrial organization, from Intel's integrated manufacturing to Qualcomm's "fabless" design-focused approach.

### Texas Instruments

Texas Instruments designs, manufactures, tests, and sells analog and embedded, specialized semiconductors across several applications.<sup>77</sup> During the timeframe subsumed by the analysis, they included industrial production, automotive transportation, personal electronics,

---

<sup>77</sup> Texas Instruments acquired National Semiconductor in 2011, which added about 12,000 analog chip patents to its portfolio. The network datasets only have coverage on patents and their citations and other traits until 2006, however; therefore, I omit these.

communications equipment—including mobile phones—and enterprise systems. The company pioneered Digital Signal Processors (DSPs), specialized chips that became the engine of the digital media and communications revolution. A key patent embodying this breakthrough is US Patent Number 4,577,282, granted by the USPTO in 1986 for the architecture of its landmark TMS320 family of processors (Caudel and Magar 1986). This invention was essentially a specialized computer on a single chip, optimized for a single, crucial task: processing real-world analog signals.

The DSP was groundbreaking because it acted as a high-speed translator between the analog world (like sound waves or radio signals) and the digital world of computers. This innovation made it economically feasible to build the first generation of mass-market digital devices. DSP therefore became the heart of early digital cell phones, high-speed modems, and the complex controllers inside computer hard disk drives.

For other technology firms, Texas Instruments' DSP architecture, particularly the TMS320 series, established a key benchmark for high-performance signal processing in applications like audio, telecommunications, and embedded systems. Companies across the electronics industry sought to license this technology or design products that aligned with it. For example, Nokia relied on Texas Instruments' DSP cores in its digital mobile phones throughout the 1990s, using them to convert voice signals into digital form and to manage the modulation schemes required for 2G networks.<sup>78</sup> The technology's outside influence is reflected in the frequent citations of patents related to the TMS320 architecture, which served as foundational prior art for subsequent innovations in digital devices and helped solidify Texas Instruments' prominent position in the embedded processing market for many years. US Patent 4,577,282 was cited 65 times between 1976 and 2006. Compared to the mean of 9.4 and median of 5 citations across the 2.5 + million patents in the dataset, this puts the TMS320 patent comfortably in the top few percent of all patents (well above the 95th percentile).<sup>79</sup>

## Intel

In Chapter 2 of the book, I discussed Intel's role in spearheading the PC revolution by sustaining "Moore's Law" and relentlessly pushing CPU performance. In that vein, Intel secured a series of landmark patents, including one from 1996 that was crucial for the performance of its Pentium processors. While it did not introduce superscalar execution itself, US Patent Number 5,584,001 on branch prediction in superscalar, out-of-order cores, describes a branch-target buffer that

---

<sup>78</sup> A modulation scheme is the set of rules for encoding digital information (bits of 1s and 0s) onto a radio carrier wave for wireless transmission. Think of a radio station's broadcast frequency as a pure, constant tone (the carrier wave). To transmit a song, the station must vary, or "modulate," that tone in sync with the music. For 2G digital networks, the primary method was a type of frequency modulation called Gaussian Minimum-Shift Keying (GMSK), which represented a digital "1" by slightly increasing the carrier wave's frequency and a "0" by slightly decreasing it. The DSP was the specialized hardware that performed the complex, real-time calculations required for this encoding and decoding process.

<sup>79</sup> The descriptive statistics about patents in this section are from the original NBER patent level source dataset, rather than the patent data aggregated by company I outlined above.

maintains speculative as well as actual branch histories and computes a one-cycle prediction, with mechanisms (including a return-stack buffer) to recover on mispredictions (see Hoyt et al. 1996).

To understand this patent's significance, it helps to define these architectural concepts. A “superscalar, out-of-order” processor is an advanced design that can execute multiple instructions in a single clock cycle and can even reorder them to keep its various components busy. The primary challenge for these complex designs is a “branch”—an “if-then” decision in a program that creates uncertainty about which instructions to execute next. Waiting for the correct path to be determined creates a “pipeline stall”, a costly delay where the processor's power is wasted. To solve this, processors use “branch prediction” to make an educated guess about the branch's outcome and proactively start executing instructions down that path, a technique known as “speculative execution”. The innovation in this patent was a more sophisticated method for making these guesses, using a “branch-target buffer”—a small memory cache on the chip—to keep a detailed history of both actual past outcomes and the success of its own speculative predictions, allowing it to learn and improve its accuracy over time.

To better grasp its contribution, think of a modern processor as a factory with multiple, parallel assembly lines, known as a “superscalar” design. The factory's goal is to keep these lines constantly moving at top speed. However, a computer program is full of “if-then” decisions, or branches, which create a problem: the factory manager must guess which path the product will take down the assembly line before the decision is final. A wrong guess forces the entire line to shut down and restart, wasting precious time. This is a “pipeline stall.”

The Intel patent described a revolutionary new system for “branch prediction” to solve this very problem. It was like giving the factory manager a supercomputer that kept a detailed history of all previous orders and used that data to predict the outcome of the next “if-then” decision with incredible accuracy. This system was so advanced that it even remembered its own past predictions to improve its future guesses. When it did guess wrong, it had an ultra-fast recovery mechanism to clear the line and switch to the correct path with minimal disruption. By drastically reducing these pipeline stalls, this innovation produced the dramatic leap in real-world speed that was essential for running the demanding, graphics-heavy software of the era, most importantly the “Windows 95” user interface.

For other semiconductor firms, the branch prediction patent's influence was immense—it set a new, non-negotiable standard for high-performance chip design: not just another step along the semiconductor path, but the definitive prior art that shaped the trajectory of microprocessor design for a decade. Rivals like AMD could not compete without developing their own, equally sophisticated branch prediction systems. They were forced to either license Intel's technology, or invest heavily in designing around the patent, but were heavily inspired by its pioneering approach.

Consider that, between 1976 and 2006, this patent received 48 forward citations. That places it well above the mean of 9.4 and the median of 5 for all patents in the dataset, putting it in the 98<sup>th</sup> percentile of forward citations. In other words, while most patents languished with only a

handful of references, Intel's branch prediction design became part of the relatively small club of patents that subsequent inventors routinely had to cite.

### Qualcomm

Finally, we come to Qualcomm—one of the first so-called fabless firms, along with others such as Britain's ARM, to focus purely on designing chips that are fabricated in foundries located in Taiwan and South Korea—was one of the most important semiconductor firms of the smartphone era. While, as I explain in the Chapter 6 of the book, it is most famous for contributing to the SEPs that helped run almost all those mobile digital devices by 2025, in the pre-2007 era, the timeframe of this analysis, it designed critical modem chips that allowed mobile phones to communicate with base stations and the wireless cellular infrastructure of the 2G and 3G era: Countless patents embody Qualcomm's contributions to the semiconductors that powered that period's modems, antennae, batteries, and other essential components inside mobile phones.

Consider the company's most cited patent during that period, US patent number 5,103,459, which the USPTO awarded to Qualcomm in 1992 for its so-called Code Division Multiple Access (CDMA) technology (see Fischer and Smith 1992). While I will elaborate on later advances to this invention in Chapter 6 of the book, for now readers should note that this invention is a revolutionary method for allowing multiple mobile phones to share the same radio spectrum simultaneously without interfering with one another. Unlike older methods that divided the airwaves by frequency or time slots, CDMA gives each user a unique digital code. This is like being in a crowded room with many conversations happening at once, but each is in a different language; a listener who knows a specific language (the code) can pick out their desired conversation from the background noise. Moreover, this "spread spectrum" technique was more secure and offered greater capacity than previous standards, making it a foundational technology for the 2G and 3G mobile networks that brought wireless communication to the masses.

The CDMA patent amassed 917 citations before 2007, earning it a spot in the 99.99<sup>th</sup> percentile, far above the mean of 9.4 and the median of 5 citations—essentially among the very top handful of patents in the entire U.S. system during that thirty-year period. This extraordinary citation record underscores how Qualcomm's CDMA invention did more than solve a technical bottleneck: it set the terms of competition for an entire generation of wireless communication. Whereas most patents are cited only sparingly and quickly fade into obscurity, Qualcomm's CDMA patent became the canonical reference point across sectors, from handset makers to infrastructure providers.

### **Formal Network Analysis of the Semiconductor Network**

While Figure S2.3 reveals that the horizontal network possesses a hierarchical structure, exhibiting the type of extreme inequality in citation patterns that characterizes scale-free networks, I now turn to a more formal analysis. It will help me evaluate several important features besides whether it is scale free, such as how well technology seems to flow through it and what type of technology is acquired by more peripheral firms from the network's hubs.

### *Preliminary Tests for Evaluating the Network's Structure*

Let's begin with the network's average path length, which is the statistical equivalent of the famous "six degrees of separation" concept.<sup>80</sup> It measures, on average, how many citation "hops" it takes to get from any firm to any other. A small number indicates a "small world"—where the average path between any two individuals in a large network is much shorter than one would intuitively expect—and information and technology can diffuse very quickly. The analysis reveals that the average shortest path between any two connected firms in the semiconductor network is just 1.79. This incredibly low number suggests a highly efficient and tightly interconnected ecosystem where any given firm is, on average, less than two steps away from any other.

With the average path length established, we can now test if the network exhibits the "small-world" property characteristic of a scale-free network. A common test compares the network's diameter—the longest shortest path between any two nodes—to the theoretical prediction for a scale-free network, which scales with the logarithm of the number of nodes, or  $l \sim \log(N)$ . In this network of 127 firms, the observed diameter is 5. The theoretical prediction is  $\log(127)$ , which is approximately 4.84. The close match between the observed reality (5) and the theoretical prediction (4.84) provides strong evidence that the semiconductor network does indeed have the "small-world" structure of a scale-free network.<sup>81</sup>

Next, we can measure the network's clustering, which is essentially a measure of its "cliquishness." In a social network, it answers the question: "Are my friends also friends with each other?" In our patent network, it asks: "Are the firms that cite the same technology leader also likely to cite each other?" A high clustering coefficient means the network is full of tight-knit, collaborative neighborhoods, while a low coefficient suggests a more open structure where information flows through central hubs rather than within cliques.

More formally, the network's various clustering coefficients measure the dependence between edges in different, overlapping ways. There are three clustering coefficients of interest. First, the local clustering coefficient of a node  $n$  is the probability that two randomly chosen neighbors of  $n$  are themselves adjacent: the fraction of pairs of neighbors of a node that are also each other's neighbors. Second, the overall clustering coefficient is the ratio of existing links connecting a node's neighbors to each other to the maximum possible number of such links. More specifically, it is the fraction of length-2 paths that are closed with a triangle, which is calculated as the number of observed transitive relations divided by the number of possible transitive relations in the network. Third, the average clustering coefficient is the overall clustering coefficient at the

---

<sup>80</sup> The "six degrees of separation" is the theory that any two people on Earth are connected by a surprisingly short chain of social acquaintances. The idea originates from a famous 1960s experiment by psychologist Stanley Milgram, who found that, on average, a letter could be passed between any two random strangers in the U.S. through a chain of only about six intermediaries.

<sup>81</sup> While there are 4,906 paths between nodes that measure 2 (61.32 percent of the total), there are 164 paths that measure 3 (2.05 percent of the total).

node level (the local clustering coefficients) averaged across all the nodes. Each of these measures range from 0 to 1.

How do these coefficients help ascertain whether the semiconductor network is scale free versus random? First, a random network has independent edges, which implies both overall and average clustering coefficient close to 0. For the semiconductor network, however, the average clustering coefficient is .61, and the global clustering coefficient is .49. What this suggests is that citation patterns are not idiosyncratic but follow a generic, structured behavior where firms cluster around the same central sources of innovation, a hallmark of a scale-free network.

The clustering coefficients provide our first strong evidence that the network is not random. A random network would have a clustering coefficient close to 0, but for our semiconductor network, the average clustering coefficient is a high 0.61. This indicates a significant amount of local structure, in that firms are forming clusters by citing the same central sources of innovation. Specifically, many of the semiconductor firms in this network cite the same semiconductor firms, but the citing firms do not necessarily cite each other, and the cited firms do not necessarily cite the citing firms. In social network terms, this is less like a group of friends who are all friends with each other (a closed clique), and more like a group of students who all seek to befriend a network influencer. The low local clustering scores for the main hubs like Texas Instruments (0.30) and Intel (0.26) reinforce this, suggesting they are open hubs broadcasting to a wide audience, not the center of a closed community.

This brings us to the network's node degree distribution. Recall that a node's degree is the total number of its connections. If a node boasts a greater number of "degrees" it is, *ceteris paribus*, more influential; the average degree over all nodes is the average degree of the network. For a scale free network, what we expect is a few nodes with high degree and most of the nodes with very low degree. Specifically, we can seek to ascertain a network's node degree distribution by evaluating whether the semiconductor network is a scale free network that follows a power law (Barabasi and Albert 1999). More specifically, that the probability that a node has  $k$  links is proportional to  $k^{-\beta}$ , such that  $P(k) \sim k^{-\beta}$ .

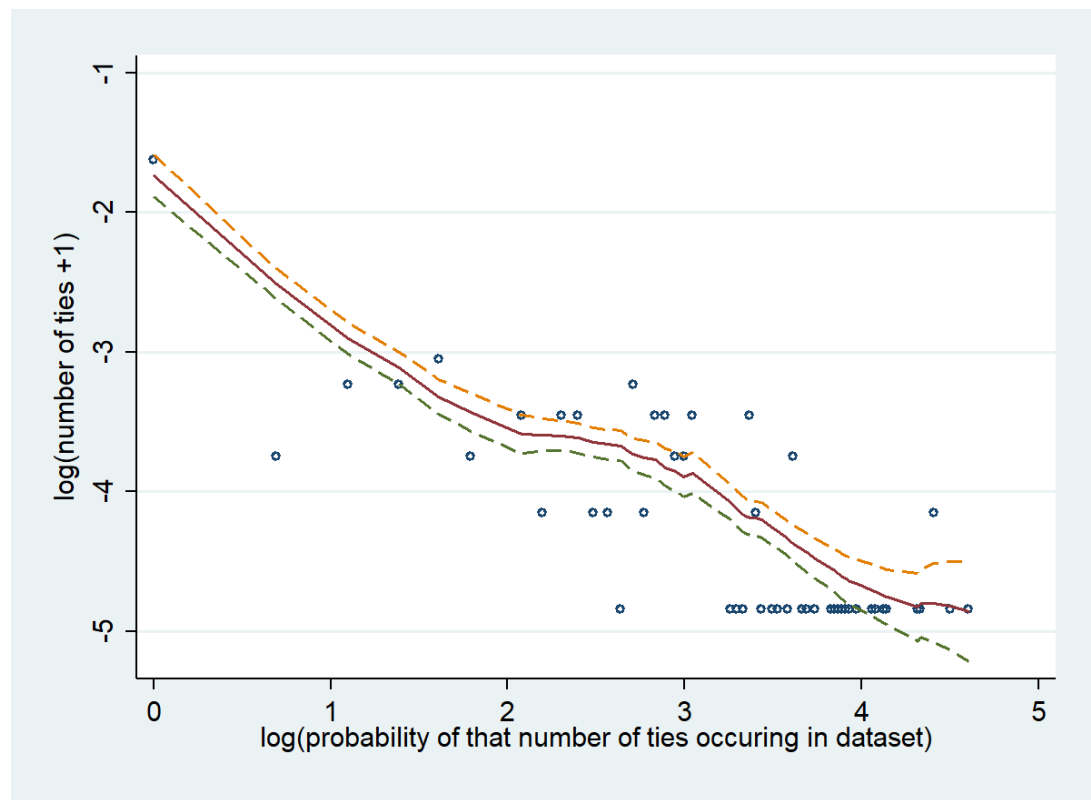
Before I evaluate this question in a systematic manner, first consider some facts. In this network, most firms are not cited all that much by each other. There are several firms that are not cited at all. Specifically, 25 firms. There are also 47 semiconductor firms that are cited by five or fewer semiconductor firms. Therefore, the indegree centralization, the number of incoming links, is only .64 (the outdegree centralization, or out-going links, is .55). Yet, there are a lot of forward citations overall, even if they are unevenly distributed.

#### *Testing Whether the Distribution of Citations Obeys a Power Law*

Now that these preliminaries are out of the way, I can formally assess whether the horizontal network abides by the same Power Law dynamics identified in the cross-sectional data. As established in the previous section, the key test involves examining the log-log plot for a linear relationship and formally testing the goodness-of-fit for the distribution's tail. I therefore proceed with that analysis.

I therefore proceed in a few steps. First, the log-log plot of the probability degree distribution in Figure S2.4 provides strong visual evidence for the hypothesis. The predicted line becomes increasingly linear in the right tail of the distribution, suggesting that the power law kicks in at higher values. This visual intuition is supported by a high R-squared of .85 for the relationship. While highly suggestive, this is not dispositive proof. Therefore, to formally test the goodness-of-fit and locate the precise threshold, I employ the Urzúa (2020) technique as I did previously when addressing this same issue for the cross-sectional dataset. The formal analysis corroborates the visual intuition perfectly. After testing various percentiles, I fail to reject the hypothesis that the data follows a power law at the 75th percentile of the distribution (equal to 29.1 ties), with a p-value of .11. This confirms that the tail of the distribution is indeed governed by a power law, allowing us to now move on to formally identifying the network's leaders.

**Figure S2.4 Power law fitted to Probability Degree Distribution**



*Notes: The predicted line for  $\log(\text{number of ties} + 1)$  is estimated using a running line smoother. I estimate each smooth by employing a backfitting algorithm and a running-line smoother for the  $\log(\text{probability of the observed number of ties occurring})$ . The smoother is a linear function for each observation of that variable and 95 percent confidence intervals are calculated from standard errors for each smooth.*

*Formal Identification of Horizontal Network Leaders*

Several statistics corroborate the idea that Texas Instruments is the central player in this network. It ranks first in influence, with 99 other semiconductor firms—a staggering 78% of the entire network—citing one or more of its patents. In total, during the analysis period, over 2,800 different Texas Instruments patents were cited more than 40,000 times by its semiconductor peers alone.<sup>82</sup>

Beyond raw citation counts, several network centrality measures confirm Texas Instruments' pivotal role as the network's primary intermediary. The first of these is betweenness centrality, which identifies the critical "bridges" in a network. Specifically, the betweenness centrality of a node is the number of shortest paths among all other nodes that pass through this node, normalized by the maximum number of paths that a given node could lie on between pairs of other nodes.<sup>83</sup> Think of it like a major hub airport: a node has high betweenness if a greater number of the shortest paths between other nodes must pass through it. The data reveals that Texas Instruments is the main hub of this network, with a massive betweenness score of over 1,800. This is even more striking when compared to the median score for a typical firm, which is just 1.5, and the 59 firms that have a betweenness centralization score of 1 or less (46.5 percent of all semiconductor firms). This shows that Texas Instruments is not just a source of innovation; it is the chief conduit through which technology is transferred between other firms in the industry.

Texas Instruments' status is also corroborated by its network-leading eigenvector centrality score of 0.20. This metric measures influence based on the principle that it's not just how many connections you have, but who you are connected to. In other words, a node's centrality is weighted by both its number of neighbors and those neighbors' own network importance. The most famous example of this is Google's original PageRank algorithm: a link to your website from an authoritative source like a major university is worth far more than a link from an unknown blog. Similarly, a patent citation from an influential firm contributes more to a company's score than a citation from a small, peripheral firm. Texas Instrument's top score therefore indicates that it is not just widely cited, but that it is cited by other important and central players in the industry.<sup>84</sup>

Moreover, we can look at closeness centrality. This metric identifies which nodes are in the best position to spread information quickly to the entire network. Specifically, it is the inverse of the average shortest distance between a node and all other nodes in the network.<sup>85</sup> Think of it like finding the optimal location for a national distribution center; you want the spot that has the shortest average travel time to all other cities. In this network, Texas Instruments again has the

---

<sup>82</sup> Its total forward citations across all its patents, irrespective of what type of firm cites it, are 151,373.

<sup>83</sup> When there are multiple shortest paths between two nodes, I weight each shortest path as a function of the inverse of the number of shortest paths between them.

<sup>84</sup> I calculate this score after excluding the isolates because the score depends on the number of connections between nodes.

<sup>85</sup> The formula is  $1/(\text{average distance to all other nodes})$ . I define the distance between unconnected nodes as the maximum distance found in the network+1.

highest closeness centrality score (0.728). This indicates that its central position allows a new technology or piece of information originating from Texas Instruments to diffuse through the entire semiconductor industry faster than it could from any other point.

Beyond its central position, we can now test the kind of technology flowing from a network hub like Texas Instruments. As Hall et al. (2001) argue, firms whose patents receive highly general forward citations are most likely introducing GPTs, while firms that cite those patents are assimilating them. Patent generality scores range from 0 to 1, with values approaching 1 indicating the patent is cited broadly across many different technological domains (suggesting general-purpose technology characteristics), while values approaching 0 indicate citations are primarily from within the same narrow technological field (suggesting specialized applications). For example, a patent cited equally by 10 different technology classes would earn a Generality Score of 0.90, while a patent cited only within its own technology class would score 0. A high generality score awarded to a patent's forward citations means it makes a broader impact across several applications: it influenced follow up innovations across a variety of fields.<sup>86</sup> Thus, within a single industry, this measure can map the flow of technology from sectoral leaders to peripheral firms.

To validate the generality score as a reliable proxy for a patent's technological breadth, I conducted a regression analysis linking it to the number of a patent's inventive claims—a measure of its scope.<sup>87</sup> The results confirm the score's validity. The relationship is positive and highly significant: a one standard deviation increase in claims corresponds to a 63% of a standard deviation increase in forward citation generality ( $p < .001$ ), with an impressive R-square of .40.

With this validated measure of patent generality at hand, I can directly assess how knowledge flows. The core question is whether citation patterns exhibit homophily (firms citing others with similar, specialized technology) or heterophily (firms seeking out different, foundational technology). I expect the latter because, as already outlined, if peripheral firms are indeed seeking out foundational GPTs, then firms with specialized, low-generality patents should be more likely to cite innovation hubs like Texas Instruments, which possess foundational, high-generality patents. If this "opposites attract" dynamic exists, then it should show up as a strong

---

<sup>86</sup> To better understand why that is the case, consider how the patent generality index is constructed. It is based on the idea of Herfindahl concentration and capturing the percentage of citations received by a patent belonging to different patent classes out of the total universe of patent classes (Hall et al. 2001). Specifically, patent generality is calculated as:  $\text{Generality} = 1 - \sum_i (s_i)^2$ , where  $s_i$  represents the share of citations received from patent class  $i$ .

<sup>87</sup> A patent's claims are the detailed specifications of the invention's building blocks. The number of claims may indicate a patent's "scope" or "width" (Hall, et al. 2001: 23). The mean number of claims in each patent is 13.6 and the median is 13.8. Very few firms have many claims on their patents. The standard deviation is 13.5 claims per patent. The company whose patent portfolio has the greatest number of claims is Scanner Technologies Corp., which manufactures semiconductor machinery, specifically, 2D and 3D coplanarity inspection technology for electrical circuit boards, with 126 claims (on average) per patent.

dyad level correlation between network ties and the absolute difference in the generality scores of the citing and cited firms using a QAP approach.<sup>88</sup>

The results of the correlation test confirm the hypothesis exactly as expected. For Texas Instruments, there is a statistically significant negative correlation (-0.20,  $p < 0.001$ ) between a firm's similarity to Texas Instruments in patent generality and its likelihood of citing the latter. In simple terms, this means that "opposites attract": firms with specialized, low-generality patents were the most likely to cite the foundational, high-generality patents of Texas Instruments. Conversely, when looking at ties between firms other than Texas Instruments, there was no statistically significant relationship. This provides strong evidence that the flow of knowledge in the network is not random; peripheral "spoke" firms are specifically seeking out the "hub" for its general-purpose technology.

In short, Texas Instruments is not only the semiconductor network's citation leader but the network's key intermediary. Aside from generating major innovations that are cited by other semiconductor firms, Texas Instruments seems to be a chief conduit of information transfer between other nodes in the network. And the information that is being transferred appears to be GPTs flowing from Texas Instruments to peripheral firms.

#### Exploring Other Leaders in the Semiconductor Network: Intel and Qualcomm

Starting with Intel, I now explore the leadership roles of that storied semiconductor firm in the horizontal network and will then also do the same for Qualcomm. The data confirms its role as the network's second most influential hub, trailing only Texas Instruments across nearly every key metric. In terms of raw influence, it is the second most cited firm within the semiconductor network, with 89 other firms citing its patents. Its broader impact is demonstrated by its 102,436 total forward citations, ranking it third in the entire dataset.

Intel's importance as a key intermediary is confirmed by its betweenness centrality score of 1555.86, the second highest in the network. This massive number, which dwarfs the network's median score of just 1.5, proves that Intel, like Texas Instruments, acts as a critical bridge for technology transfer, with over 1,500 of the network's shortest communication paths flowing through it.

Its leadership status is further cemented by its other centrality scores. Intel boasts the second-highest eigenvector centrality (0.19), indicating it is highly connected to other influential firms, and the second-highest closeness centrality (0.728), confirming it is in a prime position to spread innovation efficiently throughout the network.

---

<sup>88</sup> I obtain these correlation coefficients by first creating a new dataset where the unit of analysis is each dyadic pair and then estimating a QAP regression with 50 permutations of the rows and columns (simulations). These permutations allow us to preserve any dependence among elements of the same row or column while also eliminating any relationship between the dependent variable and the independent variable, with the null hypothesis that, for each dyad pair, there is no systematic correlation between a directed tie and the absolute distance in the generality of patents' forward citations.

Just as in the analysis of Texas Instruments, I next examine whether Intel occupies a similar position as a transmitter of general technologies within the semiconductor network. Applying the same dyad-level correlation test between network ties and an artificially generated network based on the generality of firms' forward citations (using the QAP approach), I find a comparable pattern. A node is more likely to send a tie to Intel when it differs from Intel in terms of patent generality: the correlation coefficient is  $-0.20$  ( $p < .001$ ). By contrast, a node is neither more nor less likely to form a tie with firms other than Intel based on differences in generality (correlation =  $-0.04$ ,  $p = 0.30$ ). This mirrors the results obtained for Texas Instruments and supports the interpretation that Intel, like Texas Instruments, acts as a conduit for general-purpose technological knowledge, diffusing broad, cross-domain innovations from the network's core to its periphery.

Finally, we reach Qualcomm, the last of the "Big Three" semiconductor firms in the network. At first glance, its leadership position is less obvious than that of Texas Instruments or Intel. With only 39 semiconductor firms citing it, Qualcomm ranks just 19<sup>th</sup> in direct network ties. Moreover, the 22,120 total forward citations to its patents by both semiconductor and non-semiconductor firms make it the 6<sup>th</sup> most cited semiconductor firm. However, a deeper at the quality of its patents and the nature of its connections reveals its unique and critical role.

Take the fact that Qualcomm exhibits the highest average patent generality score, 0.65, among the top 20 cited firms, substantially exceeding the non-Big-Three average of 0.55. This exceptional generality score suggests its technologies enjoy broad applicability across several domains. And the QAP-style correlation analysis I undertook for both Texas Instruments and Intel strengthens this conclusion: Qualcomm displays the strongest negative correlation ( $-0.24$ ,  $p < 0.001$ ) among the Big Three for citation patterns based on patent generality, surpassing both Texas Instruments ( $-0.20$ ) and Intel ( $-0.20$ ). This means firms citing Qualcomm are the most dissimilar to it in patent generality, confirming that specialized firms seek out Qualcomm specifically for access to foundational, broadly applicable wireless technologies.

Qualcomm's strategic network position further validates its leadership status through several complementary metrics. It places in the top tier for critical centrality measures: 22<sup>nd</sup> in betweenness centrality (70.39, top 17% of all firms), 19<sup>th</sup> in Eigenvector centrality (0.14), and 20<sup>th</sup> in closeness centrality (0.565). These scores indicate Qualcomm serves as an important intermediary for knowledge flows and maintains connections to other influential nodes.

Perhaps most revealing is Qualcomm's unique relationship with other network leaders: every single firm citing Qualcomm also cites Texas Instruments (correlation of 0.50), while its correlations with Intel (0.44) and Texas Instruments (0.50) are notably lower than the Texas Instruments-Intel correlation (0.65). This pattern suggests Qualcomm doesn't duplicate but rather complements other leaders' technologies—firms need both Texas Instruments' broad semiconductor capabilities and Qualcomm's specialized wireless innovations.

The gap between Qualcomm's within-network citations and its 22,120 total forward citations (6<sup>th</sup> overall) further indicates its influence extends well beyond traditional semiconductor boundaries into telecommunications, consumer electronics, and software sectors. This bridging function,

combined with its foundational CDMA patents that revolutionized wireless communication, positions Qualcomm as an essential network leader whose influence transcends simple citation metrics to shape entire technological trajectories. This is a topic we now explore further ahead.

### *Indirect Connections between the Big Three*

It is also important to understand how these three leaders relate to each other. We can do this by measuring how similar their "audiences" are. In a social network, we would consider two influencers to be similar if they are followed by the same group of people. In our patent citation network, we do the same thing: we measure the similarity between two leaders (like Intel and Texas Instruments) by calculating the correlation between the lists of firms that cite them. This tells us if they are influential in the same technological neighborhood.

In network parlance, to find out if two nodes are connected to the same "alters" we calculate how correlated the tie vectors of two nodes are with each other. Specifically, we calculate the correlation between the vector of incoming ties of node  $i$  and the vector of incoming ties of node  $j$ .<sup>89</sup> The correlation between Texas Instruments and Intel's alters is .65. In fact, firms that cite Intel but don't cite Texas Instruments only include four firms.<sup>90</sup> And Firms that cite Texas Instruments but do not cite Intel include 15 firms.<sup>91</sup> Meanwhile, the correlation between Texas Instruments and Qualcomm is .50 in terms of the firms that cite both firms. There are no firms that cite Qualcomm but do not cite Texas Instruments. There are 60 firms that do not cite Qualcomm but do cite Texas Instruments. Finally, the correlation between Qualcomm and Intel is .44 in terms of the firms that cite both firms.

### *Technology Flow through the Horizontal Network*

We can now move onto the related topic of the flow of technology through the network. A relative lack of cliques is beneficial for technology exchange, as it represents fewer speed bumps in the way of information flowing between network nodes. To investigate how "cliquish" the network is, we can look at its transitivity, which can be thought of as the "friend of a friend is also my friend" principle applied to patent citations. A network with high transitivity would be full of closed, triangular relationships and self-contained cliques.

---

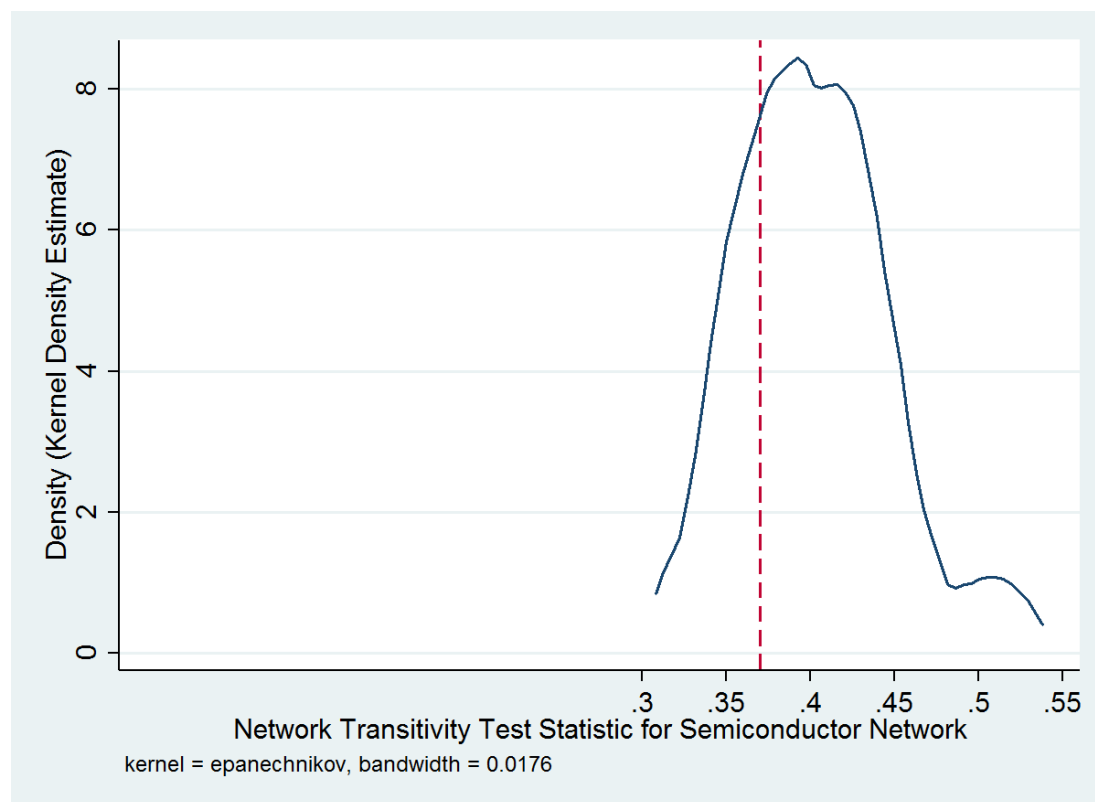
<sup>89</sup> It is also the case that these three firms are tightly linked to each other: Intel and Qualcomm both cite Texas Instruments, and Texas Instruments cites these two firms' patents as well. Intel patents cite 2,597 Texas Instruments patents. Qualcomm patents cite 101 Texas Instruments patents. Texas Instruments patents cite 1,592 Intel patents and Qualcomm patents cite Intel patents 71 times. Finally, Texas instruments patents cite 283 Qualcomm patents and Intel patents cite Qualcomm patents 61 times.

<sup>90</sup> They are Color Kinetics Inc., Integrated Circuit Systems, Mosys Inc., Netsilicon Inc., and Photonics Inc.

<sup>91</sup> Advanced Analogic Tech, Aviza Technology Inc., Energy Conversion Dev., Evergreen Solar Inc., Hittite Microwave Corp., Logic Devices Inc., Microtune Inc., Monolithic Power Systems Inc., Stratos International Inc., Techwell Inc., Tessera Technologies Inc., Trident Microsystems Inc., Triquint Semiconductor Inc., Universal Display Corp., and White Electronic Designs Corp.

More formally, the network's transitivity is the overall probability of adjacent nodes being interconnected. is not abnormally high. Figure S2.5 positions our network's transitivity score within the distribution generated by simulating 50 random networks that share the same structural parameters (127 nodes and 2,347 ties), revealing that our observed transitivity aligns with what chance alone would produce—a transitivity score of 0.37 for a network with a density of 0.15. Moreover, local clustering coefficients reveal that there are only 12 firms exhibiting “citation clique” behavior operationalized as a clustering coefficient equal to 1.<sup>92</sup> The conclusion is that the semiconductor network is not abnormally cliquish. While a few small, perfect cliques exist, the overall structure is open, suggesting that knowledge flows through the central hubs rather than getting trapped in isolated, self-referential subgroups.

**Figure S2.5 Semiconductor’s Network Transitivity Score**



*Notes: the figure plots where the network’s transitivity score versus the distribution of transitivity scores for 50 (simulated) random networks with the same network structure: the same number of nodes and ties as the semiconductor firm network.*

This structural openness is not merely a technical detail; it is the economic safeguard against economywide stagnation. When this diffusion mechanism is broken—when “Superstars” become isolated islands rather than central hubs—the result is a “Dual Economy” of the type that plagued

<sup>92</sup> These firms are Diodes Inc., DPAC Technologies Corp., Hei Inc., Hi/Fn Inc., Insilicon Corp., Logic Devices Inc., Microtune Inc., Mosys Inc., NVE Corp., Photronics Inc., PLX Technology Inc., and Sigma Designs Inc.

Japan for decades after its economic meltdown in 1989 (Katz 1998; McKinsey Global Institute 2020; Katz 2025). While Japanese giants like Toyota and Sony achieved “Global Frontier” status through fierce efficiency, their innovations failed to diffuse to the domestic service sector, which was shielded from the competitive pressures that drive adoption. Because the network remained bifurcated rather than open, Japan suffered a “High-Technology/Low-Productivity Trap” where the benefits of innovation remained locked inside a few corporate giants. The open, clique-free structure of the U.S. semiconductor network revealed above suggests that American high-tech industries avoided this fate during the Third Industrial Revolution precisely because their “Superstars” remained highly connected conduits for broad-based diffusion.

## DIGITAL DEVICES SUPPLY CHAIN NETWORK

I now transition to analyzing the second network database, which maps the electronic digital device supply chain before 2007—before the smartphone revolution launched by Apple with the release of the iPhone. It is a “vertical” network that observes connections between upstream semiconductor firms, which composed our first network, and firms operating downstream. Those include hardware manufacturers specializing in fabricating personal computers and mobile phones such as Motorola.<sup>93</sup> Other downstream firms circa this period also included firms involved in wireless telecommunications, such as AT&T, and software publishers, such as Microsoft.

This network captures a supply chain in rapid evolution, a story I touched on in Chapter 2 of the book. The era begins in 1976, with early minicomputers and personal computers—clunky, often user-assembled devices with lackluster performance but capable of basic word processing and calculations.<sup>94</sup> It ends in 2006, on the eve of the iPhone, with the industry firmly oriented towards a new generation of mobile and interconnected devices: laptops, early mobile phones, and personal digital assistants with wireless capabilities.

This evolution was driven by rapid advances across the technology stack. Improvements in integrated circuits and battery technology made smaller, more powerful devices possible, while the integration of modems connected laptops and mobile phones to the nascent internet. These hardware enhancements were facilitated by a constant flow of technology from upstream semiconductor firms to downstream device makers and were complemented by parallel advances in telecommunications and software.

---

<sup>93</sup> Motorola, which I treat as a downstream hardware manufacturer to capture the era of rising cellular communication between circa 1980 and 2006 in the runup to the introduction of smartphones, also contributed semiconductor innovations during this period (see Suntech Display Technology 2019). I note that in 2012 Google bought Motorola’s hardware division (Motorola Mobility) and its almost 25,000 strong patent portfolio and then sold Motorola to Lenovo in 2014 after keeping most of the acquired patents.

<sup>94</sup> These machines were smaller and more affordable than mainframes. They were outfitted with input-output devices (e.g., teleprinters), basic memory, and could run programs using Fortran or BASIC.

To capture this dynamic, I constructed a bipartite network that maps only the vertical ties between the upstream semiconductor firms and the downstream hardware, software, and telecommunications firms. This structure explicitly excludes the horizontal, within-industry ties from the previous analysis to focus exclusively on the flow of technology up and down the supply chain. A link is therefore only recorded when, for example, a computer manufacturer cites a semiconductor firm, or a semiconductor firm cites a software publisher.

Moreover, I have omitted two high tech economic sectors from this network database. The first is Data Processing and Hosting. There are 61 firms in this sector, and these firms include prominent merchant payment companies, such as Paypal Holdings Inc., and First Data Corporation.<sup>95</sup> The second is Computer Systems Services. This category includes Custom Computer Programming Services, Computer Systems Design Services, Computer Facilities Management Services, and Other Computer Related Services. There are 206 firms in this sector. The most prominent ones are IBM, Unisys Corp., Novell Inc., and Lucent. Other noteworthy firms are Perot Systems and TeraData Corp.<sup>96</sup>

The reason for excluding these two industries is that these sectors are only indirectly involved in the supply chain I care about here: the pre-smartphone and pre-Fourth Industrial Revolution electronic digital devices that could connect to the internet and, to a certain extent, used wireless networks, including cellular ones, to do so. For the most part, firms in the excluded industries manufactured peripheral equipment and/or provided business to business services during that era that were not critical to the functionality of desktops, laptops, proto-smartphones, such as the Blackberry—produced by Research in Motion Inc., it was first launched in 1999—and other pre-2007 personal assistants and precursors to tablet computers such as Palm Pilots.

This focus is also a reflection of the technological era. The pre-2007 digital world was defined by “on premises” computing, where companies stored, processed, and transmitted data on their own internal servers and private intranets.<sup>97</sup> While the seeds of the modern cloud were planted in

---

<sup>95</sup> Compared to other high-tech sectors, it is not clear that patents play a major role in the Data Processing and Hosting sector. If I exclude Xerox Corporation, which has 13,649 patents and 145,616 forward citations, as well as Harris Corporation, which is a defense contractor that produces wireless equipment, radios and antennas for the US military, the firms in this industry have very few patents between them, with only 189 total and an average of 3.2 each.

<sup>96</sup> Compared to Data Processing and Hosting, it does appear that patents are central to firms operating in Computer Systems Services. Between them they boast a huge number of patents, with a total of 681,417; IBM holds a staggering 77 percent of the total. Moreover, IBM was very important to the personal computer revolution. While these facts seem to suggest that I should have included this sector in the dataset, it is very hard to disentangle IBM’s role in the computer supply chain from its role providing business services as, over time, IBM shed the hardware aspects of its business. In the runup to 2007 it had begun to commit to AI, cloud computing, data management and facilitating merchant payment services. And it sold its personal computing division to Lenovo, a Chinese firm, in 2005. Most important, including this sector in the dataset makes no material difference to the overall patterns I discuss below.

<sup>97</sup> Both before and after 2006, some of these omitted firms help big companies run internal servers, sustain local area networks (LAN), and manage their data.

2006 with the launch of Amazon Web Services (AWS), the business analytics and cloud-based services that are critical to the Fourth Industrial Revolution were not yet central to the digital device supply chain of the period this analysis covers, which was centered on.

### **The Vertical Network's Basic Characteristics**

The basic features of this supply chain network reveal a structure that is both sparse and efficient. The network is not dense: 289 firms are connected by only 2,186 ties, only about 3% of all possible connections. The low reciprocity score of 0.29 further indicates that knowledge flows are often one-way, "unrequited" transfers from one layer of the supply chain to another. Despite this sparseness, the network is highly efficient. With an average path length of just 2.14 and a diameter of only 5, it exhibits the classic "small world" property, meaning that technology and information can travel from any one firm to any other in just a few short steps.<sup>98</sup>

To gauge a firm's importance in this supply chain network, we can measure the degree to which it intermediates relationships between other firms. As discussed previously, this can be measured in three ways: Betweenness, Closeness, and Eigenvector Centrality. Unlike the clear hierarchy of the semiconductor network, the leadership picture in this vertical network is much murkier, suggesting a more complex web of technology transfer up and down the supply chain.

The firms with the highest betweenness centrality numbers are all hardware manufacturers. The top five are, respectively, Concurrent Computer Corp. (14,698.98), Steelcloud Inc. (13,203.08), Compaq Computer Corp. (11,473.91), Palm Inc. (7,735.457), and Motorola Inc. (2,5642.012). Other notable firms include Apple Inc. (211.778), ranked 25<sup>th</sup>, Texas Instruments Inc. (111.823), ranked 49<sup>th</sup>, Intel Corp. (0.104), ranked 196<sup>th</sup>, and Qualcomm Inc. (0), ranked 221<sup>st</sup>. To keep things in perspective, the network's mean level of betweenness is 263.568, and the median is 4.458.

The betweenness centrality scores reveal a surprising pattern. Unlike in the horizontal network, the key intermediaries here are not the major semiconductor firms but a group of specialized hardware manufacturers. Firms like Concurrent Computer Corp., Compaq, Palm, and Motorola dominate the top of the rankings, with scores in the thousands. This is a striking contrast to the major semiconductor hubs: Texas Instruments ranks just 49<sup>th</sup>, while Intel and Qualcomm have scores near zero.<sup>99</sup> This finding suggests that in the vertical supply chain, it was the downstream

---

<sup>98</sup> It is also important to note that a significant number of firms exist outside the main network of knowledge flows. In total, 68 firms are not cited by any other firm in the supply chain and, of these, 13 are semiconductor firms that are complete isolates, neither citing nor being cited by any other company in the dataset.

<sup>99</sup> The top five firms by betweenness centrality are all hardware manufacturers: Concurrent Computer Corp. (14,698.98), Steelcloud Inc. (13,203.08), Compaq Computer Corp. (11,473.91), Palm Inc. (7,735.46), and Motorola Inc. (2,542.01). For comparison, other notable firms ranked much lower: Apple Inc. was 25<sup>th</sup> (211.78), Texas Instruments Inc. was 49<sup>th</sup> (111.82), Intel Corp. was 196<sup>th</sup> (0.10), and Qualcomm Inc. was 221<sup>st</sup> (0). The network's mean betweenness centrality was 263.57, while the median was just 4.46, highlighting the extreme influence of the top hardware firms.

device makers who acted as the critical bridges, connecting the innovations from the upstream chip designers to the broader market.

The closeness centrality scores reinforce the "murkier" leadership picture in this vertical network. Unlike the previous metric, which was dominated by hardware firms, this measure of who can spread information most quickly is distributed across all sectors. The top of the rankings includes a mix of software publishers like Verity Inc., semiconductor firms like Advanced Analogic Tech, and hardware manufacturers like Compaq.<sup>100</sup> This finding suggests that while hardware makers may have acted as the primary structural bridges, firms from all parts of the supply chain were effectively positioned to be efficient broadcasters of technological information.

Finally, eigenvector centrality reveals the network's prestige hierarchy. This metric measures influence based on the principle that "it's not just who you know, but who you know knows"; a citation from an influential firm like Intel is worth more than one from a small, peripheral company. When we apply this "prestige-weighted" measure, the top of the rankings is dominated by the most famous and influential technology titans of the era, including giants like AT&T, Sun Microsystems, Motorola, Apple, Intel, and Microsoft.<sup>101</sup> This confirms that, while specialized firms acted as important bridges and broadcasters, the overall technological gravity of the supply chain still centered on these established industry leaders.

### **Core downstream firms in the network**

There are several ways to identify and measure the core downstream firms in the pre-2007 electronic digital device network. In this section, I seek to do so; as well as evaluate how these firms stack up against, and operate in relation to, the so-called Big Three (upstream) semiconductor firms. To do so, I consider firms' total citations, their sheer logistical influence, and the degree to which firms play an intermediary role in the network. Like when I was describing and analyzing the semiconductor network, I again lean on contextual knowledge of the industry to focus on three downstream firms in particular: AT&T, Apple, and Motorola.<sup>102</sup> As I outline directly below, AT&T's dominance across several key statistics makes it a natural candidate as a network leader. And while Motorola ranks fourth in terms of ties with other firms

---

<sup>100</sup> The top ten firms by closeness centrality are: Verity Inc. (0.653), Advanced Analogic Tech (0.641), Midway Games Inc. (0.622), Maxim Integrated Products (0.599), Microtune Inc. (0.560), Optical Communication Products (0.523), I2 Technologies Inc. (0.509), Microchip Technology Inc. (0.507), Zilog Inc. (0.492), and Compaq Computer Corporation (0.488). The network's mean closeness centrality is 0.43 and the median is 0.44.

<sup>101</sup> The top ten firms by eigenvector centrality are: AT&T Inc. (0.296), Sun Microsystems Inc. (0.280), Compaq Computer Corp. (0.269), Motorola Inc. (0.257), Apple Inc. (0.219), Intel Corp. (0.167), Texas Instruments Inc. (0.155), Microsoft Corp. (0.149), Gateway Inc. (0.131), and Advanced Micro Devices (0.128). The mean eigenvector centrality for the network is 0.043 and the median is 0.035.

<sup>102</sup> There are other important nodes in the network I will not focus on here; these include Microsoft, which ranks ninth with 43 ties; Gateway Computers, which ranks fifteenth with 28 ties; Oracle, which ranks 22<sup>nd</sup> with 20 ties; and Palm, which ranks 28<sup>th</sup> with 16 ties.

in the network with 126, and Apple ranks fifth with 103 ties, both companies created a bridge from the Third to the Fourth Industrial Revolutions by contributing innovations that culminated in the smartphone, the digital app revolution, and the widespread use of AI based on mining huge datasets.

### *AT&T*

Beginning with AT&T, the data immediately reveals its dominant position as the most connected firm in the vertical network. It has ties to 166 other companies in the supply chain, and crucially, this includes 60 different semiconductor firms that cite its patents, demonstrating its deep integration with the upstream innovation engine.<sup>103</sup> Moreover, its own patent portfolio cites all the major upstream players, including Intel, Texas Instruments, and Qualcomm.

AT&T's overall patenting record is equally impressive. During our 1976-2006 analysis period, alone, it amassed nearly 12,000 patents cited over 200,000 times. And this represents only a fraction of its historical innovative output, which includes a staggering 30,000 patents since its inception in 1885.

The reason for AT&T's exceptional position is simple: Bell Labs. Although officially a telecommunications company, AT&T inherited the legendary in-house research and development laboratory after its 1980s breakup.<sup>104</sup> This meant that in addition to its long-distance business, it retained control over the patents for some of the 20<sup>th</sup> century's most foundational electronic inventions, including the transistor, giving it a unique and powerful position that bridged the worlds of telecommunications and semiconductors.<sup>105</sup>

AT&T was also critical in creating and launching the physical cellular network that would power the mobile revolution. Beginning in the 1970s, its engineers at Bell Labs designed the Advanced Mobile Phone System (AMPS), which enveloped the United States in a grid of hexagonal "cells," each anchored by a base station that could send and receive messages over radio frequencies. These base stations acted as the crucial bridge, connecting the wireless signals from the new mobile phones to the existing national telephone system.<sup>106</sup>

A single patent, U.S. Patent 4,672,658, illustrates the kind of foundational technology AT&T was producing (see Gilhousen et al. 1987). This groundbreaking patent, which boasts a sky-high

---

<sup>103</sup> These include Intel Corp., Texas Instruments Inc., Advanced Micro Devices, LSI Corp., National Semiconductor Corp., Micron Technology Inc., Altera Corp., and Qualcomm Inc.

<sup>104</sup> After several Baby Bells were spun off in the 1980s, it also retained its long-distance service and the Western Electric Co., its manufacturing subsidiary.

<sup>105</sup> Bell Labs contributed several milestone technological developments during the 20<sup>th</sup> Century besides the transistor: the laser, solar cells, fiber optics, and satellite communications. Six Nobel prizes, including for the invention of the transistor, were awarded to Bell Labs researchers. AT&T also made contributions to computer software, systems engineering, audio recording, and digital imaging.

<sup>106</sup> Interference was ruled out because two adjacent cells operated at different frequencies; mobile phones were able to seamlessly switch between frequencies as they moved from one cell to another.

generality score of 0.84 and was cited over 240 times, describes a novel wireless local network—a direct precursor to modern Wi-Fi. It outlined a system where multiple users in a local area could connect wirelessly to each other and to an external network using a new kind of routing device. This innovation required a combination of new semiconductor, hardware, and software designs, making it a foundational piece of prior art for the interconnected world to come. And while AT&T had other key patents that drove innovation and technology transfer in the pre-2007 era, I now turn to investigating Motorola's role as one of the most important downstream firms in the electronic digital device supply chain during this period.

### *Motorola*

Motorola, a storied communications company, spearheaded the commercialization of mobile phones before the digital era was in full swing. In 1984 it developed the first commercial mobile phone, the DynaTAC, which used AT&T's AMPS cellular system. Nicknamed the "Brick," this iconic phone, which retailed for the equivalent of over \$30,000 today, became a ubiquitous status symbol and ushered in the 1G analog standard. This commercial breakthrough was not just a marketing feat; it represented Motorola's unique position as a systems integrator, combining its own radio engineering expertise with foundational network technology from AT&T and crucial semiconductor components from upstream suppliers.<sup>107</sup>

Motorola also played a pivotal role during the 2G era, which began in 1987 and was marked by the advent of the Global System for Mobile Communication (GSM) common standard.<sup>108</sup> The GSM standard allowed for interoperability across national borders, transmitted communication digitally, and allowed users to send texts. Moreover, it introduced so-called SIM cards, which allowed consumers to switch phones while retaining their phone numbers and data. Motorola was a key contributor to this standard, contributing several of the 1,300 essential patents that allowed mobile phones to use the network in an interoperable manner. It also manufactured several iconic 2G phones. For example, in 1992 it introduced its 5.9-ounce MicroTAC Ultra Lite phone, which was the world's lightest mobile phone at the time.

### *Apple Computer*

For its part, while Apple is today best known for pioneering the modern smartphone with the introduction of the iPhone in 2007, for the purposes of this analysis I will focus on Apple's role as a key downstream firm within the personal computer supply chain. To be sure, I will explore

---

<sup>107</sup> As a first-generation analog device, the DynaTAC's internal electronics were a dense collection of custom chips. At its heart was a microprocessor (from Motorola's own 6800 series) that managed the phone's logic. It contained ROM chips to store the operating software and a small amount of RAM for storing a few dozen phone numbers. Instead of a single digital "modem chip," its radio functions were handled by a complex suite of analog and mixed-signal semiconductors responsible for modulating and demodulating the analog voice signals for the AMPS cellular network.

<sup>108</sup> This was not the only standard used in the United States. The other was Code Division Multiple Access (CDMA), which was used by Sprint and Verizon, and was not compatible with GSM.

that story and its role in the Fourth Industrial Revolution in the book's later chapters—from its creation of the App Store and the app economy it fostered, to its development of custom mobile processors and its unique, privacy-centric approach to on-device AI. However, for the pre-2007 era covered by this analysis, the company's most significant contributions were to the Third Industrial Revolution: It was Apple Computer that played a critical role in introducing personal computing to the masses with its launch of the “Apple II” in 1977, just one year after our dataset begins. Several of the patents associated with that machine and its successors were crucial in helping disseminate cheap computing capable of performing commercially valuable functions like spreadsheets, word processing, and graphic design.


Moreover, in the run up to the introduction of the iPhone in 2007, Apple also built up a patent portfolio that would prove critical to that device. And many of those patents were inspired, or at least had to acknowledge, technologies that were developed by key players in the PC supply chain and pre smartphone cellular device era. As we shall see ahead, these include several inventions ushered in by the Big Three chip firms as well as other firms that specialize in both hardware and software. Specifically, these technologies encompass foundational wireless modem technologies from Qualcomm, digital signal processors for audio from Texas Instruments, the low-power processor architecture licensed from ARM, and crucial power management techniques pioneered by firms like Intel.

The question, therefore, is how much of Apple's innovation relied on technology that trickled down from its upstream suppliers. While that is a central theme of Chapter 6 of the book, I can preview the answer by first outlining Apple's connections to its key chipmakers before turning to a similar analysis of Motorola.

### **Paths Between Important Downstream Firms and Central Upstream Semiconductor Firms**

The most important feature of this network is the diffusion of technology from upstream semiconductor firms to downstream firms, including device makers and software publishers. The following analysis visualizes these connections, mapping the direct ties from the “Big Three”—Texas Instruments, Intel, and Qualcomm—to the rest of the supply chain.

### **Figure S2.6 Vertical Network and the Role of Texas Instruments**



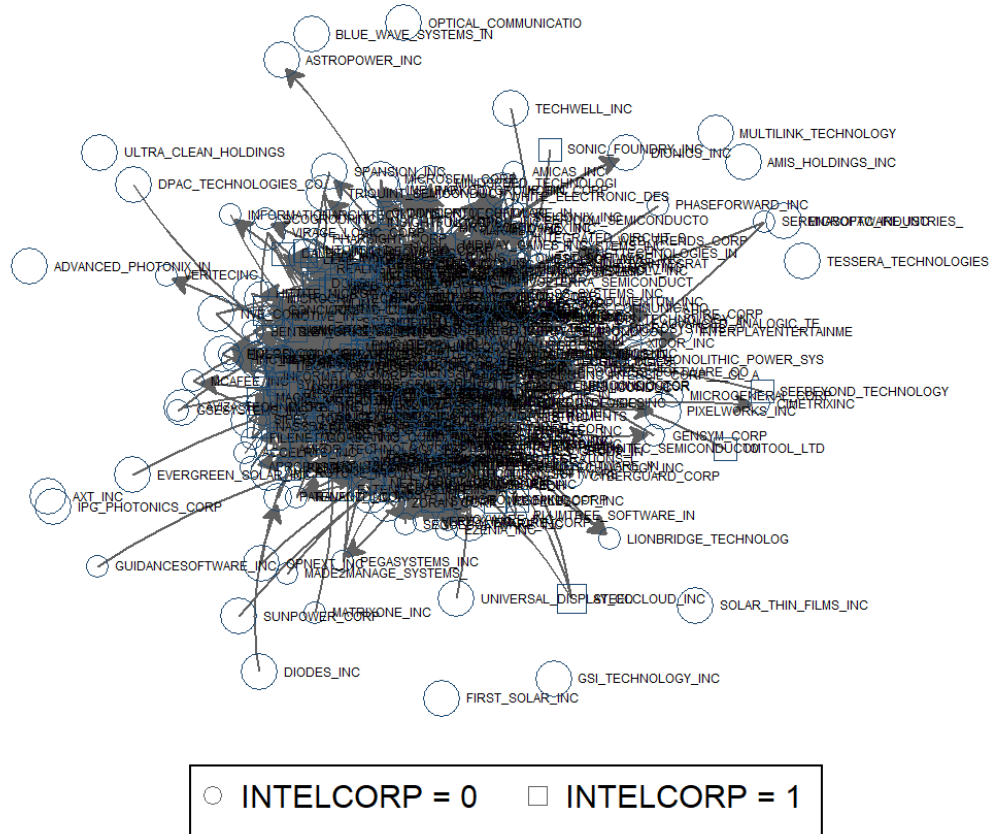
○ TEXASINSTRUMENTSINC = 0    □ TEXASINSTRUMENTSINC = 1

*Notes: This network visualization highlights Texas Instruments' connections within the vertical supply chain network of 289 firms, showing only the ties between the upstream semiconductor sector and downstream firms (hardware manufacturers, software publishers, and wireless telecommunications companies). The figure displays the 72 downstream firms whose patents cite Texas Instruments' patents between 1976 and 2006, demonstrating its role as a foundational technology provider across the digital device ecosystem. This vertical network explicitly excludes horizontal ties between firms within the same industry sector to focus exclusively on the flow of technology up and down the supply chain. Node size is proportional to the total number of connections each firm has within the broader vertical network.*  
*Sources: NBER Utility Patent and Patent Citation Data File (Hall et al. 2001); S&P Global Market Intelligence COMPUSTAT database (2025).*

Figure S2.6 reveals that Texas Instruments served as a foundational technology source for the entire downstream ecosystem. A total of 72 downstream firms cite its patents, a group that includes virtually every major technology titan of the era. The list of citing companies spans all

key sectors, from telecommunications giants like AT&T and hardware pioneers like Apple and Motorola to the dominant software firms of the day, including Microsoft, Oracle, and Adobe.<sup>109</sup>

**Figure S2.7 Vertical Network and the Role of Intel**



*Notes: This network visualization highlights Intel's connections within the vertical supply chain network of 289 firms, showing only the ties between the upstream semiconductor sector and downstream firms (hardware manufacturers, software publishers, and wireless telecommunications companies). The figure displays the 80 downstream firms whose patents cite Intel's patents between 1976 and 2006, confirming its central role as a foundational technology provider. This represents the largest downstream citation base among the "Big Three" semiconductor leaders. This vertical network explicitly excludes horizontal ties between firms within the same industry sector to focus exclusively on the flow of technology up and down the supply chain. Node size is proportional to the total number of connections each firm has within the broader vertical network.*

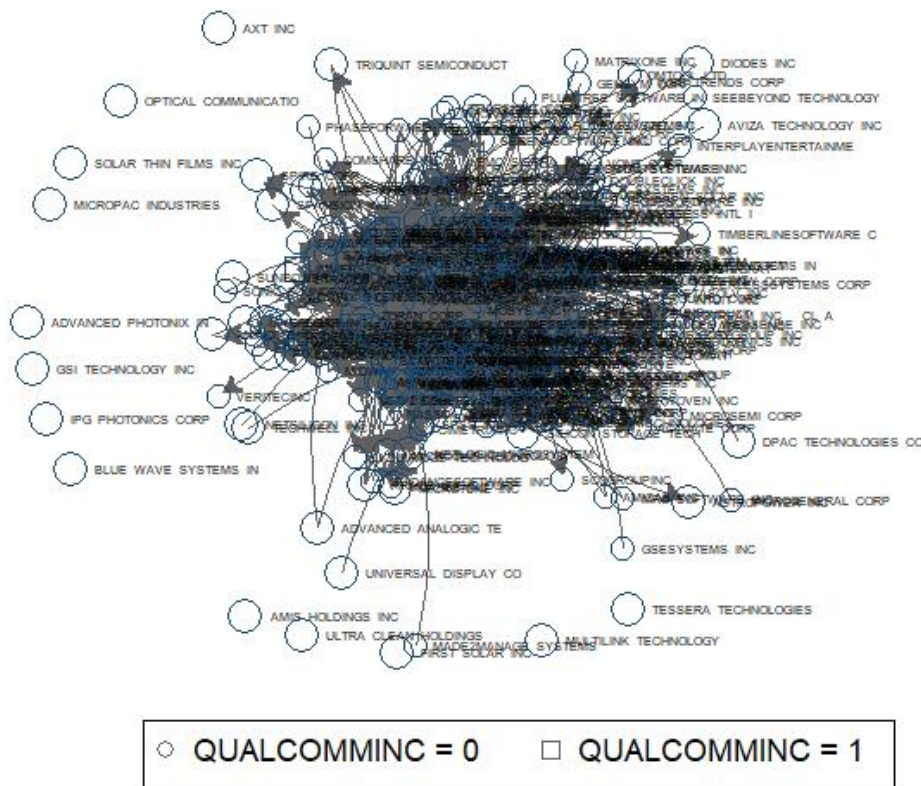
*Sources: NBER Utility Patent and Patent Citation Data File (Hall et al. 2001); S&P Global Market Intelligence COMPUSTAT database (2025).*

<sup>109</sup> The 72 citing firms are composed of 61 software publishers, 8 hardware manufacturers, and 3 wireless telecommunications firms. Other notable citing firms include hardware makers Compaq Computer Corp., Gateway Inc., and Palm Inc., and software firm Sun Microsystems Inc.

Figure S2.7 reveals that Intel was cited by an even larger group of 80 downstream firms, confirming its own central role as a foundational technology provider. The list of citing companies shows a high degree of overlap with those citing Texas Instruments, including nearly all the same industry titans as AT&T, Apple, Motorola, and Microsoft.<sup>110</sup> This demonstrates that most major downstream players relied on innovations from both top semiconductor hubs.

Finally, Figure S2.8 reveals Qualcomm's more specialized but still critical role in the supply chain. It was cited by a much smaller and more focused group of only 19 downstream firms. However, this select group included the same core industry titans that cited Texas Instruments and Intel, demonstrating that even with fewer connections, Qualcomm's foundational wireless technologies were essential to all the major players in the digital device ecosystem.

**Figure S2.8 Vertical Network and the Role of Qualcomm**



*Notes: This network visualization highlights Qualcomm's connections within the vertical supply chain network of 289 firms, showing only the ties between the upstream semiconductor sector and downstream firms (hardware manufacturers, software publishers, and wireless*

<sup>110</sup> Other notable firms that cited both Intel and Texas Instruments include Compaq Computer Corp., Gateway Inc., Palm Inc., Sun Microsystems Inc., Oracle Corp., Symantec Corp., and Adobe Systems Inc. The data also reveals firms with specific dependencies: 25 firms cited Intel but not Texas Instruments, while 17 firms cited Texas Instruments but not Intel.

telecommunications companies). The figure displays the 19 downstream firms whose patents cite Qualcomm's patents between 1976 and 2006, revealing its more specialized but still critical role in the supply chain. While Qualcomm's downstream citation base is smaller and more focused than Texas Instruments' or Intel's, this select group includes the same core industry titans, demonstrating that its foundational wireless technologies were essential to all major players in the digital device ecosystem. This vertical network explicitly excludes horizontal ties between firms within the same industry sector. Node size is proportional to the total number of connections each firm has within the broader vertical network. Sources: NBER Utility Patent and Patent Citation Data File (Hall et al. 2001); S&P Global Market Intelligence COMPUSTAT database (2025).

### **Diffusion between the Big Three Semiconductors and Apple and Motorola**

Now I turn attention to evaluating the links between the Big Three semiconductor firms and our two key downstream firms, Apple and Motorola. First, I look at patent citation patterns in which Apple cites patents held by either Texas Instruments, Intel, or Qualcomm. Second, I focus on Motorola's ties to either of these three semiconductor firms' patents. Third, I focus on ties that run in the opposite direction: ties directed *to* Apple and Motorola *from* the semiconductor firms.<sup>111</sup> This detailed, firm-level analysis will move beyond abstract network metrics to map the tangible pathways of knowledge transfer, revealing the specific dependencies and technological conversations between the most important firms at the heart of the Third Industrial Revolution.

#### *Apple citing Texas Instruments, Intel, Qualcomm*

Beginning with the downstream flow, the data shows that Apple's patents cite a total of 235 different Texas Instruments patents, demonstrating a broad reliance on its foundational technology.<sup>112</sup> The most cited of these, a highly general patent for a communication and information processing system, describes what is essentially a blueprint for a modern smart device.<sup>113</sup> This indicates that Apple was building upon some of Texas Instruments' most fundamental and broadly applicable innovations.

---

<sup>111</sup> I should also add that Apple and Motorola direct ties to each other. Apple patents cite Motorola patents 269 times. These patents have 38.7 average citations and an average generality score of .54. Motorola patents cite Apple patents 230 times. These patents have 44.8 average citations and an average generality score of .53. In terms of indirect ties, there are 60 pathways by which Motorola indirectly cites Apple and included within those 60 are Texas Instruments, Intel, and Qualcomm. Similarly, there are 48 pathways by which Apple indirectly cites Motorola; also included within those 48 are the Big Three semiconductor firms.

<sup>112</sup> Some Texas Instruments patents were cited by multiple Apple patents; for example, U.S. Patent 4,562,535 and U.S. Patent 4,688,195 were each cited by five different Apple patents. The most cited patent, U.S. Patent 5,465,401, received 174 total citations and has a generality score of 0.80.

<sup>113</sup> US Patent Number 5,465,401, issued by the USPTO in 1995, describes a "Communication system and methods for enhanced information transfer... This invention relates to communication and information storage and processing systems and more particularly to

The data shows an even deeper reliance on Intel, with Apple citing 467 different Intel patents. The most frequently cited of these was a landmark 1996 patent for a sophisticated power management system.<sup>114</sup> This technology was critical for the rise of portable, battery-operated computers. It acted as an intelligent coordinator that could automatically put different parts of a laptop into low power "sleep" states, dramatically extending battery life. This innovation was a foundational building block for the entire mobile computing revolution that would follow.

Apple's connection to Qualcomm, while smaller in scale, was highly strategic, focusing on the core wireless technologies that would define the mobile era. In total, Apple cited just seven Qualcomm patents.<sup>115</sup> However, the most important of which was a landmark 1999 patent for a novel, miniaturized antenna, a critical physical breakthrough that enabled the creation of the small, sleek handheld devices with reliable connections that were essential for the mobile revolution.<sup>116</sup> This shows that even with fewer direct citations, Apple was building upon Qualcomm's foundational innovations in wireless communication.

#### *Motorola Citing Texas Instruments, Intel, Qualcomm*

Turning to Motorola, the data reveals an even deeper and broader reliance on Texas Instruments' technology. In total, Motorola's patents cite a staggering 2,495 different Texas Instruments patents.<sup>117</sup> The most frequently cited of these, a highly general patent from 1994, describes a communication system with a touch-sensitive visual display—a clear precursor to the modern smartphone.<sup>118</sup> This demonstrates that Motorola, the leading handset maker of the era, was building its most advanced products directly on top of Texas Instrument's foundational innovations in processing and communications.

The data reveals a similarly deep connection to Intel, with Motorola's patents citing over one thousand different Intel patents.<sup>119</sup> This broad citation pattern shows that Motorola, in building its own portfolio of mobile communication technologies, consistently drew upon Intel's foundational innovations in microprocessor design and architecture.

Finally, in a striking illustration of this vertical knowledge flow, Motorola cites a staggering 1,241 Qualcomm patents, which possess a high average generality score of .68. The most cited of these is Qualcomm's foundational patent for CDMA technology (US Patent No. 5,103,459),

---

communication devices with processors, expanded memory capability, and visual displays for enhanced communication and information storage and processing" (Earle 1995).

<sup>114</sup> This is US Patent No. 5,560,022 (Dunstan et al. 1996).

<sup>115</sup> The seven Qualcomm patents cited by Apple have an average generality score of 0.65. The most cited, U.S. Patent 6,008,762, received 73 total citations and was cited by three different Apple patents.

<sup>116</sup> This is U.S. Patent No. 6,008,762 (see Nghiem 1999).

<sup>117</sup> The most cited Texas Instruments patent by Motorola, U.S. Patent 5,335,276, received 292 total citations and has an extremely high generality score of 0.89.

<sup>118</sup> This is U.S. Patent No. 5,335,276 (Earle and Birdwell 1994).

<sup>119</sup> In total, Motorola cites 1,094 Intel patents. These patents had an average of 46.5 citations each, with the most-cited patent having been referenced 241 times.

which I've already discussed above, underscoring the deep reliance of downstream device manufacturers on the core innovations of their upstream semiconductor partners.

### **Evaluating Trickle Up Technological Diffusion**

It's important to emphasize that the flow of innovation was not a simple one-way street from chipmakers to device manufacturers. Leading semiconductor firms were not just passively receiving information, but actively integrating distinct, best-in-class solutions from downstream innovators to improve their own products. Prominent examples I outlined below include power efficiency enhancing technologies from Apple and hardware-level security from Motorola. Along these lines, Qualcomm was laser-focused on absorbing foundational technologies for secure mobile communications, an area where Motorola, not Apple, was the undisputed pioneer

As illustrated by Texas Instruments' citation patterns, the story of "trickle-up" technological diffusion is quite complex. While the company cites only a single Apple patent, it references Motorola's portfolio an astonishing 3,137 times. The most cited of these Motorola patents is a highly influential invention for an unforgeable personal identification system, US Patent No. 4,993,068 (Piosenka and Chandos 1991). This particular focus suggests that upstream semiconductor firms were actively integrating security and authentication technologies pioneered by downstream hardware specialists, incorporating them directly into their own product roadmaps.

Intel's citation patterns reveal an even more nuanced story of trickle-up innovation, showing how it strategically absorbed key technologies from different downstream partners. The company highly referenced Apple's portfolio, citing its patents 924 times, with a particular focus on power management. Most tellingly, Intel repeatedly cited Apple's groundbreaking patent for conserving battery life in laptops, US Patent No. 5,167,024 (Smith et al. 1992)—a critical technology for developing the energy-efficient microprocessors needed for the mobile computing revolution. Simultaneously, Intel cited Motorola even more copiously—3,906 times—again focusing on the same foundational patent for an unforgeable personal identification system (US Patent No. 4,993,068) outlined in the previous paragraph.

Qualcomm's citation behavior, the last of the Big Three, sharply reflects its specialized focus on wireless communication. It shows minimal interest in Apple's portfolio, citing its patents only 19 times, with most references pointing again to the patent on laptop power management (US Patent No. 5,167,024)—a technology less central to Qualcomm's mobile-first mission. In stark contrast, Qualcomm cited Motorola's patents 2,277 times, delving deep into its mobile communication innovations. The most referenced of these, which was cited by 16 different Qualcomm patents, was a critical patent for a cellular privacy device, US Patent No. 4,785,463 (Janc and Jasper 1988), boasts a relatively high generality score of .78.<sup>120</sup>

### **Semiconductors to Device Manufacturers as Mediators**

---

<sup>120</sup> This patent, which describes a "digital global positioning system receiver," is a blockbuster invention that introduced a radio receiver adapted for use with the Global Positioning System (GPS) navigation system using digital circuitry.

The data also reveals that both upstream semiconductor firms and downstream device manufacturers served as essential intermediaries in the pre-Fourth Industrial Revolution high-tech supply chain. This shared role in promoting technological diffusion is evident in their strikingly similar network profiles: both groups were cited by a common audience of innovators from across the digital ecosystem, which drew on a body of bedrock patents. Such significant overlap indicates that a set of foundational technologies permeated the industry, regardless of where they originated.

I can assess this structural overlap using the same formal network analysis I applied earlier to the semiconductor leaders. Just as I measured the similarity of the "Big Three" by comparing their shared audiences, I can do again so here by identifying the correlation between the firms citing device makers versus those citing chipmakers. More formally, to find out if two nodes are connected to the same alters, I calculate how correlated the tie vectors of two nodes are with each other. Specifically, I identify the correlation of the vector of incoming ties of node  $i$  and the vector of incoming ties of node  $j$ .

This analysis reveals a dense web of indirect connections and a shared audience for key players, confirming that they were central to the same broad technological conversation. The dialogue was strongest between Intel and Apple, which shared an audience correlation of .32—four times the network average of .08. This relationship was clearly reciprocal; in the reverse direction, Intel cited 22 companies that also cited Apple. The other semiconductor leaders followed this pattern, with Texas Instruments sharing a notable audience with Apple ( $r = .13$ ) while citing 10 of its partners, and Qualcomm ( $r = .09$ ) citing six. That this shared audience included titans like Microsoft, Oracle, and Adobe confirms that the industry's most important players were closely watching innovations from both the downstream device makers and the upstream chip designers, cementing their joint status at the heart of the digital supply chain.

The same pattern of dense, reciprocal connections is equally evident with Motorola, confirming this was an ecosystem-wide structure. Dozens of indirect pathways linked Motorola to the Big Three, and crucially, the common intermediaries were the very same cohort of software and platform giants: Microsoft, Oracle, and Symantec. This consistent presence of a core group of software firms acting as a bridge for both Apple and Motorola demonstrates their unique and essential role in this "innovation commons." The reciprocal dialogue also held true, with firms like Intel citing 26 companies that were simultaneously partners with Motorola. Ultimately, the fact that the same powerful players orbited both leading device manufacturers proves that this was a single, tightly integrated network, not a series of disconnected supply chains.

### **A More Complex Dynamic: Industry Leaders as Curators**

Another question is whether industry leaders act not just as creators of their own technologies, but as curators and conduits for innovations originating *outside* their own sector. In such a scenario, a pacesetter firm might supercharge diffusion by importing a cutting-edge process from another industry and then distributing it to its own followers. To address this issue, I return to the cross-sectional patent dataset and again compare firms operating in high-tech industries versus firms outside of those sectors.

Although one might expect high-tech leaders to be particularly open to external GPTs, the patent data reveals a surprising and more nuanced reality—the "curator" role is far more common in non-high-tech sectors. While the average generality of backward citations for technological followers is similar across sectors (.35 for high-tech vs. .38 for non-high tech), the difference for leaders is stark. The average for high-tech leaders is only .39, while for non-high-tech leaders it is .56. Moreover, non-high-tech sectors with leaders scoring above .85 include Asphalt Shingle Manufacturing, Synthetic Dye Manufacturing, and Light Truck Manufacturing, indicating a high degree of openness to external technologies: leaders in industries like chemical and vehicle manufacturing are significantly more open to citing patents from outside their own fields. Therefore, while the high-tech innovation commons proved incredibly powerful at generating and diffusing its *own* autochthonous technologies, its leaders were surprisingly insular compared to those in more traditional industries, who had to look further afield for their next breakthrough.

This seemingly surprising finding offers a crucial insight into the very nature of superstar economics. The relative insularity of high-tech leaders isn't a weakness; it's a defining characteristic of their dominance during the latter part of the Third Industrial Revolution. These firms weren't just participating in the economy; they were creating a new one. Rather than importing best practices from other sectors, they *were* the best practice.

## CONCLUSION

Chapter 5 of the book charted a profound economic transformation, tracing the arc from the postwar era of renewed globalization to the rise of a U.S. economy built not on factory floors but on intangible capital. The international economic architecture championed by the United States—defined by the dollar's hegemony, open capital flows, and expanding trade networks—reshaped global commerce. It catalyzed the vertical disintegration of supply chains, allowing American firms to specialize in high-value, IP-intensive activities like design, R&D, and software development. This structural shift, in turn, paved the way for the quintessentially American phenomenon of the "superstar firm," which leveraged scalable intangible assets and superior technology to bolster productivity and capture outsized market shares and profits.

The statistical evidence presented in this section of the appendix revealed that this transformation was neither incremental nor evenly distributed. Firm-level analyses demonstrated that technological influence and productivity became increasingly concentrated among a small cohort of "innovation leaders." High-tech sectors, in particular, exhibited a Pareto-like "winner-take-most" pattern: the top 20 percent of firms accounted for roughly 98 percent of all forward patent citations, and firms at the technological frontier displayed total factor productivity levels about 70 percent higher than their sectoral peers. Patent distributions followed a Power Law dynamic, with Gini coefficients exceeding 0.9 in high-tech industries—nearly double those observed in non-high-tech sectors—underscoring the extreme concentration of inventive capacity. Network analysis further confirmed that semiconductor giants such as Texas Instruments, Intel, and Qualcomm emerged as the central hubs through which general-purpose technologies diffused both vertically across supply chains and horizontally to rival firms.

Yet this asymmetry did not signal technological stagnation or monopolization. The networks revealed a dense but open "innovation commons," where a handful of pacesetter firms

generated foundational technologies that others refined, implemented, and built upon. Knowledge flowed efficiently through short paths linking hub firms to the periphery, allowing followers to adopt and adapt the leaders' general-purpose technologies.

At the same time, the final empirical section showed that these high-tech leaders, while extraordinarily effective at propagating their own innovations, were relatively insular in adopting ideas from outside their sectors—unlike non-high-tech firms, which remained far more open to external technological borrowing. This insularity was not a weakness but a defining feature of their dominance: rather than importing best practices from elsewhere, they became the best practice, setting the technological standards and codifying the processes that underpinned the transition to a digital, intangible economy.

Ultimately, the macroeconomic and network-level evidence tell a unified story. While the global liberal order provided the fertile soil for superstar firms to grow, the patent citation networks reveal the root system through which those firms exchanged and diffused technological knowledge. Within this ecosystem, a technology trendsetter could effectively standardize and transmit general-purpose technologies to its peers, creating a decentralized yet highly productive innovation commons. As I show in Chapter 7 of the book, preserving this technological bedrock required a decisive break from the “populist-statist” past (see Chapter 3 of the book), as the legal system evolved to tolerate the market power of these central hubs in exchange for the massive efficiencies and knowledge transfers they generated.

Yet, as powerful as this decentralized mechanism proved, it had its limits. While it was remarkably effective at propagating foundational technologies within industries, it could not, on its own, provide the explicit, guaranteed, global interoperability required for the next great leap: the truly mobile internet. That challenge demanded a different model of collaboration—one based not just on the gravitational pull of market leaders but on formalized cooperation. Chapter 6 of the book explores precisely this evolution through the case of Qualcomm and the development of 4G mobile technology, where the creation of Standards-Setting Organizations (SSOs) and the licensing of essential patents on FRAND terms became indispensable tools for building the globally interconnected world that would give rise to the smartphone—and with it, the data-rich foundation of the Fourth Industrial Revolution.

## ONLINE SUPPLEMENTARY APPENDIX TO HISTORY'S MOST REVOLUTIONARY INNOVATION, SECTION 3

During the early 2020s, several algorithms at the practical edge of AI innovation could process, interpret, and act upon the vast datasets made possible by the digital platforms I investigated in Chapter 8 of the book, exercising bounded autonomy in decision-making systems ranging from content moderation to logistics and finance. Companies such as OpenAI and Google developed and deployed machine learning to enable computers to act intelligently without explicit programming. Instead of following predetermined, hardcoded rules, machine learning systems learned statistical patterns from data and adapted their outputs as they were trained and fine-tuned, marking a significant shift from traditional rule-based programming (see Russell and Norvig 2021). This nimbleness and flexibility explain why machine learning underpinned many of AI's advancements during this era (see OpenAI 2023; Google 2023).

It also explains why computer scientists and engineers were treated like superstar athletes—pursued by competing AI labs in an unprecedented bidding war. By mid-2025, multiple outlets reported nine-figure, multi-year packages for a very small number of frontier researchers.<sup>121</sup> Meta's push to build a “superintelligence” lab coincided with aggressive, multi-year offers and strategic dealmaking; notably, the company agreed to acquire a 49% stake in Scale AI for about \$14.3 billion, a move widely read as a bid to secure data-labeling capacity and talent (Weinberg 2025).

The competition for talent quickly escalated beyond just hiring individuals. In 2025, Google licensed the agentic-coding startup Windsurf while bringing over CEO Varun Mohan and staff in a deal reported at \$2.4 billion (Cai et al. 2025). Amazon had executed a similar maneuver a year earlier—hiring Adept's co-founders and dozens of researchers while licensing Adept's technology into its internal foundation-model group (Wiggers 2024a). Microsoft, for its part, circulated an internal “shopping list” of target Meta researchers and approved multimillion-dollar packages to poach them (Eriksson 2025). Apple simultaneously recruited LLM specialists from rivals—especially Google—as part of its systemwide AI push (Acton 2024). On the defensive side of the market, OpenAI rolled out retention bonuses worth several million dollars apiece and refreshed employee equity to reduce attrition risk following high-profile departures (Tong and Cai 2025).

The bottom line is that as these systems crossed key performance thresholds, the constraint shifted from ideas to execution—data, compute, and the people who could turn both into working models.

But this fevered competition for human capital also reflects a deeper truth about AI's economic significance. The willingness of corporations to pay effectively unlimited sums for scarce expertise signals that something more consequential than ordinary technological improvement was underway.

---

<sup>121</sup> To be sure, these offers were typically dominated by equity grants rather than cash sign on bonuses; viral claims of \$100 million “signing bonuses” overstated the cash component of offers (Bort and Temkin 2025).

## The AI Talent Wars Underscore AI’s Revolutionary Appeal

As I argued in Section 1 of this appendix, industrial revolutions occur when new general-purpose technologies (GPTs) herald a dramatic discontinuity rather than the seamless continuation of technological progress engendered by their broad applicability across sectors, uninterrupted technical improvement over extended periods, and strong complementarities with co-inventions that amplify their impact. In this manner they fundamentally transform both the industrial organization of production across multiple sectors simultaneously and the economy and society in general (see Bresnahan and Trajtenberg 1995).

By 2025, AI was driving a Fourth Industrial Revolution. In healthcare, AI-driven precision medicine was revolutionizing patient care by tailoring treatments based on individual genetic profiles. In manufacturing, the integration of AI with industrial automation was underwriting more efficient and flexible production lines. In finance, advanced algorithmic trading techniques were optimizing strategies and improving market liquidity. In education, AI-powered adaptive learning systems provided personalized experiences tailored to individual students’ needs.

This transformation rose out of the innovations in hardware, software, and business processes outlined in the book—themselves byproducts of what I termed in Chapter 3 of the book the “Creative Destruction Paradigm,” an institutional and policymaking framework that prioritized first-principles economic reasoning and cost-benefit analysis over other goals, such as protecting incumbents or redistributing to politically sensitive constituencies.

The Fourth Industrial Revolution was being powered by an unprecedented capital explosion, which I explore in the book, powered by an audacious moonshot that transcended traditional business metrics: the pursuit of Artificial General Intelligence (AGI)—AI systems capable of matching or exceeding human cognitive abilities across the full range of intellectual tasks, rather than excelling only in narrow, pre-specified domains. Company executives increasingly invoked AGI as the ultimate justification for infrastructure spending that defied conventional return-on-investment calculations, arguing that the transformative potential of human-level artificial intelligence warranted essentially unlimited capital deployment (Bobrowsky 2025). This AGI-centric rationale marked a departure from previous technology investment cycles, which typically anchored spending to specific product roadmaps and revenue projections.

This section of the appendix tells the story of why modern AI systems were able to evolve into a unique supply chain to begin with. I recount the historical process by which advances in both AI theorizing and practical computer science enabled foundation models trained on massive corpora to scale and undergo refinements in kind. This scaffolding supports my examination of AI as a maturing, capital-intensive stack in Chapter 9 of the book: where I explore how the distribution of power and value turn on who controls the scarce complements needed to deploy it at scale—rather than who conjures up a standalone model out of whole cloth.

This section of the appendix proceeds in five steps. I begin with a brief technical primer on modern AI—machine learning, neural networks, and the mechanics of deep learning—before turning to large language models and explaining how transformers convert text into probabilistic next-word predictions, and what that does (and does not) imply about “understanding.” I then

place these systems in historical perspective, tracing the breakthroughs that culminated in foundation models and the post-2022 LLM wave. Next, I show how these advances translated into an explosion of training compute and capital intensity, shifting the frontier from academia to a small number of well-capitalized labs. With that context in place, I explain the 2024–25 pivot toward reasoning-oriented systems—post-training, test-time compute, and tool-augmented architectures—and why they improved performance on multi-step tasks while still exhibiting sharp limits under distribution shift and adversarial conditions. I close by outlining the emerging technical and institutional paths forward, including efficiency-focused architectures, verification, and brain- and hardware-inspired research agendas.

## THE TECHNICAL NUTS AND BOLTS OF AI

There are three primary approaches to machine learning, each defined by how it processes data. Supervised learning relies on data labeled by researchers (Goodfellow et al. 2016). Unsupervised learning discovers patterns in unlabeled data by itself (LeCun et al. 2015). Reinforcement learning allows researchers to train models through rewards and penalties (Sutton and Barto 2018). While all three approaches matter for modern AI, supervised learning provides the clearest entry point for understanding neural networks, which I turn to next. I return to reinforcement learning later in this section of the appendix when discussing RLHF and reasoning models.

In supervised learning algorithms, each input is paired with a known correct output—the “ground truth” against which the algorithm checks its predictions. This pairing enables the algorithm to learn patterns and make accurate predictions or classifications on new, unseen data. By iteratively adjusting their internal parameters to minimize errors, supervised learning models become highly effective at tasks like speech recognition, sentiment analysis, and predictive maintenance in manufacturing (Jurafsky and Martin 2023).

In this way, machine learning leverages structured data to address specific real-world problems. This includes fraud detection in banking, where algorithms identify unusual patterns in transaction data (Ngai et al. 2011), and medical diagnosis, where systems analyze patient records to flag potential health risks (Esteva et al. 2019). Building upon the principles of supervised learning, neural networks introduce a layered architecture that enables the modeling of far more complex patterns (LeCun et al. 2015).

### Neural Networks

Neural networks represent machine learning’s most transformative development. Consisting of interconnected artificial “neurons” arranged in layers that simulate the human brain’s architecture, these networks are trained on massive datasets, allowing computers to classify and cluster data with remarkable accuracy (see *ibid*). They learn to recognize patterns by adjusting the connections between neurons based on experience, enabling them to tackle complex tasks ranging from image recognition to generating human-like text (Goodfellow et al. 2016).

Despite their extraordinary capabilities, neural networks remain fragile. Consider that researchers demonstrated how small modifications to pixels in an image—invisible to the human eye—can mislead neural networks into misidentifying objects entirely (Szegedy et al. 2013; Goodfellow et

al. 2014). For example, a stop sign with carefully crafted perturbations might be classified as a speed limit sign, with obvious implications for autonomous vehicle safety (see Eykholt et al. 2018).

Said brittleness reveals a core limitation, which I will investigate further ahead in this section of the appendix: these systems rely on statistical correlations rather than true conceptual understanding (Marcus and Davis 2019). That same limitation helps explain why generative AI is sometimes unreliable, prone to “hallucinating” facts—generating plausible-sounding but fabricated information—or exhibiting misalignment between a user's goals and the machine's responses. To understand why these limitations persist despite massive investments in AI development, we must first examine how these systems actually function.

### *How Neural Networks Work*

Data moves through a neural network in a process called forward propagation across three types of layers.

The input layer serves as the network's entry point. Each neuron in this layer represents one feature or characteristic of the data; for example, if analyzing houses, neurons might represent square footage, number of bedrooms, or location. This layer performs no calculations, simply receiving and passing along raw data.

The hidden layers are the "thinking" engines where the actual processing occurs. They are called "hidden" because their operations are not directly visible to the user. Each neuron here receives inputs from the previous layer, multiplies each by a weight (representing the connection's importance), and sums them. Crucially, it then applies an activation function to introduce non-linearity before passing the result to the next layer.<sup>122</sup>

To put the importance of non-linearity in perspective, think of trying to map temperature zones across a landscape using only flat planes. With linear functions, you would have to approximate mountain ranges and valleys with crude diagonal slices, missing many local variations. But with non-linear functions, you can model the actual contours of the terrain, capturing how temperature truly varies with elevation and geography. In the same way, activation functions allow neural networks to fit the complex, “contoured” reality of real-world data, rather than forcing it into rigid, linear boxes.

Neural networks often employ multiple hidden layers to detect increasingly abstract patterns. A subsequent hidden layer combines the outputs of a preceding one, enabling the network to learn sophisticated features (LeCun et al. 1998). In image recognition, for instance, the first layer

---

<sup>122</sup> An activation function is a simple mathematical rule applied to the output of each neuron to decide whether it should "fire." A key breakthrough in deep learning was the widespread adoption of the Rectified Linear Unit (ReLU). The ReLU function is a tiny "if-then" rule: if the input is negative, it outputs 0; if positive, it outputs the number itself. This simple "kink" at zero provides the non-linear element that, when repeated across millions of neurons, allows the network to model complex, "contoured" patterns (Nair and Hinton 2010).

might detect simple edges, while the second identifies shapes or objects (Krizhevsky et al. 2012). Each node is connected to every node in the previous layer, with each connection carrying a specific “weight” that determines the signal's strength (see Rumelhart et al. 1986).<sup>123</sup>

The output layer produces the final predictions or classifications (Goodfellow et al. 2016). The number of neurons here depends on the task: classification tasks (like identifying objects) use one neuron per possible category, while regression tasks (like predicting prices) might use just one neuron outputting a continuous value. The values of these nodes represent the network's confidence, calculated by converting raw output scores—called 'logits'—into probabilities using a normalization function.<sup>124</sup>


Before I explain how these networks allow “learning” to occur, it is helpful to visualize their basic structure. Figure S3.1 illustrates a generic feedforward neural network, the most common architecture. The network is organized into distinct layers of processing units called "neurons" or nodes. Data enters through the input layer, which receives raw features—such as pixel values from an image or numerical data from a spreadsheet—and passes them to the next stage without computation. The core processing happens in one or more hidden layers. Each neuron in a hidden layer receives signals from the previous layer, performs a calculation, and transmits a new signal forward. Finally, the output layer produces the network's final prediction or classification, such as identifying an object or forecasting a value. The arrows connecting the layers represent the flow of data, and each connection carries a "weight", a numerical value adjusted during training that determines the strength and influence of that connection; more on that below.

### **Figure S3.1 A Generic Feedforward Neural Network Architecture**

---

<sup>123</sup> These weights correspond to billions—or even trillions—of internal connections and are stored as matrices (grids of numbers). During training, the system must constantly multiply these matrices together. This is where the Multiply-Accumulate (MAC) operation becomes critical. A MAC operation calculates the "dot product" of a row and a column: it multiplies corresponding pairs of numbers and then adds them all up to get a single value. In mathematical terms, for a single neuron, this is expressed as:  $y = \Sigma(w_i \cdot x_i)$ , where  $w$  is the weight and  $x$  is the input. Because neural networks contain matrices with thousands of rows and columns, this simple multiplication-and-addition step must be performed quadrillions of times, making the efficiency of the underlying MAC hardware the primary bottleneck for AI speed. This is a topic I will explore in the next chapter.

<sup>124</sup> Normalization is achieved by applying the “softmax function”—a mathematical formula that converts a vector of raw scores, or logits, into a probability distribution. The formal equation is:  $\text{softmax}(z_j) = e^{z_j} / \Sigma e^{z_j}$  for all classes  $j$  (see Bridle 1990).



*Notes: This diagram shows the flow of information from left to right across three types of layers. The input layer (represented by  $x$ ) receives raw data. The hidden layer (represented by  $h^{(1)}$ ) performs intermediate calculations, extracting features from the input. The output layer (represented by  $y$ ) generates the final prediction. The vertical dots ( $:$ ) indicate that the number of neurons in each layer can vary depending on the specific application ( $n$ ,  $m$ , and  $k$  represent these variable counts).*

### *How Neural Networks Learn*

Neural networks learn much like humans refine a skill: through practice, error, and feedback. Consider how a child learns not to touch a hot stove. Through the painful feedback of a burn, the brain's biological neurons adjust their connections, strengthening the pathways that associate the glowing red coil with danger. This process of inductive learning—generalizing rules from specific examples of trial and error—is rooted in the biological principle that synaptic connections strengthen when activated together. As psychologist Donald Hebb (1949) summarized it: neurons that “fire together, wire together.” Artificial neural networks replicate this “strengthening” process mathematically.

The process begins with a forward pass, where input data is fed through the network to produce an initial prediction. This prediction is compared against the true answer to calculate the error. Armed with this error calculation, the network performs a backward pass, or backpropagation. This algorithm sends corrective feedback in reverse through the layers. The non-linear functions are vital here, as they enable the computation of necessary adjustments (gradients) to the

network's internal weights. Each weight is adjusted slightly—with larger errors prompting bigger tweaks—to minimize future prediction errors (see Rumelhart et al. 1986).<sup>125</sup> It is like tuning thousands of knobs simultaneously until the system produces accurate results.

This entire cycle—forward pass, error calculation, backpropagation, and weight updates—repeats millions of times. Gradually, as the network "tunes" these thousands of internal knobs, it learns to recognize underlying patterns in the data, enabling it to make accurate predictions on new, unseen inputs.

Figure S3.2 visualizes the complete flow of information through a neural network, using a simplified model based on the landmark LeNet-5 architecture (LeCun et al. 1998). This system was built to recognize digits from the MNIST database, a collection of handwritten digits formatted as 28×28-pixel images. The input layer consists of 784 neurons, corresponding to the total pixel count of a single image. Each pixel maps to one input neuron, holding a value from 0 (white) to 1 (black); together, these 784 values represent the complete digital image.

The network's first hidden layer focuses on edge detection, learning to identify basic structural features like lines, curves, and corners (see Zeiler and Fergus 2014). Some neurons might activate strongly when they detect a vertical line, while others respond to curves. For instance, a written "7" would activate neurons that detect a horizontal line at the top and a diagonal line going down.

The second hidden layer combines these simple features into more complex patterns. It learns to recognize geometric combinations: a horizontal line connected to a diagonal line (characteristic of a 7), a closed loop (the circular form found in 0, 6, 8, and 9), or a vertical line with serifs (characteristic of a 1). These neurons are essentially detecting digit-specific conceptual shapes.

Finally, the output layer contains 10 neurons, each representing one possible digit from 0 through 9. Each neuron produces a score; a softmax function converts these scores into a probability distribution that sums to 1.

Figure S3.2 traces the process in action when a user writes the number "7." First, the 28×28 pixel image is converted into 784 numerical values and fed into the input layer. During feature

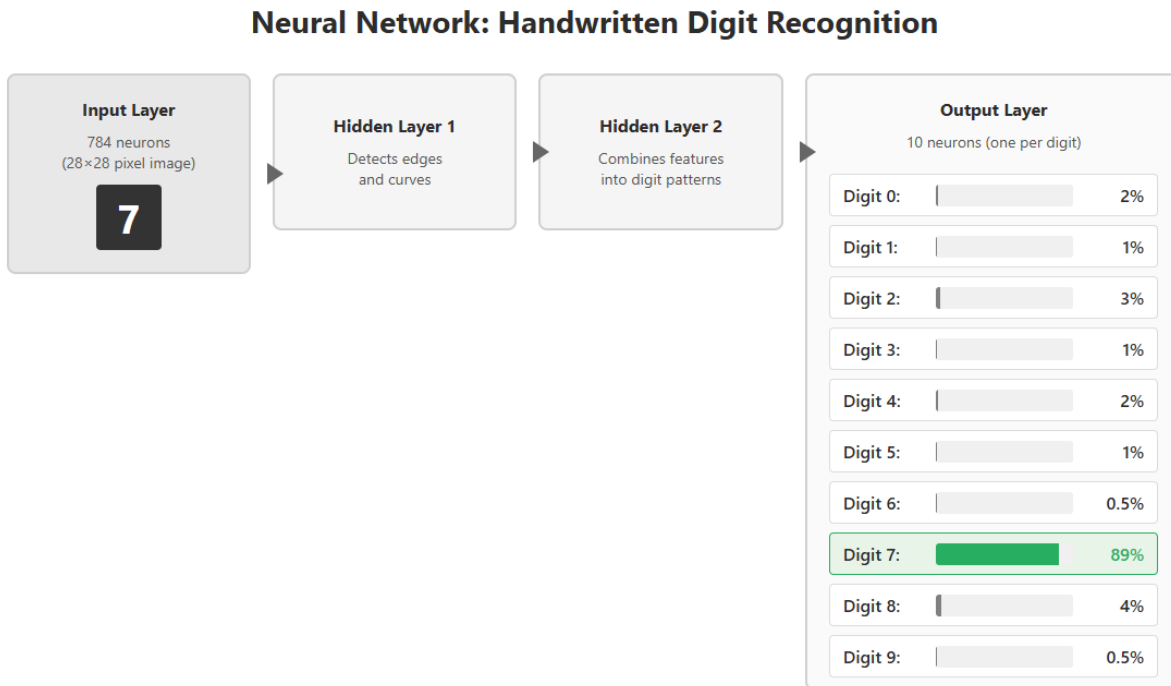
---

<sup>125</sup> Technically, backpropagation uses the chain rule from calculus to compute the gradient for every weight—a vector indicating the slope of the error curve with respect to that specific weight. This feeds into the optimization algorithm, which is the mathematical rule set determining how to update the network's parameters. The most common strategy is gradient descent. Conceptually, imagine a hiker blindfolded on a rugged mountain (the "loss landscape") trying to reach the lowest valley. The loss represents the hiker's current altitude (the magnitude of error). The gradient tells the hiker the steepness and direction of the slope directly under their feet. The optimization algorithm directs the hiker to take a step in the opposite direction of the incline—moving "downhill" to progressively minimize loss (reduce the discrepancy between the model's prediction and the correct answer). For a comprehensive overview of these dynamics, see Ruder (2016); for the seminal introduction of backpropagation, see Rumelhart et al. (1986).

extraction, Hidden Layer 1 detects the horizontal line at the top and the diagonal line, while Hidden Layer 2 recognizes this specific combination as characteristic of a "7."

The output layer then converts its scores into probabilities; for illustration, Digit 0 might receive about 2% (unlikely due to the lack of closed loops), Digit 1 about 1%, and Digit 2 about 3%. Digit 7, however, might receive roughly 89% probability—a strong match with the detected features—while Digit 8 might get about 4%. Consequently, in this example, the network selects “7” as its final prediction.

**Figure S3.2 A Neural Network for Handwriting Recognition**



*Notes: The network processes the input image through multiple layers, extracting increasingly complex features. The output layer contains 10 neurons, each representing a digit from 0 to 9. Each neuron outputs a confidence score (probability) that the input image is that particular digit. The network predicts the digit with the highest probability—in this case, it is 89% confident the handwritten input is a "7."*

*Source: LeCun et al. 1998.*

### **Explaining Deep Learning and its Practical Applications**

The architecture of stacked layers and nonlinear activation functions shown in Figure S3.1 is the foundation for deep learning, an advanced subset of neural networks that employs numerous hidden layers to capture complex and abstract features from raw data. Deep learning processes vast amounts of data and identifies patterns with extraordinary speed and precision, enabling breakthroughs in fields like natural language processing and computer vision.

Building on the basic architecture illustrated in Figure S3.1, a deep learning model contains a significantly larger number of hidden layers, allowing the network to extract increasingly abstract representations of data and perform complex tasks with unparalleled accuracy. Each of these hidden layers is designed to process and refine information hierarchically (Goodfellow et al. 2016). For instance, in image recognition, the early layers might identify simple patterns like edges or corners, while deeper layers detect more abstract concepts such as textures, shapes, and eventually entire objects (LeCun et al. 2015).

Additionally, deep learning architectures often incorporate specialized layer types tailored to the specific problem domain. For example, Convolutional Neural Networks (CNNs) handle spatial data like images by using convolutional layers to extract local features (Krizhevsky et al. 2012). Similarly, Recurrent Neural Networks (RNNs) process sequential data such as time series or text using feedback loops to capture dependencies over time (Graves 2012).

While the foundational principles of weight adjustments and error minimization remain consistent (Rumelhart et al. 1986), deep learning's expanded depth and complexity allow it to solve a range of real-world problems that were impossible for simpler networks. This powers many of the "smart" features now taken for granted, such as Amazon's Alexa, which recognizes voices and responds intelligently (Goodfellow et al. 2016). While Google's search engine historically relied on advanced statistical models, more recent developments—such as the RankBrain and BERT systems—employ neural networks to better understand context, semantics, and user intent (Vaswani et al. 2017; Devlin et al. 2018).

Google also employs neural networks to detect faces in photos—a feature embedded in its photo organization tools (Goodfellow et al. 2016). Autofill capabilities in text messaging and email rely on similar algorithms to predict user input (ibid). In hardware innovation, Google's Tensor chip integrates deep learning directly into smartphones, enabling devices like the Pixel to adapt dynamically to user preferences (Google 2023).

### *The LLM in Action: A Deeper Look*

OpenAI's ChatGPT represents perhaps the most quintessential example of modern AI capabilities. Large Language Models (LLMs) like GPT function by using deep learning techniques to analyze patterns in massive text datasets, discerning statistical relationships between words and phrases (Vaswani et al. 2017). Through a process called "attention," these models weigh the relative importance of different words in context. When processing the sentence "The animal didn't cross the street because it was too tired," the attention mechanism helps the model recognize that "it" refers to "animal" rather than "street"—a classic test of ambiguity resolution (Winograd 1972). This enables the generation of coherent and contextually appropriate responses.

To understand how a generative AI model like ChatGPT works, we can walk through the process, which is also illustrated in Figure S3.3. The entire system is an elegant, multi-stage pipeline that converts human language into mathematics and back into human language.

## The Embedding Layer: Converting Words to Meaning

The process begins with the embedding layer, which acts as a sophisticated translator that converts words into a language the neural network can understand—numbers. But it is far more than a simple dictionary. It takes a single word (or "token") and converts it into a large vector—a list of hundreds or even thousands of numbers. The magic of this layer is that these vectors represent the word's meaning and context. In this high-dimensional space, words with similar meanings end up with similar vectors. For example, the vectors for "cat" and "dog" will be close together, while the vectors for "cat" and "democracy" will be very far apart. This "understanding" is not pre-programmed; it is learned during training, as the network discovers that words used in similar contexts (e.g., "you can pet a \_\_\_") should have similar vectors.

## Transformer Layers: The Context Engine

Once the text is converted into these meaningful vectors, it is fed into the transformer layers. This is where the real "intelligence" of the model resides. Each transformer layer performs self-attention, the mechanism I introduced above. Unlike earlier models that read text strictly from left to right and often "forgot" early words—more on that below—this process allows the model to look at every word in a sentence and weigh its relationship to every other word simultaneously. For instance, when processing the sentence "The bank by the river was steep," the self-attention mechanism analyzes the word "bank" and sees its strong connection to "river" and "steep," allowing it to conclude that it means a riverbank, not a financial institution.<sup>126</sup>

These transformer layers are stacked, often dozens or even hundreds deep. This depth is crucial, as each layer builds a more sophisticated level of comprehension. Early layers might detect basic grammar, middle layers understand phrases, and deep layers grasp abstract reasoning. This hierarchy allows the model to resolve linguistic ambiguities—correctly identifying that in the sentence "The cat... living in Boston... it is sleeping," the word "it" refers to "the cat" many words prior.

## The Final Output: Generating a Response

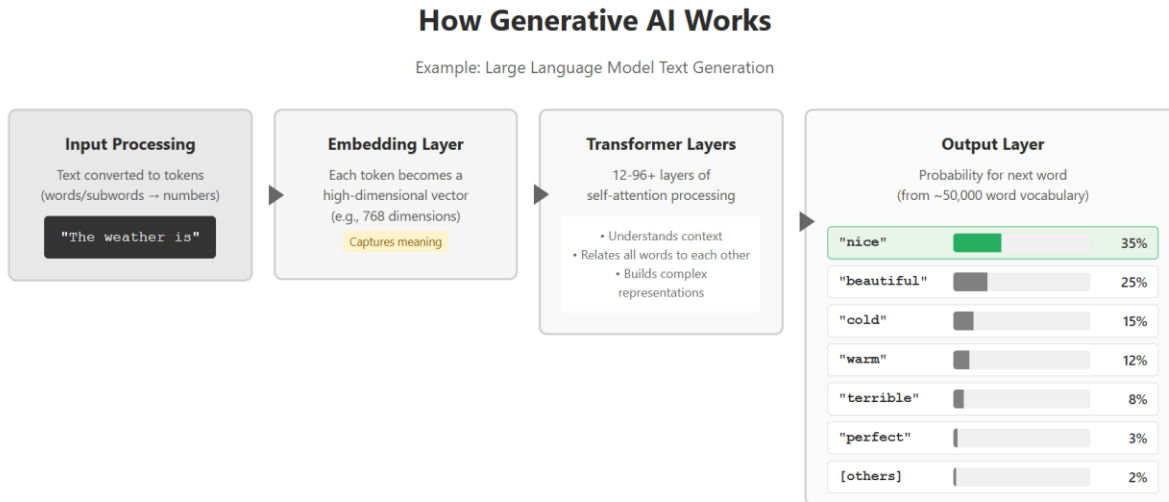
As Figure S3.3 illustrates, the input flows through all these layers until it reaches the output layer. At this point, the model has built a rich, contextual understanding of the prompt. It produces a probability score for all ~50,000 possible tokens in its vocabulary. For the prompt "The weather is," it might assign a 35% probability to "nice," 25% to "beautiful," and 15% to "cold." The model selects a word based on these probabilities, adds that new word to the input,

---

<sup>126</sup> The distinction between a "river bank" and a "financial bank" is the canonical example of polysemy (lexical ambiguity) in computational linguistics. It has served for decades as the standard illustration of how context determines meaning, appearing in everything from introductory textbooks to seminal research papers (see, e.g., Peters et al. 2018 on ELMo embeddings, which are among the first models to dynamically assign different mathematical representations to the same word based on its specific context within a sentence).

and runs the entire process again. This iterative, word-by-word process is what allows the AI to maintain a coherent train of thought.

**Figure S3.3 Explaining LLMs: Generative AI**



*Notes: Generative AI predicts one word at a time by analyzing patterns learned from vast amounts of text. Each word is converted into a numerical vector (embedding) that captures its meaning—similar words have similar vectors. The transformer layers then process these embeddings using "self-attention," allowing every word to consider its relationship with every other word in the context. Through 12-96+ layers of processing, the model builds increasingly sophisticated understanding: early layers grasp grammar, middle layers understand phrases, and deep layers comprehend abstract concepts and reasoning. The output layer produces probabilities for all ~50,000 possible next words. The model selects one (often sampling from top candidates for variety), adds it to the input, and repeats the process—generating coherent text word by word. Unlike simple classification networks that make one decision, generative AI makes thousands of sequential decisions, with each new word depending on everything that came before.*

### *Deep Learning that Goes Beyond Generative AI*

Self-driving cars illustrate how deep learning extends beyond language. These systems use 'sensor fusion'—combining cameras, radar, and LIDAR—to construct a unified model of their environment (Bojarski 2016).<sup>127</sup> The classic approach to this problem relied on a combination of

<sup>127</sup> Sensor fusion is the process of combining data from multiple different sensors—such as cameras (visual light), radar (radio waves), and LiDAR (laser light)—to create a single, unified model of the environment (see Hall and Llinas 1997). This is critical because each sensor has different strengths; cameras read signs, while radar functions in fog. The AI fuses these streams into one cohesive 3D 'world model'—an internal representation of the vehicle's surroundings that

two specialized network types. First, CNNs analyze visual data from cameras, using a hierarchical process—learning edges, then shapes, then objects—to recognize pedestrians and road signs (LeCun et al. 2015). Second, RNNs add a layer of sophistication by using their internal memory to analyze time-series data, allowing them to predict the likely future trajectory of a moving car (Graves 2012). These models are continuously refined using backpropagation on massive datasets of real-world driving scenarios (Goodfellow et al. 2016).

Leading up to 2025, companies like Waymo, Nuro, and Wayve experimented with “Visual Language Models” (VLMs) to enhance their vehicles’ ability to handle unexpected scenarios (Vaswani et al. 2017).<sup>128</sup> This marked a shift from traditional approaches that relied on pre-programmed rules. For instance, Waymo demonstrated that using Google’s Gemini AI model helped vehicles identify and respond appropriately to novel objects they weren’t specifically trained to recognize, such as an unexpected dog crossing the street (Google 2023).

Having established how modern AI systems work, I now turn to the historical trajectory that produced them—a path marked by long periods of stagnation punctuated by rapid breakthroughs.

## BRIEF HISTORY OF AI UNTIL 2025

The technical capabilities I have described did not emerge overnight. Understanding why AI reached its 2025 inflection point—and why it developed in one direction rather than others—requires tracing the field’s historical trajectory. This history reveals a pattern of uneven progress: long periods of incremental advance interrupted by breakthrough innovations that reshaped the field’s possibilities. It also reveals the growing role of compute and capital in determining who could participate at the frontier.

The history of AI development, particularly for neural networks, mirrors the pattern that paleontologists Eldredge and Gould (1972) termed “punctuated equilibrium”: long periods of relative stasis interrupted by brief bursts of rapid transformation. For AI, the stasis periods—sometimes called “AI winters”—saw incremental refinement of existing techniques. The punctuations came when breakthrough innovations, like the backpropagation algorithm in the 1980s, or the Transformer architecture in 2017, suddenly unlocked new capabilities and triggered cascades of follow-on innovation.

### **A Promising Start, Followed by False Starts**

The AI field was formally established in the mid-20th century, particularly following the legendary Dartmouth AI Conference in 1956, where researchers such as John McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude Shannon coined the term “artificial intelligence”

---

tracks objects, predicts their motion, and supports navigation decisions—that is more robust than any single sensor (see Yurtsever et al. 2020).

<sup>128</sup> VLMs are a type of AI trained on massive datasets of images paired with text. This allows the model to find statistical relationships between visual data (pixels) and linguistic concepts. In practice, it moves the self-driving car’s AI beyond simple object detection (e.g., “pedestrian”) to a deeper contextual understanding (e.g., “a pedestrian who is looking at their phone and appears distracted”), which is essential for handling novel scenarios (see Radford et al. 2021).

(McCarthy et al. 1956). While this concept traces its roots to the 19th century—with Charles Babbage’s design of the Analytical Engine and Ada Lovelace’s insight that machines could manipulate symbols to produce meaningful content—Alan Turing’s work in computation and machine intelligence from the 1930s to 1950s demonstrated how machines could mimic human thought and solve complex problems (Turing 1950).

Turing made three profound reinforcing contributions. First, he conceptualized the “Turing Machine,” a theoretical framework that represents the heart of digital computing as a simple, rule-based device that could execute any algorithmic process (Turing 1936). Second, during World War II, Turing’s decryption of the Enigma code showcased computational machines’ ability to tackle complex problems (Hodges 1983). Finally, he introduced the “Turing Test” to determine if a machine’s behavior is indistinguishable from a human’s, establishing an enduring criterion for identifying artificial intelligence (Turing 1950).<sup>129</sup>

However, early optimism soon hit a wall. In 1958, Frank Rosenblatt’s “Perceptron” offered a promising early neural network, but it was limited. When Minsky and Papert (1969) mathematically proved that single-layer Perceptrons could not solve basic logical problems, funding dried up. This triggered the first “AI Winter,” a period of stagnation that lasted until the breakthroughs of the 1980s.

## **The Deep Learning Revolution**

The advent of machine learning changed AI research and practice forever. This becomes clear when examining the foundational academic research behind AI during the post-World War II era (see Table 11.1). The trajectory of neural network development spans over seven decades, with each advance building upon previous discoveries to transform theoretical concepts into today’s powerful language and vision models.

As Table 11.1 suggests, the University of Toronto emerged as a crucial hub for deep learning, embodied by the work of Geoffrey Hinton and his students. Hinton and his collaborators fundamentally transformed how neural networks function, most notably by popularizing the backpropagation algorithm in 1986. The field experienced another leap forward in 2006 when Hinton introduced deep belief networks, which reignited interest in deep learning after years of skepticism. Perhaps most dramatically, 2012 marked a watershed moment when Hinton and his students developed AlexNet, a deep convolutional neural network that achieved unprecedented accuracy in the ImageNet Large Scale Visual Recognition Challenge—reducing the error rate by nearly half compared to the previous year’s winner and demonstrating decisively that deep learning could outperform hand-engineered approaches (Krizhevsky et al. 2012).<sup>130</sup>

---

<sup>129</sup> While Turing’s operationalization of “intelligence” as behavioral indistinguishability remains influential, modern AI systems have complicated its interpretation: today’s chatbots can fool casual interlocutors while failing at tasks any human could perform—a topic I’ll revisit further below.

<sup>130</sup> The ImageNet project, started by Fei-Fei Li in 2006, provided researchers with a huge database of images for object recognition research. As of 2010, it contained more than 14 million hand-annotated images categorized across approximately 22,000 categories. By training neural

The success of AlexNet catalyzed the widespread adoption of deep learning approaches across multiple domains, from image and speech recognition to natural language processing—but it was a winding road to get there. While AlexNet quickly revolutionized computer vision, progress in language and speech initially relied on a different architecture: Long Short-Term Memory (LSTM) networks. Introduced by Hochreiter and Schmidhuber (1997), LSTMs solved the problem of neural networks "forgetting" earlier inputs, making them the standard for large-scale neural machine translation systems, including Google Translate, for nearly a decade (Wu et al. 2016). However, LSTMs were difficult to train at scale due to their sequential nature.

The paradigm shifted decisively in 2017 when Google researchers introduced the Transformer architecture (Vaswani et al. 2017). By leveraging the "self-attention" mechanism described earlier to process vast amounts of data in parallel, the Transformer finally overcame these bottlenecks, uncorking the bottle and leading to the explosion of frontier models that populated the imagination during the AI revolution of the early 2020s (Brown et al. 2020).

### *Reinforcement Learning and RLHF*

Yet another major watershed on the road to the AI revolution was the advent of Reinforcement Learning (RL), a broad field of machine learning where an AI agent learns to make decisions by taking actions in an environment to maximize a cumulative reward signal (Sutton and Barto 2018). Q-learning is a classic method where an AI learns by building a table of values that rates the quality of a given action in a given situation, while policy gradient methods directly train a neural network to learn a "policy" (a set of rules for making decisions). Deep Reinforcement Learning (DRL) combines RL with deep neural networks. For example, "TD-Gammon" achieved superhuman performance at backgammon, demonstrating the potential for machines to learn complex strategies through self-play (Tesauro 1995). This technique was later used by Google's AlphaGo in 2016 to master Go—defeating human champions and showcasing the power of combining deep learning with reinforcement learning and search (Silver et al. 2016).

Research into RL eventually culminated in what were among the most cutting-edge AI models by the end of 2025, those that used Reinforcement Learning with Human Feedback (RLHF), a technique I will elaborate on shortly ahead (Christiano et al. 2017; Ouyang et al. 2022). Rather than relying on a simple, pre-programmed reward rule, RLHF uses a separate "reward model" trained on human preference data—reviewers rank or compare different model outputs based on helpfulness, safety, or accuracy. The main model then optimizes its behavior against this learned reward signal, iteratively improving by taking actions that score well under the reward model (Christiano et al. 2017; Ouyang et al. 2022).

### **Table 11.1 Selected Groundbreaking Research Behind Modern AI**

---

networks on these image-label pairs, researchers enabled the models to classify objects with superhuman accuracy (see Deng et al. 2009).

<b>Year</b>	<b>Milestone/Event</b>	<b>Contributors</b>	<b>Significance</b>
<b>1943</b>	McCulloch-Pitts Neuron Model	Warren McCulloch, Walter Pitts	The first mathematical model of a neuron, laying the foundation for neural networks.
<b>1958</b>	Perceptron	Frank Rosenblatt	Introduction of the perceptron, an early neural network model capable of binary classification.
<b>1980</b>	Neocognitron	Kunihiko Fukushima	Early convolutional neural network for pattern recognition.
<b>1986</b>	Backpropagation Algorithm	Hinton, Rumelhart, Williams	Made training deep neural networks efficient, revolutionizing the field.
<b>1986</b>	Parallel Distributed Processing	Rumelhart, McClelland	Introduced the concept of distributed representations and multi-layer networks.
<b>1989</b>	TD-Gammon	Gerald Tesauro	Used neural networks to achieve superhuman performance in backgammon via reinforcement learning.
<b>1995</b>	Support Vector Machines	Vapnik, Cortes	Introduced a powerful method for classification that competed with neural networks for decades.
<b>1997</b>	LSTMs (Long Short-Term Memory)	Hochreiter, Schmidhuber	Solved the "vanishing gradient" problem for time-series data, powering early speech and text AI.
<b>2006</b>	Deep Belief Networks	Hinton, Osindero, The	Demonstrated unsupervised pre-training, reigniting interest in deep learning.
<b>2006</b>	Restricted Boltzmann Machines	Geoffrey Hinton	Developed as a key building block for training Deep Belief Networks.
<b>2012</b>	AlexNet	Hinton, Krizhevsky, Sutskever	Achieved a breakthrough in image recognition and popularized CNNs.

Year	Milestone/Event	Contributors	Significance
2014	GANs	Ian Goodfellow	Introduced a new framework for training neural networks to generate realistic data.
2015	ResNet (Residual Networks)	He, Zhang, Ren, Sun	Allowed networks to scale from dozens to hundreds of layers, crucial for modern computer vision.
2016	AlphaGo	DeepMind	Defeated a world champion Go player using deep neural networks and reinforcement learning.
2017	Transformers	Vaswani et al.	Introduced the attention mechanism, the architecture behind all modern LLMs.
2018	BERT	Google AI	Set new benchmarks in NLP by using bidirectional training of Transformers.
2020	GPT-3	OpenAI	Demonstrated that scaling up language models unlocks massive few-shot learning capabilities.
2020	Diffusion Models	Ho, Jain, Abbeel	Introduced the noise-removal process that powers modern AI art generators (like Midjourney/DALL-E).

Sources: Barto and Sutton (2018); Bishop (2006); Brown et al. (2020); Goodfellow et al. (2014); Goodfellow et al. (2016); He et al. (2016); Ho et al. (2020); Hochreiter and Schmidhuber (1997); Huh et al. (2016); Jordan and Mitchell (2015); Karras et al. (2019); Krizhevsky et al. (2012); Konečný et al. (2016); LeCun et al. (2015); McMahan et al. (2017); Minsky and Papert (1969); Nilsson (2009); Pan and Yang (2010); Radford et al. (2015); Raghu et al. (2019); Russell and Norvig (2021); Silver et al. (2016); Tesauro (1989); Vapnik and Cortes (1995); Vaswani et al. (2017); Yang et al. (2019).

### More Recent Innovations

Into the 2010s, the field continued to innovate with a series of breakthroughs that expanded AI's capabilities. Generative adversarial networks (GANs), for example, introduced a framework where two neural networks compete against each other to create highly realistic synthetic images

(Goodfellow et al. 2014). By 2020, however, the frontier shifted toward diffusion models (Ho et al. 2020). Unlike GANs, diffusion models generate images by iteratively denoising random noise—a training and generation paradigm that is often more stable in practice and that underlies major image generators such as OpenAI’s DALL·E 2 (Ramesh et al. 2022).

This progress was also made possible by improvements in how models were built and adapted—especially the ability to harness larger context windows (Tay et al. 2022) and to benefit from domain-specific fine-tuning (Howard and Ruder 2018; Raffel et al. 2020). Therefore, by the time generative chatbots entered the mainstream, they could rapidly absorb industry-specific knowledge, adapt to novel conversational contexts, and—when paired with appropriate prompting and system scaffolding—sometimes course-correct in response to feedback, setting the stage for dramatic improvements in real-world performance.

Alongside these technical advancements, an emerging ecosystem of safety measures, explainability frameworks, and governance protocols helped ensure these advances operated ethically and built public trust. For example, RLHF and prompt-engineering practices improved instruction-following behavior and helped align outputs with human values and organizational policies (Ouyang et al. 2022).<sup>131</sup> Coupled with explainable AI (XAI) techniques that illuminate how models arrive at their decisions, these advances supported expert oversight as subject-matter specialists refined model outputs (Simonyan et al. 2014; Lundberg and Lee 2017).

Interdisciplinary teams that included domain experts, data scientists, and operations managers further ensured that deployed AI systems were informed by human judgment and operational constraints (Brynjolfsson and McAfee 2017; Christiano et al. 2017; Ziegler et al. 2019).

Meanwhile, advances in multimodal AI systems during the early 2020s demonstrated the field’s growing sophistication in integrating diverse types of data—text, images, audio, and even video—into more unified modeling frameworks (Baltrusaitis et al. 2019; Lu et al. 2019). Such systems enable more nuanced understanding and more natural human–AI interaction. OpenAI’s GPT-4, for example, helped mainstream the idea of chat systems that can accept both images and text as inputs (OpenAI 2023), ushering in a trend toward increasingly comprehensive AI approaches that span multiple domains and modalities.

Alongside these capability advances, the field developed new methods for responsible deployment. Major tech firms championed new privacy-preserving methods. Apple is a paradigmatic example: it promoted differential privacy and related approaches intended to reduce privacy risk while learning from user data (Apple 2017). This approach traded some model performance for user trust—a strategic choice that reflected Apple’s hardware-centric business model, in contrast to the data-hungry approaches of advertising-driven competitors.

Finally, transfer learning revolutionized development by allowing a model pre-trained on a vast dataset to be fine-tuned for new, related tasks—dramatically reducing the labeled data and compute needed to build sophisticated applications (Pan and Yang 2010). But even as these techniques improved sample-efficiency and adaptability, frontier performance increasingly

---

<sup>131</sup> On the broader landscape of AI ethics and governance frameworks during this period, see Hagendorff (2020), Mittelstadt et al. (2016), and Cath et al. (2018).

depended on sheer scale—so the practical bottleneck shifted toward the compute required to train and iterate on ever-larger models.

### **Progress in AI Research Translated into Increased Compute to Train Models**

While early breakthroughs in machine learning focused largely on language-only or single-modality tasks, as the field advanced toward frontier systems capable of combining text, images, and complex reasoning, the computational cost of progress skyrocketed. As model architectures became more complex—particularly during the deep learning era and the subsequent LLM wave—the compute required for a single frontier training run surged by orders of magnitude. In fact, since 2012, the amount of compute used in the largest AI training runs has doubled approximately every 3.4 months—a rate vastly outpacing Moore’s Law (Amodei and Hernandez 2018). Engineers at leading AI labs responded by allocating ever-increasing resources to their training exercises, turning compute from a utility into a strategic asset.

Figure S3.4 visualizes this exponential escalation by plotting the growth in floating-point operations (FLOPs) used to train prominent AI systems between 1950 and 2024.<sup>132</sup> The logarithmic y-axis highlights the exponential nature of this escalation: the roughly straight upward trend—represented by the linear OLS prediction line—underscores that each successive generation of leading models typically required orders of magnitude more compute than its predecessors. The pattern is stark. What began as trivial compute in early Perceptron-era experiments (see Minsky and Papert 1969) grew to  $4.7 \times 10^{17}$  FLOPs for AlexNet in 2012 (Krizhevsky et al. 2012),  $3 \times 10^{23}$  FLOPs for ChatGPT-3, and then to tens of septillions for subsequent LLM models (see Epoch AI 2024).<sup>133</sup>

This shift changed the industry’s political economy. When state-of-the-art systems required modest compute, academic and nonprofit labs could compete at the frontier. Once competitive training runs required hundreds of millions of dollars in hardware and energy, the center of gravity inevitably shifted toward well-capitalized industry labs (Ahmed and Wahed 2020).<sup>134</sup>

---


<sup>132</sup> A FLOP refers to a single mathematical calculation involving numbers with decimal points (e.g.,  $3.14 \times 2.5 = 7.85$ ). Unlike “integer operations” that use whole numbers, floating-point operations handle a wide range of values by allowing the decimal point to “float.” The ability to handle decimal-point calculations is crucial for scientific computing, graphics, and especially AI, where models perform trillions of such calculations during training and inference.

<sup>133</sup> The figures for ChatGPT-3 were derived from architectural analysis rather than official disclosures (see Patel and Ahmad 2023).

<sup>134</sup> This shift occurred despite significant attempts by the public sector to maintain an alternative, open research ecosystem. During the early 2020s, research consortiums and industry collaborations helped spearhead AI development; these included the Partnership on AI, a collaborative effort by major tech companies and the Allen Institute for AI (Brynjolfsson and McAfee 2022). Meanwhile, the National AI Research Institutes program received a substantial investment from the NSF to advance AI research, and the EU and Japan invested billions in sovereign AI initiatives (EU Commission 2023; METI 2024). Ultimately, however, these public investments were outpaced by the orders-of-magnitude larger capital expenditures of the private hyperscalers.

The talent wars described at this section of the appendix’s opening reflect this dynamic: the nine-figure compensation packages exist precisely because only a handful of organizations can afford frontier-scale compute—therefore, the researchers who know how to use it are extraordinarily valuable. By the same token, as Chapter 9 of the book explores, the exponential scaling of compute revealed in Figure S3.4 helps explain the industry's consolidation, structure, and its potential future trajectory into the late 2020s.

**Figure S3.4 Training Compute (FLOP) of Significant ML Systems (1950–2025)**



*Notes: The figure plots the training compute required for each model in the dataset against its publication date, with the y-axis shown on a base-10 logarithmic scale. “Training compute” refers to the approximate total floating-point operations used to train the base model (not inference compute—the resources consumed when a trained model generates responses to user queries—and not fine-tuning unless the dataset explicitly records it as the main training run). The fitted line is an OLS regression of  $\log_{10}(\text{training compute})$  on time.*

*Sources: Epoch AI (2024).*

Consequently, the technical improvement trajectory of AI technologies also showed dramatically accelerating growth: The compute explosion documented in Figure S3.4 translated directly into model scale; Figure S3.5 visualizes the “scale trajectory” of frontier AI by plotting prominent models’ release dates against their parameter counts on a base-10 logarithmic axis. To make that trend interpretable, like Figure S3.4, the latter overlays an OLS fit through the log-transformed parameter counts. The trend mirrors the growth in compute: model size grows exponentially, moving from millions of parameters to billions, and then to trillions, within a single decade.

This acceleration matters because parameter count is a useful (though imperfect) proxy for a model’s representational capacity. Parameters are the learned weights of a neural network—the

internal degrees of freedom that allow the model to store statistical regularities and compress information from its training data. As parameter counts rose, models became better able to capture complex patterns and long-range relationships, which translated into major gains in language understanding, text generation, multimodal reasoning, and instruction following (see Kaplan et al. 2020).


To be sure, parameter count alone does not determine performance. While Kaplan et al. (2020) established the initial scaling laws linking model size to capability, researchers at DeepMind subsequently demonstrated that many frontier models were “undertrained”—they had more parameters than their training data could effectively utilize. These findings, derived from their model “Chinchilla,” established the so-called Chinchilla scaling laws that shifted the field toward more balanced training regimes (Hoffmann et al. 2022).

Notwithstanding this critical insight, the figure also implicitly shows the core economic tradeoff of the AI revolution: scaling parameters dramatically increases the compute, memory, and energy required to train and deploy systems, which helps explain why progress at the frontier became inseparable from the massive capital investment in data centers and specialized accelerators that I discuss in Chapter 9 of the book.

The dramatic gains in scale had real-world implications far beyond benchmarks. As training pipelines streamlined, teams could prototype new architectures and curate fresh datasets in weeks rather than months. This created a virtuous feedback loop: more capable models produced richer outputs, which could be used to generate higher-quality training data for the next generation (Wang et al. 2023). Throughout 2024–2025, this loop intensified with “agentic” fine-tuning methods that taught models to plan and self-correct. To fuel this process, developers turned to increasingly specialized corpora—from millions of lines of code to vast libraries of user-interface logs (see Wei et al. 2022)—transforming data curation from a janitorial task into a high-leverage engineering discipline.

This shift had significant implications for market structure: firms with access to proprietary data streams, particularly the platform companies discussed in Chapter 8 of the book, held advantages that pure-play AI labs could not easily replicate. I will take this topic up again in Chapter 9 of the book.

### **Figure S3.5 Number of Parameters of Large Language & Multimodal Models**



*Notes: Data includes widely reported AI models with publicly disclosed (or credibly estimated) parameter counts and publication dates through 2025. Parameter counts refer to total trainable parameters for the base model (not fine-tuned variants), and mixture-of-experts models use the reported total parameter count (rather than “active” parameters) where both are available. The fitted line is an OLS regression of  $\log_{10}(\text{parameter count})$  on time (publication date). Source: Epoch AI (2024).*

Ultimately, however, this 'scaling era' encountered a qualitative ceiling. While adding parameters and FLOPs produced models with encyclopedic knowledge and native fluency, mere size did not automatically yield the reliable, multi-step logic required for high-stakes problem-solving. As the marginal returns on raw parameter counts began to plateau, the industry's focus pivoted toward architectures and post-training methods that prioritized depth of thought over breadth of data—and toward inference-time techniques that could trade additional 'thinking' compute for reliability.

The Chinese laboratory DeepSeek offered a striking demonstration that architectural efficiency could substitute for brute-force scale with its release of the DeepSeek-V3 and DeepSeek-R1 models. Its innovation was twofold. First, it optimized the base model using a Mixture-of-Experts design (activating only a fraction of parameters per token) and Multi-Head Latent Attention to compress memory overhead (Guo et al. 2025). Second, to unlock reasoning capabilities, they pioneered Group Relative Policy Optimization (GRPO)—a reinforcement learning technique that incentivized the model to verify its own logic chains without the massive computational cost of traditional training methods. DeepSeek’s models matched or surpassed those from OpenAI, Meta, and Anthropic on several benchmark tasks while reportedly costing only a fraction of what American labs paid for models of similar scale (Xu 2025).

This achievement served as a powerful market validation of test-time compute—a scaling law first proposed by Snell et al. (2024). While the authors had demonstrated that allowing a model to allocate more computation after a prompt is received can produce substantial gains in theory, DeepSeek operationalized this insight in practice.

## **The Rise and Consolidation of AI Reasoning Models**

DeepSeek-R1 represented a broader paradigm shift: the consolidation of AI reasoning models. These systems, unlike earlier models that typically produced a direct answer in one pass, were designed to deliberate internally—generating and checking intermediate steps before responding—so they could solve harder problems more reliably. Developers soon converged on three families of techniques that made these systems both more capable and more dependable: alignment by feedback (e.g., RLHF), deliberation strategies (e.g., chain-of-thought prompting (CTP) plus sampling/search methods), and grounding and tool access (e.g., Retrieval-Augmented Generation (RAG) and structured tool use).

First, RLHF—the technique I introduced in the historical section above—fine-tunes models using human preference data to better follow instructions and align outputs with human judgments (Christiano et al. 2017; Schulman et al. 2017; González Barman et al. 2025). The AI uses the feedback to create its own internal “reward system” and applies reinforcement learning (a method of learning through trial and error) to practice and improve until it consistently gives the answers the reviewers seek. A turning point was InstructGPT, which operationalized this pipeline at scale for instruction-following chat behavior (Ouyang et al. 2022).

Second, CTP involves training models to break down complex tasks into a series of logical, step-by-step explanations. This method leverages the model’s underlying Transformer architecture and attention mechanisms, a key feature of transformers that allows the model to focus on the most relevant parts of the input data as it processes it (Vaswani et al. 2017). In practice, many gains came not only from writing intermediate steps, but from decoding methods that sample multiple reasoning paths and select the most consistent answer (Wang et al. 2023).

Third, RAG gives an AI model the ability to dynamically access external information, like a search engine or a company’s internal database. A key enabling component in many modern RAG pipelines is dense neural retrieval—a method that represents both queries and documents as vectors in a shared semantic space, allowing the system to find relevant passages based on meaning rather than keyword matching (Karpukhin et al. 2020). This helps the AI stay current with information and produce answers that are both relevant to the user’s request and grounded in verifiable facts (Wei et al. 2022; Yao et al. 2023; Guu et al. 2020; Lewis et al. 2020).

The convergence of these three innovations matters historically because the “reasoning era” wasn’t one trick—it was an engineering synthesis of how models are trained, how they are prompted/decoded, and how they consult external resources. These advances were built on earlier breakthroughs that turned large neural language models into practical general-purpose systems: the Transformer architecture and attention (which enabled efficient scaling), empirical scaling laws linking performance to data/model/compute (Kaplan et al. 2020), and later work on

compute-optimal training that clarified how to allocate parameters vs. tokens (Hoffmann et al. 2022; see also Vaswani et al. 2017 for the underlying architecture).<sup>135</sup>

While these breakthrough techniques showed progress, they also highlighted a fundamental limitation: AI models in 2025 did not build causal maps—a deep understanding of cause-and-effect relationships. Instead, they relied on memorized heuristics (mental shortcuts or rules of thumb) and pattern-matching, which made them prone to breaking or failing unexpectedly and poor at applying knowledge to new, unseen situations (Marcus and Davis 2019). Decades of cognitive-psychology work on human reasoning (Kahneman and Tversky 1974) and more recent mechanistic-interpretability probes (tools that attempt to reverse-engineer an AI model’s internal workings) of AI models (Mitchell 2025; Méloux et al. 2025) confirmed this distinction. Numerous studies revealed that, while humans use flexible, principle-based understanding, LLMs wire up thousands of tiny “circuits” that encode memorized rules of thumb. For example, a navigation model can achieve high accuracy by recalling per-origin-destination rules (Vafa et al. 2024), and models solve arithmetic problems using narrow-band “tricks” rather than a general algorithm (Nikankin et al. 2025).

Researchers therefore went beyond RLHF, CTP, and RAG to address these shortfalls, exploring approaches that combined neural networks with more structured forms of reasoning. They experimented with neuro-symbolic hybrids: models that combine a deep-learning "perception" module, which labels objects, with a traditional symbolic logic engine, which applies "if-then" rules to make inferences (see Mao et al. 2019). They also developed program-aided models: systems that translate natural-language prompts into executable code, run that code to get a precise answer, and then convert the result back into a human-readable response (see Gao et al. 2022). OpenAI famously operationalized these academic ideas along three lines when building its pioneering “o-series” reasoning models.

### *Commercializing Progress on the Reasoning Frontier*

First, the AI pioneer adopted process supervision—rewarding intermediate steps, not just final answers—to get models to reason before responding and to verify their own steps (Lightman et al. 2023).

Second, it scaled test-time compute—the concept I introduced in the DeepSeek discussion above. By using large-scale reinforcement learning and search-like sampling to encourage longer internal deliberation, OpenAI demonstrated that spending more "thinking time" improves performance on complex tasks (OpenAI 2024a; OpenAI 2024b; OpenAI 2024c). This confirmed the economic shift described earlier: whereas the previous era focused on massive training compute, the reasoning era shifted the bottleneck to inference compute. A complementary

---

<sup>135</sup> While “parameters” refer to the learned weights that determine a model’s representational capacity, “tokens” refer to the volume of text data consumed during training—roughly analogous to the number of words read. Readers may recall that in the previous chapter I defined “token” as the basic units of text, typically words or word fragments. For the significance of the parameter-to-token ratio, see the discussion of Hoffmann et al. (2022) in the main text.

innovation was deliberative alignment—an approach in which models are given explicit safety specifications and trained to reason over them before generating a response (Guan et al. 2024).

Third, it made these models agentic—capable of autonomously executing multi-step tasks, using tools, and taking actions in the world rather than merely generating text—and multimodal in production. Both o3 and o4-mini were “trained to think longer before responding” and can choose and chain tools (e.g., web search, Python/code execution, retrieval) inside ChatGPT. They even “think with images” by manipulating visual inputs within their internal reasoning (Jin et al. 2025; OpenAI 2025c; OpenAI 2025d; OpenAI 2025e).<sup>136</sup>

### *Why This Mattered*

This distinction birthed a new theoretical architecture known as "Orchestration" or "System 2 Routing." Rather than forcing every user query through an expensive reasoning model, 2025-era systems began using lightweight "router" models to assess difficulty. Simple queries were handled by fast, cheap models (System 1), while complex puzzles triggered the 'thinking' models (System 2)—terminology borrowed from Kahneman's (2011) influential dual-process theory of human cognition (see Ong et al. 2024; Li et al. 2025). This successfully hybridized the raw scale of the LLM era with the structured deliberation of the reasoning era.

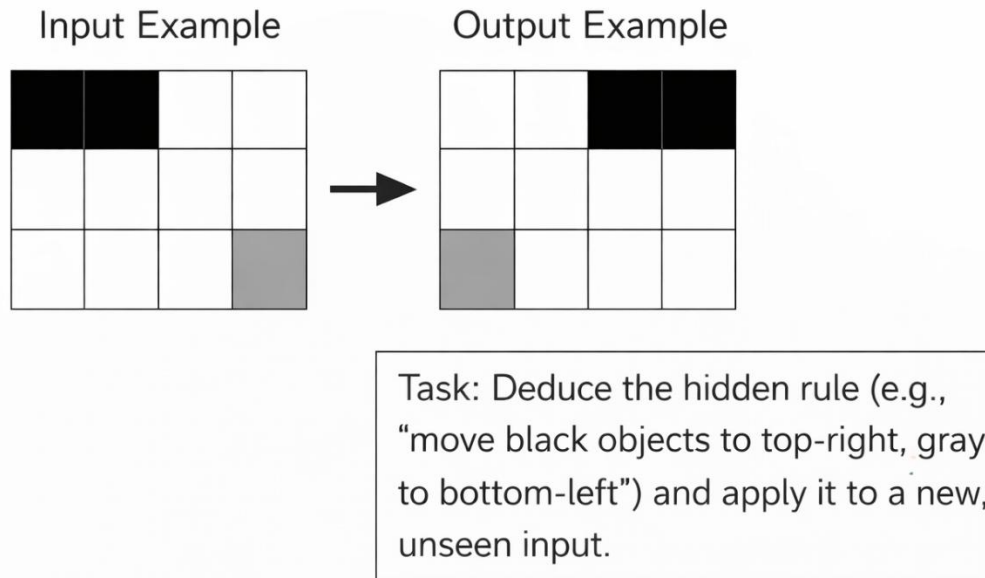
To understand the significance of these advancements, consider the ARC Prize, a competition focused on achieving AGI. The prize uses the Abstraction and Reasoning Corpus (ARC), a series of unique visual puzzles the AI model has never seen before. Each puzzle requires it to figure out the underlying logical rule and apply it—much like a human would—a key skill for true problem-solving.

As illustrated in Figure S3.6, an ARC task typically consists of a few "training" examples—pairs of input and output grids that demonstrate a specific transformation rule. The AI must infer this hidden logic (for example, "move all black blocks to the nearest corner" or "invert the shading pattern") and apply it to a final "test" input to generate the correct solution. Unlike standard image recognition, which relies on memorizing textures or shapes, this requires the model to perform abstract reasoning and generalization from very few examples.

### **Figure S3.6 Stylized Representation of an Abstraction and Reasoning Corpus (ARC) Task**

---

<sup>136</sup> This tool-using agent ability has clear roots in research that interleaves reasoning traces with actions (e.g., ReAct) and in approaches that train models to decide when to call external APIs (Yao 2023).



*Notes: The figure depicts the structure of a standard ARC challenge. On the left, the "demonstration" pairs show a consistent logical transformation (in this case, an object movement, shading inversion, or pattern completion rule). The model must identify this rule and apply it to the "Test Input" on the right to produce the correct output grid.*

*Source: Adapted from Chollet (2019).*

These puzzles are divided into different difficulty levels. On one of the harder test sections of the original ARC benchmark, called the Semi-Private Evaluation set, OpenAI’s o3-preview (a powerful, specialized version of its AI reasoning model) reached 75.7% under a compute cap and 87.5% when compute limits were lifted (Chollet 2024).<sup>137</sup> These results suggest that o3 possesses novel task adaptation abilities, which some researchers interpret as an important step toward AGI (see Kamradt 2025). However, as I discuss below, whether such benchmark performance reflects genuine generalizable reasoning—or a more sophisticated form of pattern-matching on a particular task distribution—remains contested.

As the ARC competition evolved, AI companies introduced new reasoning models. In early 2025, the original ARC benchmark was renamed ARC-AGI-1 after its creators introduced a significantly more difficult successor, ARC-AGI-2. In August 2025, OpenAI released a newer

<sup>137</sup> On the less challenging Public Evaluation set, the o3-preview model solved approximately 83% of the puzzles with less computing power and about 92% when it used more.

reasoning model, GPT-5 Pro.<sup>138</sup> While the newer model outperforms o3 while using 50-80% fewer “thinking” tokens and hallucinating less frequently (OpenAI 2025e), its performance on the ARC benchmarks revealed the limits of recent progress. On ARC-AGI-1—the original benchmark that o3-preview had largely conquered—GPT-5 scored only 65.7%, slightly behind xAI's Grok 4 (68%). The gap widened dramatically on ARC-AGI-2, designed specifically to resist the pattern-matching strategies that succeeded on the original: GPT-5 scored just 9.9%, compared to Grok 4's 16%. These results suggested that the field had not yet achieved genuine generalization.

Part of the reason for this is that AI reasoning models excelled at tool use and code generation, were increasingly able to run Python or Bash commands, iteratively tested outputs, and corrected their own errors. They demonstrate enhanced logical reasoning with fewer instances of repeating failed actions or making simple logical mistakes. Additionally, newer systems evinced adaptive problem-solving abilities, recovering from mistakes instead of getting stuck.

These models also demonstrated capacity to move beyond summarizing existing research toward brainstorming potentially novel ideas by drawing connections across fields such as physics, biology, and engineering (see OpenAI 2025f)—though whether such outputs genuinely match doctorate-level creativity or instead reflect sophisticated recombination remains debated (Bubeck et al. 2023; Knight 2023). Whatever the case, frontier models sometimes produced novel-seeming experimental approaches—from fusion-experiment variants to pathogen-detection concepts—although such proposals generally required expert vetting for feasibility and originality (Bubeck et al. 2023).

These emerging capabilities set the stage for the next phase of competition, in which the major AI labs raced to combine reasoning with autonomous action.

### **The Winter 2025 Escalation: Agents and "Deep Think"**

The rivalry between AI foundation models reached a new fever pitch in late 2025 with the convergence of reasoning capabilities with autonomous agency. Following OpenAI's summer dominance, competitors responded with systems that didn't just "chat" but actively navigated operating systems and cloud environments.

In November 2025, Google released Gemini 3.0, a model that fundamentally integrated "Deep Think" capabilities directly into its multimodal architecture rather than treating reasoning as a separate mode. Unlike previous iterations, Gemini 3.0 demonstrated the ability to natively verify its own visual outputs—correcting errors in diagrams or code generation before showing them to the user—and introduced Project Antigravity, an agentic development platform that allowed

---

<sup>138</sup> GPT-5 introduced so-called adaptive inference: rather than using the same amount of effort for every query, the model can recognize a difficult task and decide to "think longer" by dedicating more computational power to it and thus automatically adjusting its reasoning depth, “saving its strength” for “hard” questions (OpenAI 2025e).

developers to build bots capable of autonomously managing cloud infrastructure (Google 2025b).

That same month, Anthropic countered with the release of Claude Opus 4.5. While exhibiting higher latency than Gemini on standard benchmarks, Opus 4.5 established itself—according to Anthropic—as optimized for 'long-horizon' tasks, capable of executing coding workflows spanning hours without human intervention (Anthropic 2025b).

And right on cue, OpenAI responded in December 2025 with the release of GPT-5.2, an iterative but significant update to its flagship model. This version introduced "Adaptive Voice," which allowed the model's reasoning engine to operate in real-time speech without the latency lag of earlier voice modes. GPT-5.2 also deployed a new safety architecture specifically for “agentic sandboxing”—scenarios where an AI is only granted access to sensitive tools (like bank accounts or email) if its proposed plan is mathematically verified to remain within strict, pre-defined bounds (see OpenAI 2025f).

## THE FUTURE OF AGI

The reasoning capabilities of LLMs nonetheless remained rooted in statistical associations within their training data, not abstraction (Chollet 2019). At their core, these models rely on heuristics and pattern completion that can fail under tricky or adversarial challenges—and even step-by-step prompting via chain-of-thought (Wei et al. 2022) has been shown to be exploitable for strengthening attacks (Su 2024). Their generalization is often fragile—especially outside the distributions and task families they were trained and tuned for (Berglund et al. 2023).<sup>139</sup>

In stark contrast, genuine intelligence hinges on efficient skill acquisition and transfer across tasks, not just pattern fitting (Chollet 2019). The Shojae et al. 2025 systematic study of “Large Reasoning Models” underscores this difference: when a puzzle’s complexity is dialed up in a controlled way, a reasoning model’s performance eventually collapses.<sup>140</sup> Moreover, these models often cannot reliably follow explicit algorithms, suggesting fragile procedural fidelity

---

<sup>139</sup> For example, if a model is trained on a statement such as “A is B,” it often cannot infer the reverse relationship “B is A.” This failure persists even when the reverse relationship is logically equivalent to the original and occurs across different model sizes, families, and training setups (see *ibid*).

<sup>140</sup> Consider how reasoning models perform in the Tower of Hanoi, a classic puzzle game consisting of a set of discs of different sizes and three pegs. The goal is to move the entire stack of discs from a starting peg to a destination peg, following two simple rules: only one disk may be moved at a time, and a larger disk may never be placed on top of a smaller one. The puzzle is a crucial test of systematic, procedural reasoning. In low-complexity scenarios with small  $N$  values, standard LLMs often outperformed specialized reasoning models. At medium complexity, reasoning models showed modest gains. However, at high complexity, both types of models collapsed to near-zero accuracy—with performance degrading sharply around seven disks and collapsing essentially to zero by eight disks—despite an ample token budget. Shojae et al. (2025) further demonstrated that even when fed the exact solution algorithm as part of the prompt, models failed to reliably follow it, revealing they lack systematic procedural fidelity.

rather than robust algorithmic competence (Shojaee et al. 2025). To bridge this gap, researchers pursued different approaches circa 2025.

First, they turned to a “verifiable-rewards” strategy (popularized by DeepSeek) used “rule-based grading” to push models toward longer-form, self-consistent solutions without heavy human supervision (DeepSeek-AI 2025; Zhou et al. 2025). This method relies on automated systems to score a model's output based on a predefined set of rules, which is especially useful for tasks with a clear right or wrong answer like math problems or code generation. For example, the system can automatically check if a final numerical answer is correct or if a generated piece of code compiles and runs without errors.

Second, they experimented with “self-confidence/intuition” approaches (popularized by Berkeley’s Intuitor RLIF) that tried to coax better reasoning out of a model by training it to recognize when it is confident or uncertain and to identify and correct its own mistakes (Zhao et al. 2025). Researchers developed this strategy to combat reward hacking—a phenomenon where an AI exploits flaws in its scoring system to achieve high scores without actually solving problems as intended. A model might, for instance, generate a convincing but incorrect answer that satisfies a flawed grading rubric, demonstrating rhetorical skill rather than genuine understanding (Zhao et al. 2025).

Perhaps more importantly, disappointing scaling results—where successive model generations showed diminishing returns on benchmarks despite exponentially increasing compute investments—helped catalyze the 2025 pivot from a “bigger” to “better” approach centered on greater efficiency and new architectures. Remember that, unlike prior technologies that required constant human ingenuity for iterative improvements, AI can design new algorithms and improve itself. Along these lines, by 2025 AI’s ability to optimize its own underlying architecture led to dramatic advancements in model efficiency. So-called automated hyperparameter tuning acted like an army of expert chefs perfecting a recipe; instead of human engineers manually adjusting the countless settings required for training a model, an AI could rapidly test millions of combinations to find the most effective training regimen (Bergstra and Bengio 2012).

More profoundly, “neural architecture search” (NAS) allowed an AI to act as its own architect, designing the fundamental blueprint of a new, more effective AI model—discovering novel designs that might never have occurred to its human creators (Elsken et al. 2019). Together, these techniques reduced training times from weeks to hours and enabled the creation of more powerful models without a corresponding increase in computational cost (Real et al. 2020).<sup>141</sup>

Beyond AI's ability to optimize existing models, this new era of innovation was also advancing through fundamental changes in software design. Early models were "dense," meaning they used

---

<sup>141</sup> Engineers also adopted compute-optimal training—more scientific "recipes" for balancing a model's size with the amount of training data to get the most performance from a fixed amount of computing power (see Hoffmann et al. 2022)—and embraced new architectures like long-context state-space models to process vast amounts of information, which are akin to reading a long book page-by-page with a running summary rather than trying to memorize it all at once.

their entire massive network for every single calculation (Fedus et al. 2021)—an incredibly inefficient process like turning on every light in a skyscraper just to read a book in one office. The new approach is sparsity, particularly through Mixture-of-Experts (MoE) architectures. Rather than activating the entire network for every query, MoE models function like a large organization with specialized departments: a lightweight routing mechanism assesses each input and directs it to only the small fraction of “expert” subnetworks best suited to the task, leaving the rest dormant (Shazeer et al. 2017). This allows for vastly larger total parameter counts—and thus greater capacity—while keeping inference costs manageable (Shazeer et al. 2017; Fedus et al. 2021).

In addition to these efficiency gains, researchers were also creating entirely new architectural blueprints to overcome the inherent limitations of older designs. A leading alternative to the dominant Transformer architecture is the State Space Model, which is far more adept at handling very long sequences of information. Whereas standard Transformers must compute relationships between all pairs of tokens—a process whose cost grows quadratically with sequence length (Vaswani et al. 2017)—State Space Models maintain a compressed running state that updates incrementally (Gu, Goel, and Ré 2021; Gu and Dao 2023). The difference is like reading a novel while constantly cross-referencing every earlier page versus reading page by page while carrying forward a concise summary. This makes State Space Models particularly suited for very long sequences, such as entire books or genomic data (Gu, Goel, and Ré 2021; Gu and Dao 2023).

At the same time, progress was accelerating by improving how models learn and what they can do. Rather than relying on a bigger pool of data, researchers shifted focus to higher-quality data and smarter training methods like curriculum learning, which teaches a model simple concepts before moving to complex ones—much as a student learns arithmetic before calculus—improving both training efficiency and final performance (Bengio et al. 2009; Xie et al. 2023; Penedo et al. 2024). Yet for all this progress, fundamental limitations persisted.

### **Skepticism was Still Warranted**

New approaches to improving AI were not without problems. Researchers challenged the Berkeley team’s self-confidence gains on the grounds that models might “game the system” instead of engaging in genuine reasoning (Zhou et al. 2025a). While Verifiable Rewards works well in domains with clear right/wrong answers, it doesn’t directly generalize to open-ended tasks (DeepSeek-AI 2025; Zhou et al. 2025a). Indeed, “complementary system-level methods” remained essential for reliability outside of closed-form math and code. They include RAG, which as explained above grounds an AI’s answers in external sources, and tool use, which offloads computations that require precision or access to up-to-date data to external programs like a Python interpreter or a calculator (see Guu et al. 2020; Lewis et al. 2020). They also failed to reliably perform under distribution shift—situations where the data a model encounters during deployment differs systematically from its training data, as when a model trained on formal written English encounters informal speech or regional dialects (Hendrycks et al. 2021; Ovadia et al. 2019). Moreover, 2025-vintage LLMs continued to stumble when confronting tasks requiring coherent planning over many steps and “adversarial conditions”: inputs intentionally designed to trick the model (Valmeekam et al. 2023; Zou et al. 2023; Perez et al. 2022).

The upshot was that, by late 2025, the newer approaches to AI outlined above did not achieve the ability to apply cause-and-effect understanding to new situations and, to boot, they were rapidly hitting scaling limits due to rising computational costs and diminishing returns (You 2025; Kaplan et al. 2020; Hoffmann et al. 2022).

Therefore, by mid-2025, grand proclamations from OpenAI, Google, Anthropic, and others that AI agents would soon perform like PhD-level assistants—handling shopping, travel bookings, calendar management, and financial tracking—seemed overwrought.<sup>142</sup> This included committing “compounded” mistakes—errors in early steps that propagated through subsequent actions, creating cascading failures—that at times ended in failure rates as high as 70 percent (Spice 2025). Hallucinations also persisted, with agents fabricating dates, mismanaging data, and even corrupting databases when shortcuts proved “easier” than following correct procedures (ibid). Security analyses revealed that even hardened agentic systems—those designed with explicit safety constraints—succumbed to adversarial attacks between 50% and 90% of the time (Zhang et al. 2025a). When AI agents have access to sensitive tools like email, financial accounts, or code execution environments, such vulnerability rates represent serious deployment risks.

Finally, and most importantly, Zhao et al (2025b) suggested that the step-by-step problem-solving undertaken by reasoning models were closer to sophisticated pattern matching learned from training examples—not authentic problem-solving capabilities or comprehension. In other words, when confronted with problems sufficiently different from their training experiences, these models were extremely fragile (ibid).

For Marcus (2025a; 2025b), these failures aren't surprising: current LLM-based agents merely mimic surface patterns rather than truly understand tasks, so each additional step multiplies risk. He suggests that while LLMs remain valuable for coding assistance, brainstorming, and writing, they are no substitute for well-specified algorithms or symbolic systems when reliability matters.

This critique has significant implications for AI's status as a general-purpose technology: if LLMs cannot reliably perform multi-step tasks outside their training distribution, their applicability across sectors may be narrower than optimists project. In contrast, although well-specified algorithms consistently give the same, correct answer every time, they lack the flexibility to handle unseen challenges. Marcus (2025a; 2025b) therefore argues that future progress on AI, and potentially AGI, should hinge on hybrid approaches that marry deep learning with explicit symbolic mechanisms that can engage in human-like causal reasoning.

### **Avenues for Future Progress**

By 2025, evolutionary computation, inspired by natural selection, looked like a promising avenue for future AI approaches. It uses genetic algorithms to iteratively optimize solutions by simulating genetic processes, including selection, mutation, and crossover (Mitchell 1998; Back et al. 2018). During the early 2020s, researchers explored and, in some cases, applied evolutionary methods to complex challenges in engineering and robotics by “mutating”

---

<sup>142</sup> This paragraph draws on Marcus (2025a).

candidate solutions and selecting the most effective ones to “maximize fitness” (Such et al. 2017). Going forward, this same approach may help design AI systems themselves: evolutionary search can discover neural network architectures automatically—sometimes at very large scale—without relying on handcrafted designs (Real et al. 2017).

An even more fundamental shift was happening at the hardware level with neuromorphic computing. This approach abandons traditional chip design in favor of creating processors that mimic the architecture of the human brain. Unlike a traditional computer that shuttles data between separate memory and processing units in an energy-intensive process, neuromorphic systems aim to co-locate memory and computation to reduce data-movement costs and improve energy efficiency—potentially enabling capable AI inference on constrained, battery-powered devices (Merolla et al. 2014).

Similar breakthroughs were on the horizon during the early 2020s in quantum machine learning. It combines quantum computing principles with AI to potentially accelerate certain classes of optimization, sampling, and linear-algebra subroutines under specific assumptions (Biamonte et al. 2017).<sup>143</sup> In late 2024, Google and others reported progress toward practical quantum error correction—reducing error rates and improving fault-tolerance in experimental systems—marking a meaningful step toward more reliable quantum computation.<sup>144</sup> By 2025, quantum neural networks, leveraging phenomena like superposition and entanglement, could process multiple states simultaneously (Rebentrost et al. 2014; Schuld et al. 2015), offering advances in optimization and pattern recognition tasks that were challenging classical AI systems.

Taken together, these innovations signal a potential end to the era where brute-force scaling was the only driver of progress. Yet, capitalizing on these disparate advances—combining the efficiency of neuromorphic chips, the adaptability of evolutionary search, and the optimization power of quantum mechanics—requires more than just new components; it demands a unifying architectural logic. As I will explore in the last section of this appendix, the most promising template for integrating these technologies may lie in treating neuroscience findings as literal

---

<sup>143</sup> Recall the “blindfolded hiker” analogy used earlier to describe classical training (see footnote [X]). While a classical computer is like that hiker—painstakingly feeling its way down the mountain one step at a time and risking getting stuck in a small valley (a “local minimum”)—a quantum computer behaves differently. Leveraging superposition, it acts more like a flood of water that fills the entire landscape at once, or a hiker capable of “tunneling” through the mountain itself. This allows it to locate the true lowest point (the global minimum) across massive, high-dimensional datasets exponentially faster than classical methods (see Harrow et al. 2009).

<sup>144</sup> Quantum bits (qubits) are notoriously fragile; even slight temperature changes or stray electromagnetic waves can cause them to lose their data (a process called “decoherence”). The breakthrough involved grouping many unstable “physical” qubits together to form a single, robust “logical” qubit. Historically, adding more physical qubits introduced more noise than it solved. Google’s achievement was demonstrating that for the first time, increasing the size of the error-correcting code actually reduced the error rate, crossing a critical threshold known as the “break-even point” for fault tolerance (see Neven 2024).

design specifications—a “Phase Two” of AI development that could fundamentally alter the industry's trajectory.

But before speculating on that distant horizon, we must confront the immediate reality. Understanding AI requires more than grasping its technical trappings; it demands that we comprehend its industrial organization, its political-economic consequences, and the regulatory paths available to govern it. The latter chapters of the book do just that.

## CONCLUSION

Over the last decade, the architectural and training advances outlined in this section of the appendix were increasingly purchased with compute. As models grew more capable, the resources required to train frontier systems rose sharply, shifting the locus of advantage toward actors who could marshal capital at scale. This created a stark industrial bifurcation: a small number of well-resourced labs produced “foundation models,” while the rest of the economy merely adapted them for specific purposes.

The economics of this divide became the central axis of industry stratification. Companies bet that whoever first achieved AGI would capture value that dwarfed any conceivable infrastructure investment, just as the first nuclear power commanded decisive strategic advantage. This "winner-take-all" framing justified expenditures that no conventional discounted cash flow analysis could support. It also fueled a frantic race toward vertical integration and a web of circular deals that sharply departs from the organizational structure of the late Third Industrial Revolution.

Chapter 9 of the book examines how these technical realities—massive compute requirements, talent scarcity, and capital intensity—translated into the AI industry’s evolving market structure. It details who controls AI's scarce complements and how that control shapes the distribution of wealth and power in the emerging AI economy.

## ONLINE SUPPLEMENTARY APPENDIX TO HISTORY'S MOST REVOLUTIONARY INNOVATION, SECTION 4

In 2025, both the US and China viewed leadership in AI as a strategic imperative. Both nations recognized that the true prize—AGI—required massive, coordinated efforts across academia, industry, and government labs. Each wove AI into a broader industrial-policy framework designed to secure self-sufficiency, lock in competitive advantages on the path to general AI, and control military applications. China's state-directed "national AI teams" and open data mandates accelerated model development, while U.S. alliances with European and Japanese foundries (the "trusted-foundry" network) sought to build a geopolitically secure supply chain for the next generation of nodes (3 nm and below). On the defense front, both sides were busy embedding AI in command-and-control, logistics, surveillance, and cyber-operations.

This escalating rivalry was increasingly compared to the U.S.-Soviet space race, with the 2022 release of ChatGPT widely seen as a "Sputnik moment" that jolted Beijing into action (Chin and Huang 2025). The resulting "AI Cold War" fast became a central driver of national policy in both countries. In Washington, it fueled fears of "authoritarian AI," while Beijing was driven by the conviction that failing to keep pace will allow the U.S. to permanently cap its resurgence as a global power (ibid.).

But perhaps the more salient fear driving the AI Race was America's fear that it was in decline and could fall behind a nimbler, rising power, harkening back to previous episodes in history. Consider that the British were once worried about the rise of the Netherlands in the 17th Century due to financial innovations such as liquid securities markets, which birthed the Dutch East India Company and the growth of a global trading Empire that encroached upon the British sphere of geopolitical influence, including in North America (Israel 1989). The crown responded by adopting mercantilist policies such as the "Navigation Acts," which were aimed at bolstering British traders at the expense of their Dutch counterparts (Zahedieh 2010). Similarly, the U.S. was worried about the rise of Japan in the 1980s, and this motivated the Reagan administration to enact tariffs on Japanese exports and devalue the U.S. dollar.<sup>145</sup>

This sense of rivalry was not just economic but deeply military. The Pentagon, increasingly worried about China and Russia's own AI advancements in their command-and-control systems, was racing to integrate AI into its core defenses (Hirsh 2025). This included new generations of drone ships under the "Replicator" program and a contract with Scale AI for its "Thunderforge" initiative, designed to use AI for theater-level planning. This urgent push to "go faster" and "take risks" formed the direct backdrop for a new national AI strategy (ibid.).

Along these lines, the Trump Administration's "AI Action Plan," released in July 2025, set out an ambitious four-pillar roadmap to cement U.S. leadership in AI through deregulation,

---

<sup>145</sup> These worries faded after Tokyo's 1990 stock market crash, its subsequent economic collapse, and Japan's inability to return to its former economic glory after thirty years of stagnation (Frieden, 2006).

infrastructure spending, and export promotion, alongside developing guardrails around fairness and security (see Trump 2025a).<sup>146</sup>

The first pillar focuses on unleashing innovation through deregulation and investment, directing agencies to identify and roll back unnecessary regulatory burdens that slow AI development, from streamlining environmental reviews for new data-center sites to fast-tracking permitting for edge-computing facilities. It also proposes boosting federal R&D budgets, including domestic manufacturing incentives, and expanding AI-focused training programs, apprenticeships, and immigration pathways for top talent.

The second pillar aims to accelerate the building of AI infrastructure and supply chains, recognizing that inference workloads now outstrip training in energy demand. The plan calls for federal-local partnerships to expand grid capacity, modernize transmission lines, and incentivize data-center construction by exploiting abundant sources of hydrocarbons, nuclear energy, and renewables. It also recommends the US to onshore chip manufacturing, defends against supply-chain disruptions through strategic stockpiles, and encourages U.S. firms to share full-stack AI solutions—from chips through software—with allied partners.<sup>147</sup>

The third pillar strengthens international alliances and AI-related exports by adapting the existing National Technology Transfer Framework, which governs how sensitive U.S. technology is shared with other nations. Under the plan, the Departments of Commerce and State are directed to develop comprehensive AI export packages that could include specialized GPUs, pre-trained models, and technical expertise. The aim is to deepen ties with allies by providing them with trusted technology, thereby promoting a U.S.-aligned approach to AI governance. To establish these norms globally, the U.S. pledges to lead plurilateral efforts with like-minded countries in bodies such as the U.S.-EU Trade and Technology Council, the OECD, and the International Telecommunication Union. In practice, this means working to reach consensus on specific safety benchmarks for testing models, defining criteria for trustworthy AI based on principles like fairness and transparency, and creating interoperable model-exchange formats, allowing AI to work seamlessly across different international platforms.

The fourth pillar embeds guardrails, ethics, and security measures. It orders agencies to audit federal AI deployments for ideological bias, mandate third-party red-team testing—where

---

<sup>146</sup> In January 2025, President Trump rescinded President Biden’s October 2023 Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence, originally issued to establish guardrails for testing advanced AI models and to direct federal agencies to develop, evaluate, and monitor AI systems for security, bias, and safety (see Biden 2023). Trump’s own Executive Order—Removing Barriers to American Leadership in Artificial Intelligence—prioritized minimizing regulatory constraints and promoting U.S. innovation and global AI leadership (see Trump 2025b).

<sup>147</sup> The Trump Administration test ran this approach at the 2025 APEC summit in South Korea, where it sought to persuade Asia-Pacific partners to adopt American chips, software, and open-source models. It also drummed up support for several export-finance mechanisms through the Export–Import Bank and DFC to underwrite AI deals to help US companies like Nvidia and OpenAI to capture markets in healthcare and other sectors (see Ramkumar 2025a).

independent security teams simulate attacks to find flaws before malicious actors do—and codify transparency requirements using the National Institute of Standards and Technology’s voluntary “AI Risk Management Framework” as a guide for managing potential harms. New directives also aim to integrate AI-specific threat modeling, a process for anticipating unique attacks like poisoning a model’s training data, into the Cybersecurity and Infrastructure Security Agency’s (CISA) Critical Infrastructure Resilience guidelines. Finally, the pillar seeks to expand public-private information sharing on AI-born vulnerabilities, a new class of flaws created by the model’s own logic, such as the potential to generate false information or leak sensitive training data.

## **China Strikes Back**

In response to the Trump administration’s AI Action Plan, Chinese Premier Li Qiang unveiled his country’s own counterproposal at the 2025 World Artificial Intelligence Conference in Shanghai. He called for a new multilateral organization headquartered in that city open to all countries and international organizations. The proposal emphasizes securing AI as a global public good rather than the exclusive domain of major economies or cloud providers (Baptista and Cash 2025; Zhang 2025). It calls on deepening global collaboration on AI R&D, pooling talent and data resources, and building shared governance mechanisms addressing standards, ethics, and risk management—without imposing any single governance model.

China’s proposal also includes enabling secure federated learning across borders, harmonizing privacy and cybersecurity protocols, and offering regulatory safe harbors to empower innovators in less-developed markets (Zhang 2025; TechNode 2025). Government officials framed the initiative as complementary to existing UN and OECD efforts, particularly emphasizing inclusion of developing countries through capacity-building programs designed to bridge the digital divide. Beijing advocated co-developing shared protocols and international consensus on AI ethics, model safety, and data governance, explicitly referencing coordination with established standards bodies such as ISO/IEC, and inviting participants to co-author technical guidelines collaboratively (see JTC 1 n.d.; Wikipedia 2018).

## **Circa 2025, Who Led the Race?**

In the global race to achieve AGI, the US held several decisive advantages. In 2025, it boasted the greatest number of AI startups, the most private-sector funding for AI incumbents and startups, and significant AI model development and deployment. More generally, the U.S. boasted a large, vibrant economy, deep and highly liquid capital markets buttressed by the dollar’s status as the global reserve currency, and a political system capable of self-correction. Despite the nativist political leanings and rhetoric of the second Trump Administration, and aggressive efforts to ramp up immigration enforcement, the U.S. continued to attract global talent in the form of high-skilled workers in the high-tech sector. It had also achieved near energy independence, reducing a key strategic vulnerability, and diversified its supply chain, including access to raw materials such as rare earth minerals.

To put these American advantages in perspective, it’s helpful to consider the Rosenbach et al (2025) Critical and Emerging Technologies Index (CETI), a comprehensive ranking that scores

countries on eight key "pillars" of AI strength. By breaking the AI ecosystem into discrete elements, it shows not just who leads overall, but exactly where each nation's advantages lie.

The index captures several dimensions. It measures Compute, which captures access to high-performance chips (GPUs, TPUs) and the datacenter infrastructure needed to train and serve large models. It evaluates Algorithms through the quality and volume of foundational AI research—new model architectures, training techniques, and breakthroughs published in top conferences. The Data pillar captures the quantity, variety, and openness of AI-ready datasets—everything from consumer-behavior streams to industrial telemetry. Human Capital assesses the stock of AI talent: researchers, engineers, and graduates working in academia, industry, and government labs. The index also counts AI Firms, measuring the number and diversity of domestic AI companies and startups, as well as their scale and market reach. Applications looks at real-world deployments across sectors (healthcare, finance, manufacturing) and the depth of integration into existing workflows. Safety & Governance reviews regulatory frameworks, standards bodies, and best-practice guidelines that ensure responsible, secure AI development and use. Finally, International Collaboration measures the extent of cross-border partnerships, researcher exchanges, and participation in global R&D initiatives.

When one scores and weight all eight pillars, in 2025 the United States led on six. These were Compute, Algorithms, AI Firms, Applications, Safety & Governance, and International Collaboration—reflecting its dominance in advanced hardware, cutting-edge research, a vibrant startup ecosystem, broad technology uptake, mature regulatory standards, and extensive global partnerships.

This makes sense. American innovators—OpenAI, DeepMind-backed ventures, Anthropic, and Meta's research labs—operated in an ecosystem rich with the world's top AI talent, massive venture-capital backing, and a culture that prized rapid iteration. Its vibrant open-source AI community produced several initiatives—Hugging Face's Model Hub, Meta's release of Llama, and Stability AI's Stable Diffusion—that drove a flood of high-quality, transformer-based models into the public domain, complete with permissive licenses and tooling integrations. This openness accelerated innovation by enabling thousands of startups, research labs, and mid-sized enterprises to fine-tune and deploy foundation models without the multi-million-dollar R&D budgets required to train from scratch. It also insulated the U.S. ecosystem from single-vendor lock-in, democratizing access to cutting-edge AI and dispersing technical know-how.

On the hardware side, U.S. hyperscalers and chipmakers dominated the leading process nodes (5 nm and below) and specialized inference-accelerator design. Research institutions had pioneered low-precision quantization techniques (e.g., 4- and 8-bit arithmetic) that cut power draw by over 50% with minimal loss in accuracy.


The US also led China in robotics. Fueled by Elon Musk's high-profile ambitions for Tesla's Optimus, startups such as Apptronik (backed by Google DeepMind), Figure AI, 1X Technologies, and Agility Robotics collectively raised well over a billion dollars in 2025 to develop practical bipedal machines for factories, warehouses, healthcare, and eventually homes.

If these advantages seem like an American *fait accompli*, it is only because I've so far neglected other important aspects of the race to AI supremacy. Consider that China topped the CETI index on its other two pillars. With massive, centralized datasets from both private-sector platforms and state-sponsored initiatives, China outstripped the US in sheer volume and consistency of training material under the Data pillar. For Human Capital, China graduated more AI-focused PhDs each year and employed a larger aggregate pool of AI specialists—even though US researchers still led many top-tier academic venues. Because Data and Human Capital together account for a large share of the overall Index score, China's leadership on those dimensions substantially narrowed—and in some scenarios even eliminated—the gap versus the US, despite America's edge across six other pillars.

Another way to do justice to China's AI and high technology accomplishments is by evaluating its shares of publications falling in the top 10% most-highly cited for nine core computing, telecommunications, and semiconductor-related fields. Figure S4.1 compares China to the US across these categories over the 2019–2023 period. In every one of these categories, China outpaces the U.S., with its strongest lead in "Wide and ultrawide bandgap semiconductors" where China holds 42.6% compared to the U.S.'s 17.6%, and its narrowest margins in "Advanced integrated circuit design and fabrication" at 24.4% versus 22.5% and "Natural language processing" at 24.8% versus 24.1%.

While China's highest citation share appears in "Machine learning" at 36.5% and "Adversarial AI" at 31.1%, reflecting rapid growth in foundational AI research, the US displays relatively stronger positions with shares above 20% in "High-performance computing" and "Adversarial AI." However, it falls below 15% in areas like "Advanced radiofrequency communication" and "AI hardware accelerators." These patterns underscore China's expanding leadership in core AI, ML, and semiconductor disciplines over this period.

**Figure S4.1 China vs. US Share of Top-10% Highly Cited Papers in Advanced Tech Areas**



Notes: Percentages are derived from five-year aggregates of the share of highly cited papers per country for nine key computing-, AI-, telecommunications, and semiconductor-related categories. These are Advanced radiofrequency communication, which encompasses research on next-generation wireless links, such as 5G-and-beyond technologies and millimeter-wave transceivers; High-performance computing, which involves the design and deployment of large-scale computing systems—supercomputers, exascale clusters, and associated middleware—for scientific, engineering, and AI workloads; Wide and ultrawide bandgap semiconductors, which focus on materials such as silicon carbide (SiC) and gallium nitride (GaN) that operate at high voltages, temperatures, and frequencies, enabling applications in power electronics, RF amplifiers, and electric-vehicle inverters; Advanced integrated circuit design and fabrication, which covers techniques for architecting, laying out, and manufacturing modern chips, including advanced process nodes, 3D-IC assembly, design-for-manufacturability, and EUV-lithography-enabled patterning; Advanced data analytics encompasses methods for large-scale data processing and insight extraction, including statistical learning, real-time streaming analytics, graph analytics, and big-data infrastructure such as Hadoop and Spark; AI algorithms and hardware accelerators, which involve the development of specialized neural-network architectures, compiler and runtime support, plus custom silicon including GPUs and TPUs, and ASICs to speed up AI training and inference; Machine learning, which represents core ML model research spanning supervised, unsupervised, and reinforcement-learning methods, with focus on model architectures, optimization algorithms, and benchmark evaluations; Adversarial AI, which studies ML vulnerabilities such as adversarial examples and poisoning attacks, along with defenses including robust optimization and certified guarantees to harden models against malicious inputs; Natural language processing, which encompasses algorithms and systems for understanding and generating human language, including language modeling, machine translation, question answering, and conversational agents.

Source: Gaida et al. (2023).

Taking a step back from the headline numbers: by 2025, China's strategy had evolved into a deliberate bid to leverage its state-led model to move faster than the U.S. and exploit its history of leapfrogging American-led technologies, as it did when TikTok redefined the U.S.-dominated social media landscape (Chin and Huang 2025). Consider the DeepSeek model's release in early 2025, which was competitive with US models but factor magnitudes cheaper to train, "jolted American competition into overdrive" and became the focal point of Beijing's national AI plan. Premier Li Qiang reportedly told officials, "China finally has a model it can be proud of," and President Xi Jinping personally convened a meeting with DeepSeek's founder to mobilize the industry (ibid.).

### **Contextualizing the 2025 AI Race**

In this section of the appendix, I trace how the U.S. government adopted a more activist, albeit often ad hoc, approach to innovation in the wake of the Chinese challenge to its leadership in innovation and AI. That strategy included several measures that would have been verboten during the CPD era. A more interventionist stance, while not abandoning the core tenets of the CDP, elevated the strategic importance of semiconductors and A.I. For example, the 2022 CHIPS and Science Act invested roughly \$280 billion in domestic semiconductor manufacturing, R&D hubs (via the National Semiconductor Technology Center), and workforce development—explicitly to counter China's chip drive and ensure U.S. firms could design, build, and secure the AI accelerators needed to power both commercial and defense applications. Washington, D.C. also experimented with various export controls on advanced technologies, particularly those related to AI and semiconductors, and greater scrutiny of foreign investment in sensitive sectors, as well as both import and export taxes.

Meanwhile, Europe struggled to develop more effective strategies for technological leadership, reflecting the stormy outlook it faced circa 2025. In the past 15 years leading up to that moment, its economies largely stagnated; Germany's output was only about 1% above its 2017 level that year, while the U.S. was up about 19%. Europe's share of global GDP (in current USD) slid from roughly one-third to about one-quarter since 2005, and per-capita income and household wealth growth fell well behind the U.S. (Luhnow and Fairless 2025). The continent's slippage shows up beyond spreadsheets: in 2025 the U.S. executed 100-plus orbital launches, China 40-plus—Europe only four, after a year of relying on SpaceX to loft critical infrastructure (ibid).

To boot, the continent also faced serious structural headwinds. These included an aging population shrinking the future workforce, welfare states under strain, higher tax burdens, and red tape that has doubled in volume since 2010. Energy prices—crucial for compute-heavy AI—were multiples of U.S. industrial power costs in 2025 (roughly 3× in Germany), undercutting siting for data centers and advanced manufacturing (ibid).

Moreover, the EU's fragmented market, risk-averse culture, and heavy regulatory burdens meant that it had fallen significantly behind the US and China in the AI race. In 2025, Europe lacked anything near the deep bench of firms, both established and startups, populating the AI supply chain and delivering innovative models and applications. Therefore, Europe produced far fewer scale players, drew a fraction of U.S. venture capital, and struggled with fragmented markets and rules—factors that weighed on the translation of research talent into globally dominant firms

(Luhnow et al. 2025; Fairless and Luhnow 2025). By contrast, the U.S. benefited from robust, deep VC markets, streamlined regulation, strong stock-option culture, and CHIPS Act support that re-energized semiconductor manufacturing.

Indeed, to be totally frank: in 2025 Europe was militarily weak, had a stagnant economy with an almost non-existent AI sector, moved sluggishly due to the constant need for compromises baked into the European Union's Byzantine political structure and possessed a relatively open economy that was extremely vulnerable to tariffs and other trade and geopolitical disruptions (see Authers 2026).

Meanwhile, China sought to obtain global leadership not only in A.I. technologies such as surveillance and facial recognition (Wei 2021), but also in renewable energy technologies (solar panels, wind turbines, EVs). Indeed, it explicitly pursued technological self-sufficiency through major investments in advanced manufacturing, robotics, semiconductors, and aerospace (Lee 2018). China's state-led "Made in China 2025" and "self-reliance" campaigns channeled \$500 billion in annual R&D, aggressive VC funds, and integrated supply-chain scaling. Chinese integrated circuit startups raised tens of billions of dollars in financing (Mims 2025c; Gehlhaus et al. 2023). State-owned enterprises were able to subsidize R&D for relatively long periods, further supporting this catch-up effort (Mims 2025c). Beijing also made massive public investments in infrastructure, including its 5G network, and asserted increased control over private investment (Naughton 2021).

Relegating attention to AI, Chinese high-tech firms were able to rapidly scale by exploiting access to vast datasets, the country's vast talent pool, and an enormous domestic market. In a bid to deploy AI across education, transportation and healthcare, including in diverse applications that ranged from fraud detection to crime-trend prediction to military hospital tele-health systems, Beijing pursued a top-down approach that compelled public sector players to exploit vast troves of government data using Chinese created and run open-source models.

Beijing's AI+ initiative, set out in the 2024 Government Work Report, explicitly aims to embed AI across manufacturing, services, and the public sector to raise productivity (Li 2024). This effectively turned infrastructure, factories, hospitals and public offices into living AI testbeds (see Northrup 2025). On the roads, regulators granted the first approvals for public trials of advanced autonomous-driving features to nine automakers, shifting from small pilots to broader deployment (Reuters 2024b). And beyond megacities, national "smart agriculture" competitions—run with FAO technical guidance—channeled AI into greenhouses and extension services, a visible sign of diffusion into rural practice (FAO 2021).

Similarly, circa 2025, a steady grafting of models into services people already use proceeded at a rapid clip in China. Inside the super-app WeChat (Weixin), Tencent has begun beta-integrating outside and in-house models so that AI search, summarization, and planning could happen in everyday chat, payments, and mini-programs (Reuters 2025d). On the consumer side, Alibaba's Quark emerged as a mass-market assistant, with a health model benchmarked at Deputy Chief Physician level on Chinese medical exams (Chen 2025). In clinical settings, approvals of AI-enabled medical devices moved from pilot to routine use; by mid-2023, China's regulator had cleared 59 AI medical devices, with radiology the leading application (Liu et al. 2024).

This rest of this section of the appendix compares the different approaches pursued by the United States, China, and the E.U. during the early 2020s and is especially keen to identify their similarities around increased mercantilism, protectionism, and statism. While the United States retained some elements of the CDP, its economic model increasingly became more state directed, populist, and mercantilist. Meanwhile, the EU was positioned to employ comprehensive, top-down regulation and China showed strong signs it would continue to consolidate its state-centric approach.

#### THE GLOBAL CONTEXT: DEGLOBALIZATION, PROTECTIONISM, AND A.I.

In the United States, after 9/11, both Republican and Democratic administrations recognized that an unregulated global financial system posed security risks, as terrorists and rogue states like North Korea could easily move money across borders (Helleiner and Kirshner 2014). In response, the Treasury Department transformed the U.S. dollar's global role into a system of power through several mechanisms. The Treasury developed increasingly sophisticated sanctions capabilities, progressing from targeting specific individuals and banks to eventually being able to isolate entire countries from the global financial system (Zarate, 2013).

This evolution represented a deliberate, bipartisan effort to reshape the global economy around U.S. security interests. What began as targeted interventions against terrorism gradually developed into a comprehensive system of economic leverage, with officials across administrations systematically building upon the tools they inherited from their predecessors.

Global economic integration began to erode significantly after the 2008 financial crisis, both in terms of trade volumes and cross-border financial flows (Hopewell, 2021).<sup>148</sup> The rise of China, growing geopolitical tensions, and the intensifying race for technological leadership, especially in AI and semiconductors, prompted a significant shift that accelerated in the early 2020s. Companies and governments shifted their focus from pure cost optimization to economic security, leading to increased onshoring, nearshoring, or “friend-shoring” of key technologies.

During the Obama administration, Washington, D.C. began to impose targeted tariffs on Chinese goods in retaliation for “dumping.” In September 2009, President Obama imposed a 35% tariff on Chinese tires after determining that import surges were harming domestic manufacturers, a measure that gradually decreased over three years before expiring in 2012 (Hufbauer and Lowry 2012). The administration later implemented significant anti-dumping duties on cold-rolled, hot-rolled, and corrosion-resistant steel (USTR 2015). By 2016, the Obama administration was enforcing 184 orders specifically targeting steel imports and had filed multiple WTO complaints challenging China's export restraints on rare earth minerals and duties on American cars (Scott 2016).

Under President Trump's first term, US policies around trade and capital flows took a stronger turn towards protectionism. The first Trump administration implemented export bans on advanced chips to specific Chinese firms, imposed tariffs on goods from various countries,

---

<sup>148</sup> Trade volumes decreased by 15 percent between the third quarter of 2008 and the second quarter of 2009. This was the steepest fall of world trade in recorded history and the deepest fall since the Great Depression (see WTO 2009).

foremost China, and increased the use of dollar-based sanctions (Congressional Research Service 2021). Unsurprisingly, world trade relative to global GDP fell by five percentage points between 2008 and 2019.

President Biden indulged in protectionism and economic nationalism as well. He kept Trump’s tariffs in place and adopted his own set of export controls: his administration significantly expanded technology restrictions, implementing three rounds of semiconductor export controls between 2022-2024.

In response, countries around the world, including the US, increased trade barriers as they tried to beat back a torrent of electric cars, industrial metals, chemicals, and other manufactured goods that threatened their homegrown industries (see Douglas and Fairless 2025). South Korea and Vietnam imposed tariffs on Chinese steel. Indonesia imposed duties on nylon used in packaging imported from China and other countries. India levied anti-dumping duties of 30% on Chinese steel and 25% on chemicals. Perhaps the most promiscuous version of retaliation was against Chinese EVs. Both the US and EU imposed steep tariffs on Chinese EVs. More generally, Russia introduced similarly high tariffs on all imported vehicles. Moreover, several countries mirrored the US and enacted export controls on strategic goods and actively promoted domestic industries through substantial subsidies (Baldwin and Freeman 2022; Bown 2020).

### **Globalization’s Freefall**

Taken together, these protectionist policies and the reactions to them took an appreciable toll on globalization. Had they continued their 2000–2008 trajectory (CPB Netherlands Bureau 2024), world trade volumes would have been nearly twice as high by the end of 2024. Between 2010 and 2025, annual volume trade increases averaged just 2.7%, far below the pre-2008 pace—and even as the global economy expanded by roughly 30% since 2008.<sup>149</sup> The plight of cross-border financial flows during this period was similar. While banks’ global cross-border claims reached \$41 trillion by the third quarter of 2024, up 3.4% year-on-year (BIS 2024), this growth barely kept pace with inflation. Adjusting for global GDP growth, major banks’ foreign exposure at the end of 2024 represented a smaller share of economic output than the \$30.4 trillion recorded in early 2008.

But it was not until 2025 that globalization truly crashed. On April 2, 2025, the Trump administration unveiled sweeping tariffs targeting over 90 nations, declaring a “national emergency” under the International Emergency Economic Powers Act (IEEPA) to bypass congressional approval.<sup>150</sup> The policy introduced a 10% baseline tariff on all U.S. imports effective April 5, 2025, with higher reciprocal tariffs for approximately 60 countries—including China (34%), Vietnam (46%), and Cambodia (49%). Combined with existing 20% tariffs on

---

<sup>149</sup> Ignoring comparisons to the “globalization counterfactual” for a moment, global trade in goods reached a peak in 2022; trade in services did so in 2024 (UNCTAD 2024); that year overall global trade notched a record \$33 trillion.

<sup>150</sup> The administration framed these measures as necessary to rectify \$1.1 trillion in annual U.S. trade deficits and combat “nonreciprocal” practices like value-added taxes (VATs) and currency manipulation.

Chinese goods, China's total tariff rate exceeded 60%.<sup>151</sup> While Trump inherited an average effective tariff rate of around 2%, these so-called Liberation Day tariffs lifted it to 23% and covered far more economic activity than the infamous 1930 Smoot-Hawley tariffs; moreover, the latter tariffs only raised the average rate by 6 percentage points, albeit this eventually became 19 percentage points in real terms (see IP 2025).

Moreover, Trump's "Liberation Day" was met with swift retaliatory measures by US trading partners: Chinese counter-tariffs of 34% on all U.S. goods, €18 billion in levies by the EU, and 25% tariffs on U.S. vehicles imposed by Canada. Taken together, these protectionist measures essentially dismantled the post-WWII rules-based trade system, replacing it with a fragmented regime that prioritizes unilateral penalties over multilateral cooperation (Barath et al. 2025).<sup>152</sup>

### **The Return of Vertical Integration and American Reshoring**

As national security concerns increasingly took precedence over aspirations for deeper globalization during the first quarter of the 21<sup>st</sup> Century, this had a profound effect on the American economy. Even before the Covid-19 pandemic, policymakers questioned the resilience of American supply chains (Farrell and Newman 2019). After the pandemic, there was a concerted effort to reshore chipmaking to the US and away from Taiwan, South Korea, and Japan (Fuller et al. 2020). Tighter immigration rules and export controls hampered cross-border collaboration and slowed technological diffusion on the global scale (Villasenor 2025).

Vertical reintegration across major industries went beyond chips; during the early 2020s, many of the tightly integrated global supply chains that blossomed during the CDP era, as documented earlier in the book manuscript, began to unravel (Alicke et al. 2023). General Motors invested heavily in battery plants. Tesla moved downstream into battery production Gigafactories and direct agreements with mining companies for raw materials. Apple increased self-reliance by designing its own chips, while Amazon built a comprehensive logistics network including fulfillment centers, delivery vehicles, and Amazon Air. Samsung began to manufacture many of its own components, from displays to memory chips. Boeing increased in-house parts manufacturing and maintenance services. PepsiCo acquired many bottling facilities. Disney vertically integrated media operations, creating content through its studios and distributing it through Disney+. And, as Chapter 9 of the book showed, this even included the AI stack itself.

A wave of American reshoring was driven by a convergence of technological and economic changes. Rapid advances in automation, AI, and digital logistics began to erode the core economics of offshoring. These forces were amplified by narrowing energy cost gaps and shifting market dynamics. As a result, U.S. firms increasingly found domestic production and regional supply chains not only feasible but economically advantageous.

---

<sup>151</sup> Exemptions were given for USMCA-compliant Canadian and Mexican goods, pharmaceuticals, semiconductors, and energy imports.

<sup>152</sup> In terms of export controls, the second Trump administration blacklisted an additional 80 entities vis-à-vis Biden, including over 50 Chinese companies involved in AI and quantum computing technologies.

First, the rapid advancement of automation and AI began to erode the core economic logic of offshoring. This was especially the case for routine, low-skill “pick-and-place,” inspection, and simple assembly operations that had been traditionally sent overseas. As the unit labor cost benefit of offshoring shrank, domestic production became more viable. Therefore, U.S. firms announced a record number of reshoring manufacturing jobs in 2022, citing automation as a key driver (Reshoring Initiative 2023).

In parallel, new technologies like additive manufacturing (3D printing) enabled a shift toward on-demand, on-site production, with investments in the sector accelerating rapidly in the early 2020s (Bramberger et al. 2022). Furthermore, digital logistics—using real-time tracking, digital-twin modeling, and automated warehousing—made geographically closer supply chains more cost-effective by reducing the need for large, precautionary inventories (Manenti 2021).

This technological shift was not confined to manufacturing; the AI-driven automation of services also had a direct impact on offshoring. By 2023, AI-powered virtual assistants were already saving global enterprises billions of dollars in customer service costs (Gatford 2025). Major outsourcing firms began embedding AI into their platforms, reducing the need for remote human agents for standard inquiries.

These factors were compounded by domestic economic advantages. This included the widening energy cost gap; in 2023, U.S. benchmark natural gas prices were roughly one-third of those in Europe, significantly narrowing the cost differential that once favored offshore sites (see EIA 2024; Eurostat 2024). Indeed, manufacturing companies planning to relocate supply chains to North America cited not only automation but energy costs and market proximity as their primary reasons (Reshoring Initiative 2023).

### **Digital Balkanization**

Beyond tariffs and export controls, the drive toward deglobalization leading up to 2025 and beyond also extended to data governance. The major economic blocs—the U.S., EU, and China—developed divergent technology standards and data policies, contributing to increasingly splintered tech ecosystems (O’Hara and Hall, 2018). Data sovereignty rules proliferated, with governments mandating local storage of personal or sensitive data. China, for example, began to treat data as a national resource, imposing strict localization rules (Yang, 2021). This created higher compliance costs for A.I. developers and limited access to data and key technologies critical to the A.I. supply chain. It also forced global companies to maintain distinct A.I. models for different markets.

### **AMERICAN INDUSTRIAL POLICY**

Economic populism was practiced by both the first Trump administration and the Biden administration in the form of dismissing budget constraints, eschewing cost-benefit analysis, ignoring tradeoffs and implementing pro-cyclical fiscal policies (see Magistro and Menaldo 2023). For example, President Biden ran the economy hot by passing a massive, \$1.9 trillion stimulus measure in 2021 that accelerated inflation and did not improve real living standards: the economic recovery from the COVID-19 crisis began in mid-2020, and real GDP growth was a

relatively strong 5.6% in the first quarter of 2021, before the American Rescue Plan funds had been digested by the economy.<sup>153</sup> In December 2024, the unemployment rate was roughly 4%, above the 3.5% before the pandemic; inflation remained above its target; and inflation-adjusted wages were barely above pre-pandemic levels, as the entire increase in real wages had taken place in 2020—indeed, real wages had fallen since January 2021.<sup>154</sup>

Similarly, beginning with the first Trump administration, the US also resurrected the type of technological mercantilism practiced by erstwhile great powers.<sup>155</sup> Like during the Third Industrial Revolution, the federal government showered major A.I. supply chain players with subsidies and tax incentives and direct R&D funding in areas such as quantum computing and semiconductors. Unlike that era, policymakers pursued a proactive, coordinated industrial policy. To be sure, in Chapter 10 of the book I outlined how private sector “hyperscalers” like Google and Microsoft continued to make substantial AI investments. This is only part of the story: the federal government doled out muscular financial and logistical support through defense funding and research grants.

Moreover, the Trump and Biden administrations wielded various economic weapons and tools in their strategic approach. Tariffs were used aggressively against both allies and adversaries, while sanctions served as a core mechanism to influence geopolitical outcomes. Export controls restricted sensitive technologies, particularly in semiconductors and AI. The global role of the U.S. dollar and control of the SWIFT international payments system provided Washington with significant leverage, though adversaries began developing workarounds to these financial constraints.

### **The Transition from the CDP to Mercantilism**

---

<sup>153</sup> On all these points see Furman (2024).

<sup>154</sup> Consider that from 2020 to 2024, average real wage growth for workers in every income group was slower than it was from 2014 to 2019. Rapid real wage growth, especially for low-income workers, began in 2014, when the unemployment rate was around six percent, but drastically slowed when the unemployment rate fell below four percent in 2022.

<sup>155</sup> A Hamiltonian defense of protectionism holds that tariffs on imports can benefit infant industries by helping manufacturers overcome both entrepreneurial risk aversion and the advantages foreign firms receive from their governments. This strategy proves especially effective when protection enables a domestic industry to achieve economies of scale by serving both home and foreign markets—implying tariffs on competing imports while exploiting importers’ freer trade regimes—allowing it to become internationally competitive once it substantially reduces its average costs (see Pack and Westphal, 1986). However, by 2025, this rationale has largely become obsolete for the U.S., which has evolved into a post-industrial service economy characterized by high productivity, high wages, and already-competitive advanced manufacturing in sectors like aerospace and technology. The contemporary challenges to American manufacturing employment stem primarily from competition with low-cost overseas labor and increasing domestic automation, rather than any fundamental inability to achieve economies of scale in developing industrial sectors.

Beginning with the first Trump administration (2017-2021), the U.S. economic model became increasingly state-directed and mercantilist. This included restrictions on Chinese investment in the American economy, along with bans against outbound FDI headed to China and certain U.S. goods exported to China. In August 2018, the U.S. Government passed the Foreign Investment Risk Review Modernization Act; it was (at least) partly intended to reduce Chinese FDI in areas that are deemed sensitive to U.S. national security. Exploiting export controls built up by the Commerce Department's Bureau of Industry and Security under the Obama administration, the U.S. government blocked Chinese tech giant Huawei from federal contracts and prevented American chipmakers from supplying it with essential components.

The Trump administration also aggressively tinkered with international supply chains for high-tech goods. It pressured non-Chinese telecommunications providers like Nokia and Ericsson to relocate their supply chains outside China, citing security concerns. Some government officials even called for nationalizing critical infrastructure like 5G networks.

American policymakers also expressed serious concerns about Chinese companies recruiting talent from Silicon Valley. Foreign executives operating in China worried about IP theft by their own employees, a view that was shared by the U.S. government. Critics argued that companies like Huawei serve as potential backdoors for Chinese government data access, while some analysts warned that China's 5G influence could enable it to shape telecommunications laws and regulations globally, potentially imposing its own version of internet governance. These concerns led to the Federal Communications Commission (FCC) designating both ZTE and Huawei as national security threats, banning their equipment from American wireless networks, and ending related federal subsidies.<sup>156</sup>

Beyond technology concerns, American politicians raised several red flags about China's mercantilist economic strategy when justifying the US's own version. It encompassed Chinese tariffs on American imports, Beijing's alleged currency manipulation, and its subsidies for state-owned enterprises and flooding of international markets with cheap industrial goods like steel and solar panels. It also included China's alleged IP theft. While these grievances fueled President Trump's so-called trade war with China, which ushered in 10% tariffs on all Chinese goods, his administration also pointed to its desire to narrow American trade deficits—a populist rallying cry that had some grounding in veritable economic problems.<sup>157</sup>

---

<sup>156</sup> This strategic push was further intensified in September 2025, when the administration closed a major loophole in its trade blacklist. The new rule automatically subjected any majority-owned subsidiary of a blacklisted company, such as Huawei, to the same U.S. trade restrictions (Ramkumar 2025b). This move, however, also highlighted what some industry observers described as an "inconsistent" U.S. policy: While national-security hawks in the administration championed the new rule, other officials had previously allowed the continued export of some high-performance AI chips to China, reflecting an ongoing internal conflict between security concerns and commercial interests (ibid.).

<sup>157</sup> Obstfeld and Rogoff argued that the 2008 Global Financial Crisis was exacerbated by a global savings glut, driven by China's trade surplus and accumulation of U.S. treasuries, which depressed long-run interest rates (Obstfeld and Rogoff 2009). This supposedly encouraged the creation of high-yield but riskier asset classes, including mortgage-backed securities and

Moreover, the impact of U.S.-China economic integration on American labor was substantial, particularly in the rustbelt regions like Ohio, which historically housed heavy industry including steelmaking, automobiles, appliances, machinery, and chemicals. While Freeman suggested global trade only modestly reduced employment and wages among U.S. low-skilled workers (Freeman, 1995), Acemoglu and colleagues estimated that increased import competition following China's 2001 WTO accession led to losses of 2.0 to 2.4 million U.S. manufacturing jobs between 1999 and 2011 (Acemoglu et al., 2016). Autor, Dorn, and Hanson noted these effects concentrated in traditional manufacturing areas, with remarkably slow labor market adjustments to trade shocks, resulting in persistent income reductions and employment insecurity, especially for unskilled workers (Autor, Dorn, & Hanson, 2013).

Other reasons for American mercantilism and protectionism arose in the latter part of the first Trump administration. During the COVID-19 pandemic, there were calls for the U.S. to develop resilience to potential supply chain shocks, especially has become apparent, particularly in critical sectors like medical equipment, semiconductors, and critical minerals. Additionally, the premise that economic integration would lead nations to become more responsible global actors had seemed to prove false, particularly regarding China and Russia. Policymakers also voiced concerns that China could achieve dominance in strategic industries such as rare earths, semiconductors, and solar panels, which could potentially make the U.S. economically and militarily dependent on Beijing.

Finally, China Hawks argued that China sought to use its dominance over some high-tech hardware and software to foist its ideological, and potentially totalitarian, version of the internet on the global community (see Coughlan 2020). This included China's dominance over 5G—and its associated standards, platforms, and patent pools, with potential upstream influences over global telecommunications laws and regulations. Thus, Huawei and TikTok were characterized as political entities—an extension of an assertive, authoritarian state bent on spreading propaganda and using international standard setting boards to hijack national laws and promote surveillance (Rosenberger 2020a; Rosenberger 2020b).

### **President Biden's Overt Industrial Strategy**

The Biden administration's industrial strategy marked a significant break from Democratic economic policy since President Carter and continuing under Presidents Clinton and Obama, which had embraced market mechanisms and incremental welfare expansions. Instead, it pursued a more muscular government role in direct spending, industrial policy, and regulation. Biden's "Build Back Better" agenda envisioned large-scale federal investment, revived domestic manufacturing, new climate initiatives, and reduced reliance on global supply chains.

Unsurprisingly, therefore, the Biden administration continued and expanded many of his predecessors' economic policies. This included retaining Trump's tariffs on imports from various countries, especially China. Biden also implemented comprehensive sanctions targeting China's

---

collateralized debt obligations. When American homeowners defaulted on variable-rate mortgages, these assets' values plummeted, devastating bank balance sheets and triggering a credit crunch leading to the Great Recession, though this causal chain remains debatable.

high-technology sector, restricting access to A.I. chips designed by American companies and manufactured by TSMC and Samsung, while also blocking access to chip design software. The Biden administration blocked exports of AI-capable chips in October 2022, restricting Nvidia's A800 chip in 2023, and adding 140 Chinese firms to the Commerce Department's Entity List by December 2024. It also introduced export limitations on advanced memory chips and semiconductor manufacturing equipment, which prevented China from obtaining chip-making tools from Dutch, American, and Japanese sources, as well as restricting access to expertise in design, equipment, and manufacturing processes.

The Biden administration also established new controls on FDI, regulating both U.S. investments in China and Chinese investments in America, while limiting Chinese tech platforms' ability to collect data from U.S. consumers.<sup>158</sup> These actions built upon policies initiated under the Obama and Trump administrations, particularly the decision to exclude Huawei from the U.S. 5G network infrastructure.

The CHIPS and Science Act (henceforth the CHIPS Act), a cornerstone of Biden's industrial policy, aimed to strengthen American semiconductor independence by reducing reliance on both Taiwan and China. The legislation supported domestic production through various incentives, including subsidies for chip fabrication and tax credits for companies establishing U.S. manufacturing facilities using American-made equipment. The CHIPS Act allocated \$39 billion in semiconductor subsidies, including up to \$8.5 billion for Intel. The Act also invested significantly in research and development to advance semiconductor technology.

The Inflation Reduction Act (IRA) represented another major Biden initiative, encompassing broad investments in sustainable technology and domestic manufacturing. It focused on minerals, green energy manufacturing, batteries, renewable energy production, energy efficiency, hydrogen, biofuels, and carbon capture. The legislation provided tax credits for locally manufactured battery cells and modules, including Buy American requirements for raw materials. This was complemented with subsidized electricity. It also included significant labor provisions, requiring companies who secured subsidies to establish union-run apprentice programs and adopt union-scale wages. The Inflation Reduction Act expanded the DoE's lending authority to over \$400 billion for clean energy projects. By 2025, intensive construction efforts had begun in regional manufacturing clusters in states like Ohio, Michigan, North Carolina, and Arizona.<sup>159</sup>

## **President Trump 2.0**

---

<sup>158</sup> Under President Biden, the Committee on Foreign Investment in the United States, which reviews bids for foreign investment in U.S. companies, adopted a much more Draconian approach. For example, it informed the Biden administration's decision to block Japan's Nippon Steel from acquiring U.S. Steel in 2024.

<sup>159</sup> Researchers projected that U.S. emissions will be roughly 17% lower by 2050 than if there were no IRA. However, others estimate that a carbon tax of \$12 a ton would yield the same emission reductions as the IRA *without* relying on costly subsidies that cannot be scaled up and introducing economic distortions, such as crowding out private investment.

While President Biden’s muscular attempts at industrial policy yielded mixed results, early in the second Trump administration an “America First” approach came into view. It was rooted in trade protectionism and similar forms of mercantilism, continued support for American semiconductor manufacturing, scaling back direct support for green energy, and generally oriented towards deregulation.<sup>160</sup> The second Trump administration also emphasized domestic A.I. development free from heavy regulation.

During the Trump administration’s first month in power, Vice-president J.D. Vance delivered a speech during a European AI conference that broke sharply from previous international talks about AI. Rather than emphasize A.I. “guardrails” or “equity,” Vance focused instead on “opportunities.” He explicitly distinguished the U.S. position from the themes of conferences such as the U.K.-hosted A.I. Safety Summit that highlighted potential catastrophic risks of advanced A.I. Vance criticized the European Union’s Digital Services Act for its sweeping regulations on disinformation, suggesting that it stifles free expression and hampers technological progress. In short, he projected a more risk acceptant American attitude and urged European countries to follow suit.

Though President Trump publicly attacked the 2022 Chips Act and instead advocated 100% tariffs on foreign produced chips to promote domestic semiconductor manufacturing, his administration did not abandon the act’s core subsidies. In an effort to reduce regulatory burdens and court both domestic and foreign capital, it housed a so-called Investment Accelerator within the Commerce Department (known as the USIA) to oversee and enhance the existing Chips framework.<sup>161</sup> The executive order bringing this project into existence emphasized securing

---

<sup>160</sup> The Biden administration’s attempt to spark a “manufacturing renaissance” through subsidies and strict procurement policies yielded mixed results. On one hand, major investments in semiconductor plants were promised, including an Intel plant in Columbus, Ohio; a Micron chip plant in upstate New York; and a TSMC plant outside Phoenix. Automotive companies, including Honda and GM, also made substantial EV and battery plant investments in the Midwest. Factory-related construction spending reaching \$108 billion in 2022, primarily focused on semiconductors, EVs, and batteries. On the other hand, overall manufacturing output and employment did not see a broad rebound, and investment in industrial equipment stagnated. Unsubsidized sectors struggled as input costs and the dollar’s value rose in tandem with interest rates, which drove up the prices of materials and equipment, wages for construction and factory workers, and borrowing costs for entrepreneurs. Biden’s climate legislation provided large corporate subsidies for renewable energy and domestic green-tech production but offered fewer benefits to lower-income households than a carbon tax with rebates would have provided. Unionization rates fell below ten percent in 2024 for the first time ever. Infrastructure investment faced significant challenges too. Despite ambitious plans for upgrading roads, bridges, and broadband, the reality on the ground was less spending on real infrastructure due to soaring construction costs, labor constraints, and slow permitting. On all these points see Lawrence (2024).

<sup>161</sup> On all these points see Secreto (2025), who argues that this rebranding allowed President Trump and "America-first" conservatives to support semiconductor incentives without crediting President Biden or Commerce Secretary Gina Raimondo, the Chip Act architects. According to Ip (2025), the second Trump administration’s approach also allowed the president to leverage

"better deals" and declared that regulatory relief—not just cash subsidies—must accompany these types of initiatives.<sup>162</sup>

The neo-mercantilism applied to other policies too. Besides imposing onerous tariffs on imported chips and personally determining and adjusting tariff levels on imports from nearly every country worldwide, the second Trump administration also saw to it that the federal government require Nvidia and AMD to remit 15% of their China AI chip sales to the U.S. Treasury—effectively making the U.S. a business partner in their operations—as a condition for export licenses. This complemented Washington, D.C.'s ability to secure "golden shares" and board representation in strategic transactions like Nippon Steel's acquisition of U.S. Steel and the Pentagon becoming MP Materials' largest shareholder with a 15% equity stake to control rare earth supply chains (see Ip 2025).

This brand of state centric capitalism subject to the changing whims of the sitting executive also governed the ins and outs of Nvidia's next-generation scaled-down chip tailored for the Chinese market, the B30A. While President Trump had at first hinted that he might allow its sale (in contrast to his ban on top-tier Blackwell chips), the White House changed course and blocked its export (Liu 2025).<sup>163</sup> The move was a major blow to Nvidia, which had already sent samples of the chip to Chinese customers (ibid.).

Trump further intervened in more idiosyncratic ways. He urged Intel CEO Lip-Bu Tan to resign via a Truth Social post over his investments in China, only to reverse course and praise him as a "success" after a White House meeting; issued repeated executive orders to extend TikTok's U.S. operations despite a law requiring its Chinese owner to sell or cease activities; instructed Walmart to "EAT THE TARIFFS" and absorb costs without passing them to customers in response to their warnings about price increases; and signed an executive order mandating pharmaceutical companies to match U.S. drug prices to those in other industrialized nations, with threats to impose them if not complied with.<sup>164</sup>

Even before the inauguration of the second Trump administration, industry leaders began to bypass traditional lobbying in favor of direct relationships with Trump, often making financial

---

domestic chipmakers' dependence on federal subsidies to pressure their leadership; for example, when Trump demanded Intel CEO Lip-Bu Tan's resignation over alleged China ties.

<sup>162</sup> Tariffs alone could not bridge the steep cost difference between building a chip fab in Phoenix versus Taipei. Even with funding secured, in 2025 delayed permits and labor shortages represented massive obstacles: TSMC's Arizona campus, a \$65–\$100 billion multi-fab project, received \$6.6 billion in direct federal funding and \$5 billion in loans, but still faced multi-year delays. To address cost challenges, the government finalized over 30 major CHIPS Act grants, allocating funds to industry leaders including Intel, Micron, Samsung, and TSMC, with awards tied to construction and production milestones and a new wave planned for specialty suppliers. To tackle permitting backlogs, the USIA coordinated agency action and sped up the path to approval for these projects. On all these points see Stansbury 2025.

<sup>163</sup> Unlike the older H20 chip, which was largely limited to inference, the B30A was powerful enough to be clustered for large-scale model training.

<sup>164</sup> On all these examples see Bendavid and Bhattarai 2025.

commitments or symbolic gestures to win his favor. To that point, circa early 2025, Trump claimed personal oversight of approximately \$1.5 trillion in foreign investment pledges (Ip 2025)—including TSMC’s promise to invest \$100 billion (see Mickle 2025). He also “secured” pledges by US firms to invest billions more. This included Apple’s promise to invest over \$100 billion after CEO Tim Cook presented Trump with a golden plaque (Bendavid and Bhattacharai 2025) and Micron’s \$150 billion (Mickle 2025).

## US Strategy Towards AI Standards

During the early days of the second Trump administration, the federal government signaled it would continue to rely on American leadership in semiconductors, large-scale cloud infrastructure, and frontier AI labs to sustain de facto technical baselines for AI safety, security, and interoperability (Carrillo et al. 2025; TechNode Feed 2025; JTC 1 n.d.).<sup>165</sup> Consider that, by virtue of their scale, U.S. hyperscale platforms, chipmakers, and AI developers could rely on their preferred model formats, API specifications, and safety benchmarks as industry defaults: As of 2025, U.S. operators accounted for roughly 54% of global hyperscale data-center capacity (Synergy Research Group 2025). While necessary to establish AI related standards, however, this market dominance by US firms was not sufficient, however.

Like earlier episodes in the digital revolution, when U.S. technologies such as TCP/IP, USB, 4G, and cloud security protocols became global norms, the National Institute of Standards and Technology, or NIST, again played a pivotal role in nudging private firms to set standards using methods and strategies it preferred. During the early 2020s, both its voluntary AI Risk Management Framework and the U.S. AI Safety Institute (AIS I) created a powerful incentive for tech companies to adopt guidelines around AI, not by legal force, but through market access (NIST 2023; NIST 2024).<sup>166</sup> For example, for an AI developer to sell a diagnostic tool to a hospital or a logistics model to the Pentagon—key public-sector and critical-infrastructure markets—they would need to provide standardized documentation, such as model cards that detail the AI’s training data, performance, and limitations (a documentation norm). They would also have to submit their models for independent evaluation against common benchmarks for safety, security, and bias (a testing protocol), often conducted through the AI Safety Institute. Furthermore, vendors had to build in mechanisms for ongoing monitoring and auditing to prove their systems remained reliable after deployment (an assurance practice).

Finally, and similarly, rather than treaty-based consensus, in the early 2020s Washington, D.C. was also deepening coordination with partners in more official ways. First, the U.S.–EU TTC’s

---

<sup>165</sup> To be sure, in 2025 the second Trump administration was actively rewriting parts of the rulebook it had inherited from the Biden administration, including narrowing NIST guidance and rebranding safety standards.

<sup>166</sup> Established within the NIST, the U.S. AIS I is the primary government body responsible for evaluating the risks and capabilities of the most advanced AI models. It develops technical guidelines, conducts evaluations like red team testing, and facilitates research to create a common understanding of AI safety and security. It brings together over 200 members from industry, academia, and civil society, and works closely with international partners, such as its counterpart in the United Kingdom, to reach convergence over safety standards and research.

Joint Roadmap on Evaluation and Measurement Tools for Trustworthy AI and Risk Management was building a shared repository of metrics and methods intended to feed directly into international standards (NIST 2022; Hasselbalch 2022; USTR 2024). Second, the G7 Hiroshima AI Process established voluntary guiding principles and a developer code of conduct that partners were able to embed in procurement and audits (G7 2023). Third, after the U.K.'s 2023 Bletchley Park AI Safety Summit, Washington and London signed a 2024 memorandum of understanding linking NIST's new AI Safety Institute with the U.K. AI Safety Institute. The agreement commits both governments to share red-teaming data sets, co-develop evaluation metrics, and accept one another's test results in national procurement. By creating a common library of benchmarks that feeds directly into ISO/IEC SC 42 work, the deal internationalizes U.S. technical baselines through mutual recognition (DSIT 2024).<sup>167</sup>

## CHINA'S INDUSTRIAL POLICY

Circa 2025, China was deeply integrated with the global economy. China's trade connectivity encompassed the vast network of bilateral and multilateral trade agreements, logistics infrastructure, and supply-chain relationships that link it to markets around the world. China served as the single largest trading partner for over 120 countries across every region, from Southeast Asia and Africa to Latin America and Europe and had signed free-trade agreements with more than 20 economies, including ASEAN, Chile, and Switzerland, reducing tariffs and non-tariff barriers on hundreds of product lines.

In terms of regional trade blocs, since early 2022 the Regional Comprehensive Economic Partnership (RCEP) covered 15 Asia-Pacific economies and accounted for nearly a third of global GDP. By harmonizing rules of origin and cutting tariffs, RCEP further cemented China's role at the heart of East Asian supply chains. After joining the WTO in 2001, China steadily expanded market access commitments, integrating its manufacturers and exporters into global value chains.

---

<sup>167</sup> To create a unified global framework for artificial intelligence, the International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC) established a joint subcommittee, ISO/IEC JTC 1/SC 42, as the focal point for international AI standards (ISO/IEC 2025). Through a "one country, one vote" consensus process involving over 60 national bodies, this group produces a toolkit of definitions, metrics, and management systems that underpins global AI safety, security, and interoperability. It operates through focused working groups covering the entire technology stack, from foundational concepts and big data to trustworthiness, use cases, and computational approaches. This work has produced a suite of flagship standards, including ISO/IEC 22989 for foundational terminology, ISO/IEC 23894 for risk management to complement national frameworks like NIST's, and the certifiable ISO/IEC 42001 for AI management systems, which allows organizations to prove their development processes are trustworthy. Ongoing projects, such as the multi-part ISO/IEC 5259 series, are further developing granular rules for data quality and governance. To ensure seamless integration, the committee coordinates with other standards bodies focused on areas like software engineering (SC 7) and health informatics (TC 215).

By the early 2020s, China's logistics and port infrastructure was extensive, with state-owned enterprises holding stakes or concession rights in over 100 ports across more than 60 countries—ranging from Piraeus in Greece and Hambantota in Sri Lanka to Djibouti and Felixstowe in the UK. This maritime footprint accelerated Chinese exports and provided preferred docking for its merchant fleet. High-speed rail links, such as China-Europe "iron silk road" freight trains, and trans-Asian highways slashed transit times for Chinese goods bound for Europe and Central Asia, diversifying transport options beyond maritime routes.

The Digital Silk Road component featured Chinese firms such as Huawei and ZTE building or upgrading 4G/5G networks in dozens of countries, often bundled with favorable financing, while investments in submarine cables and sovereign data centers expanded China's digital reach—enabling cross-border e-commerce, cloud services, and influence over information flows.

The Belt and Road Initiative (BRI), launched in 2013, became China's flagship global-infrastructure and connectivity project, often described as the largest state-led development program in modern history. By 2025, over 150 countries and 30 international organizations signed BRI cooperation agreements, covering roughly two-thirds of the world's population, with Chinese banks and policy lenders had extended \$1–1.5 trillion in loans and investments for BRI projects.

The initiative's core components included the "Overland Silk Road Economic Belt," which sought to build highways, railways, and pipelines linking China through Central Asia to Europe via major corridors like the China-Kazakhstan-Russia rail line and China-Vietnam highway upgrades. The "21st-Century Maritime Silk Road" focused on port construction and maritime logistics from Southeast Asia through the Indian Ocean to East Africa and the Mediterranean, with key ports including Gwadar in Pakistan, Hambantota in Sri Lanka, and Piraeus in Greece. The Digital Silk Road emphasized telecom networks, satellite systems, and e-commerce platforms to knit together China's digital ecosystem with partner economies.

### **A Short History of China's Mercantilism**

Paradoxically, however, since its accession to the WTO in December 2001, Beijing methodically sheltered China's domestic market behind high tariffs, opaque licensing rules, and steep state-owned-enterprise subsidies—even as it has poured subsidized goods into global markets. In the aftermath of the 2008 financial crisis, China doubled down on this approach: its 12th Five-Year Plan (2011–2015) earmarked hundreds of billions for key industries, and in May 2015 the government formally launched its "Made in China 2025" blueprint. By 2019, China was spending roughly 2.3 percent of GDP on direct industrial support—about ten times the OECD average of 0.2 percent.

Those policies propelled Chinese firms to world-leading positions in shipbuilding (since the early 2010s), drone manufacture (rapidly from 2014 onward), advanced electronics (over the 2010s and early 2020s), and bulk pharmaceuticals (notably during the 2020 COVID-19 scramble). Having achieved such scale, Beijing then "weaponized" its dominance: in September 2010, it abruptly halted rare-earth exports to Japan and, in 2019–20, threatened Europe's drug supply by restricting precursor chemicals; beginning in 2021, Chinese-made combat and

commercial drones swarmed Taiwan’s air defenses; and from late 2021 into 2023, heavily subsidized electric vehicles poured into Europe at rock bottom prices Western automakers could not compete with.

At home, the Golden Shield (“Great Firewall”) project—rolled out beginning in 2003—barred foreign news and social-media platforms, while Confucius Institutes (first established in 2004) propagated Beijing’s preferred narratives on campus. Abroad, Chinese intelligence services exploited open Internet architectures to steal IP (Operation Aurora in 2009, ongoing “hack-and-lead” campaigns in the 2010s), embed malware in critical infrastructure (notoriously the 2014 “Cloud Snooper” intrusion), and pump state-sponsored propaganda through thousands of fake accounts and bot networks.

This strategy of state-backed dominance, market flooding, and "weaponized" export controls was not limited to just one or two sectors. By 2025, Beijing had established a chokehold on several critical global supply chains, following a playbook that President Xi Jinping himself described as the need to "tighten the dependence of international industrial chains on our country" (Kubota 2025).

The lithium-ion battery supply chain is a prime example. While Chinese firms CATL and BYD became the world's top two battery producers by 2025, China's true dominance lay in processing the raw materials. By 2025, Chinese suppliers produced 79% of the world's battery cathodes and 92% of the anodes, while also controlling most of the refined lithium, cobalt, and graphite. Having secured this dominance, Beijing began imposing export licenses on battery manufacturing technology to prevent rivals from catching up (ibid.).

This pattern was repeated in mature semiconductors. While the U.S. focused on restricting advanced AI chips, China spent billions to capture one-third of the global capacity for older, "mature" chips, which during the early 2020s remained essential for cars, electronics, and defense. It also established control over key minerals like gallium (99% of global production) and Germanium, placing export restrictions on them in 2023 (ibid.). This leverage was put on full display in 2025 when Beijing blocked the export of mature chips from a single Dutch company, Nexperia, whose products were processed in China. The move, a clear retaliation against U.S. trade policy, forced global automakers like Honda to shut down factories within weeks, a vulnerability that was only resolved in the November 2025 trade truce (ibid.).

A similar chokehold exists in pharmaceuticals, where China had become the primary source for the active pharmaceutical ingredients (APIs) for common drugs like acetaminophen (Tylenol) and ibuprofen (Advil). By 2025, even generic drugs manufactured in India were often made with APIs imported from China. This dependence was highlighted in 2020, when the official Chinese news agency warned that if Beijing restricted medical exports, the U.S. would be "plunged into the vast ocean of coronavirus" (ibid.).

### **The Toolkit of State-Led Innovation**

Since coming to power, President Xi Jinping made “self-reliance” in key technologies a national imperative, arguing that China must “not be forced to beg others for technology” if it is to

safeguard its economic and national security (Spegele 2025). Beijing's stated overarching goal was to close the frontier-technology gap with the United States across areas such as semiconductors, artificial intelligence, robotics, and biotechnology. To do so, China maniacally and methodically marshaled an array of state-directed, application-focused industrial-policy tools.

Beijing's approach was anchored in two flagship plans: the 2017 New Generation Artificial Intelligence Development Plan (AIDP) and the broader "Made in China 2025" roadmap (see Webster et al. 2017). AIDP lays out milestones for world-class AI research, domestic chip design, and mass deployment across governance, industry, and defense. "Made in China 2025" seeks to replace foreign technology in ten strategic sectors—including semiconductors, robotics, electric vehicles, and biomedicine—boosting the self-sufficiency rate year by year. This obligates Beijing to back national champions (Baidu, Alibaba, Huawei, SenseTime, Cambricon) through subsidies, grants, and preferential procurement, driving an estimated \$85 billion investment wave in AI firms and chip startups (Castro and McLaughlin 2024). Militarily, China was integrating AI into drone swarms, autonomous ground vehicles, and command-and-control systems, aiming to close any gap with U.S. forces (Webster et al. 2017).

To fund these ambitions, Beijing mobilized vast resources. Chinese R&D spending reached about \$500 billion in 2024—roughly triple its 2012 level—and rivaled U.S. outlays on a purchasing-power-parity basis (OECD Statistics 2025). China launched massive investment vehicles to finance rising technologies. Notably, the National Integrated Circuit Fund (the "Big Fund") was established to invest in semiconductor and AI hardware firms (Yahoo Finance 2025). Similarly, provincial governments and state banks channeled loans and equity into AI startups that align with national priorities (Beraja et al. 2025). From 2000 to 2023, China's government-affiliated venture funds poured nearly \$200 billion into some 9,600 AI startups, while local investment arms back challengers like Zhipu AI alongside private giants such as Alibaba and Tencent (Castro and McLaughlin 2024).

This spending was enhanced by a suite of generous fiscal measures to spur corporate R&D, including tax relief, preferential rates, and accelerated write-offs. The centerpiece was a super-deduction for R&D spending, in place since 2007, allowing enterprises to deduct more than 100 percent of their qualifying outlays from taxable income. Initially set at a 150 percent deduction, this volume-based R&D tax allowance was raised to 175 percent in 2018 and, as of 2023, stood at 200 percent for most companies (Dezan Shira and Associates 2024). Additional incentives included accelerated depreciation for R&D assets, preferential corporate income tax rates for "High-and-New-Technology Enterprises," and VAT and customs relief on imported R&D machinery. Collectively, these incentives make China's R&D tax subsidy among the world's largest, sharply enhancing the after-tax return on innovation investments (ibid).

This capital was directed toward "national champions" like Huawei, CRRC, and Sinopec, which received \$248 billion annually in industrial subsidies. Beijing sought to dominate manufacturing in green technology and cultivate self-sufficiency in semiconductors and keep pace with the US in A.I., quantum computing, and other cutting-edge technologies (Lee, 2018; Naughton, 2021). It did so by doling out generous grants, preferential loans—between 2020 and 2024 state-

controlled banks lent an additional \$1.9 trillion to industrial borrowers—land-use rights, and undertaking state-mandated mergers (Casanova and Miroux 2018).

This strategy extended to building a massive, state-controlled domestic infrastructure, such as high-speed rail and renewable energy capacity. Beijing also turned to using huge government procurement contracts to endow Chinese firms with inimitable advantages, such as reaching economies of scale. Further, China used the Belt and Road initiative to secure access to energy-related natural resources and weaponized supply chains, using its control of critical minerals and drone batteries to pressure rivals. Finally, Beijing explored alternatives to the U.S. controlled SWIFT payments system by exploring digital-currency platforms (e.g., mBridge).

### **The "Self-Reliance" Endgame: Gauging China's Success**

These policies reflected a dual motivation—to hedge against U.S. export controls and to ensure durable leadership in technologies of the future—and have yielded concrete gains in narrowing the Sino-American innovation gap.

China's bid for self-sufficiency and technological leadership accelerated rapidly after it emerged economically battered from the COVID-19 pandemic.<sup>168</sup> This domestic-first industrial strategy, combined with weak consumer spending, meant China began replacing foreign suppliers with local ones across sectors. State-owned firms faced directives to swap out foreign software under the "Delete A" ("Delete America") initiative, while private companies were pressured to source domestically, effectively squeezing out imports.

This strategy produced tangible results. China's share of global patent filings in AI and 5G standards rose sharply, and its electric-vehicle exports surged by nearly 60 percent in 2024 alone (Spegele 2025). By 2023, Chinese robot makers supplied nearly half of the country's industrial-robot installations, up from near-zero less than a decade earlier, allowing factories to automate *en masse*. In semiconductors, China's self-sufficiency in GPUs was projected to soar from 11 percent in 2021 to 82 percent by 2027. By 20205, state-led projects had built 53 percent of global ship-building capacity, up from 8 percent in 2002.

During this time, China had also increased its share of global manufacturing from 6% in 2000 to 32%, with its factory output exceeding the combined manufacturing of the United States, Germany, Japan, South Korea, and Britain.<sup>169</sup>

In turn, Chinese exports rapidly accelerated, rising 13.3% in 2023 and 17.3% in 2024. Electric carmaker BYD, for example, was slated to bring new factories online capable of producing twice as many cars as Volkswagen's Wolfsburg facility. China also built more petrochemical refinery capacity between 2019 and 2025 than Europe, Japan, and South Korea created since World War II.<sup>170</sup>

---

<sup>168</sup> This paragraph draws closely on Douglas and Leong. 2025.

<sup>169</sup> All these facts are from Bradsher (2025).

<sup>170</sup> By 2025, the results of China's subsidy-heavy approach were mixed in terms of efficiency and innovation outcomes. Looking at Chinese industrial subsidies (across sectors including tech)

This dominance provided Beijing with significant geopolitical leverage. China strategically restricted exports of critical minerals and rare earths—especially exports of terbium, scandium, and yttrium, key raw materials for EVs and defense systems—to gain an edge over trade negotiations (Kenderdine and Meidan 2021).<sup>171</sup>

In the June 2025 London framework to restore the U.S.–China trade truce, China used its export-licensing regime as the central bargaining chip, agreeing to only six-month licenses for key rare-earth elements and magnets. This cap ensured Beijing reserved itself the right to renew restrictions if broader negotiations were to stall, turning its refining dominance into a durable tool of geopolitical influence.

### **China’s IP Protections**

Despite its aggressive, state-directed industrial policy, the Chinese government also took a page out of the United States’ Creative Destruction Paradigm during the early 21<sup>st</sup> Century. By 2025, China had steadily improved its IPR for over two decades. It joined all major international IP conventions.<sup>172</sup> It also waged an extensive anti-counterfeiting and anti-piracy campaign and created additional enforcement capacity in the form of IP affairs departments (Yang 2009).

After a blockbuster 2007 WTO ruling, China took steps to liberalize the individual ownership of state-funded patents. In 2020, a new foreign investment law and implementing regulations went into effect, which made stronger commitments to protecting foreigners’ IPRs, including trade secrets and patents, and displaced the Chinese-Foreign Equity Joint Venture Law, the Chinese-Foreign Cooperative Joint Venture Law and the Wholly Foreign Owned Enterprise Law, a trio of laws imposing joint ownership requirements on western firms that justified forced technology transfers.

On the IP enforcement front, China made noteworthy strides. As Nguyen (2010) points out, especially after 2001, IP owners were able to successfully use the judicial system to enforce their rights. China boasted specialized IP courts that moved with alacrity and relatively low litigation costs, at least compared to the U.S. (see Morinville 2018). It also bolstered IP enforcement by eliminating pockets of judicial antipathy towards foreign IP and created oversight bodies and regional IP courts (see Weightman 2018). Indeed, China’s patent enforcement process was

---

from 2007–2018, Branstetter et al. (2023) found that while subsidies skyrocketed—a seven-fold increase from \$4 billion to \$29 billion in annual subsidies to listed companies—these handouts had little to no effect on boosting firms’ productivity. In fact, the data suggested subsidies often flowed to less efficient firms and failed to spur significant innovation or efficiency gains.

<sup>171</sup> In 2025, China was the world’s topmost producer of minerals such as graphite, gallium, and germanium that are critical to semiconductors, EVs, and defense systems.

<sup>172</sup> These are the World Intellectual Property Organization (WIPO), the Berne Convention for Protection of Literary and Artistic Works (copyright), the Universal Copyright Convention, the Paris Convention for the Protection of Industrial Property (patent and trademark), the Patent Cooperation Treaty, the Agreement on Trade-Related Aspects of Intellectual Property Rights, and the Madrid Agreement for the International Registration of Trademarks (Greguras 2007).

tailored to benefit inventors through negotiated settlements.<sup>173</sup> The result was that between 2006 and 2011, foreign companies brought 10% of patent infringement cases in China and won over 70% of them (Love, Helmers, and Eberhardt 2016).<sup>174</sup>

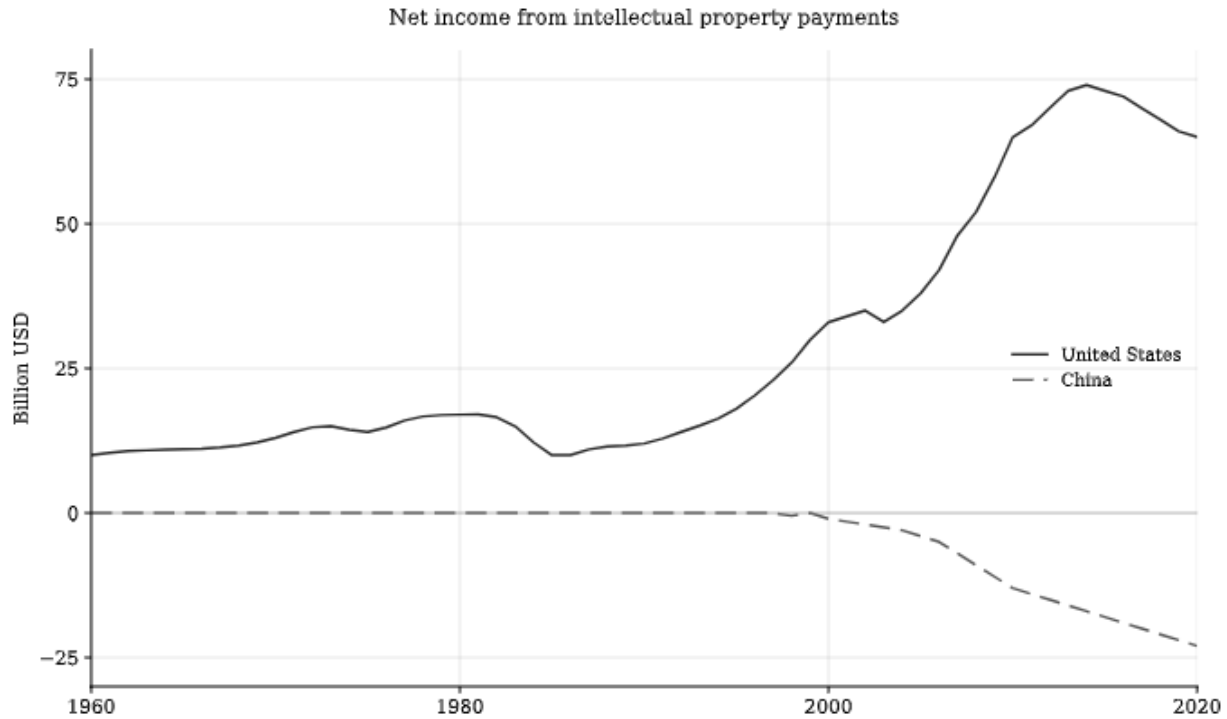
Finally, and contrary to conventional wisdom about China's disrespect for IP, leading up to the 2020s Chinese companies acquired foreign technology from the U.S. and other industrialized countries through copious patent licensing. Chinese companies operating in sectors such as transportation, energy, and robotics paid top dollar to foreign patent holders to gain access to technology from the industrial frontier: Japanese and American firms received billions of dollars in royalties in exchange for these licenses (Taplin 2018). In 2019, alone, China paid over \$34 billion to the rest of the world for the legal use of IP. The U.S. accounted for roughly 23% of this amount (World Bank 2020; OECD 2020). Figure S4.2 shows that China's royalty payments to the U.S. grew dramatically faster than its GDP over the last two decades, echoing the substantial improvements in IP protection described above (see also Lardy 2018).

**Figure S4.2: Payments for and income generated from intellectual property.**

---

<sup>173</sup> There were three main reasons for this. First, patent owners were legally obligated to notify infringers before launching a lawsuit. Second, patentholders were incentivized to sign a nondisclosure agreement, under which they were able to divulge technical information and discuss disputed issues. Third, because the likelihood a Chinese court would impose an injunction on an infringer was relatively high—for example, in 2018, injunction rates averaged around 98% (Weightman 2018)—this increased the odds that parties sued for infringement would try to seek a negotiated solution that culminated in royalty payments for patentholders.

<sup>174</sup> Moreover, in the vast majority of lawsuits MNCs brought against Huawei for stealing trade secrets, for example, the parties have reached out-of-court monetary settlements or the MNCs were awarded monetary damages (Taplin 2018). To be sure, these are not the same as an injunction issued against Huawei from selling products that use infringed upon IP. But it's not nothing either.



Notes: To create this variable, I subtracted income generated through the sale and licensing of IP from payments made to acquire it. Values are expressed in constant 2010 USD (normalized for inflation using the Consumer Price Index). Rather than presenting raw annual data points, which would appear as irregular year-to-year fluctuations, I apply a smoothing technique that preserves the essential long-term trends while enhancing visual clarity and interpretability: I apply cubic spline interpolation to the US series and quadratic interpolation to China's to create smooth curves between observed values. This approach offers advantages like calculating moving averages in that it filters out short-term noise to reveal underlying patterns—such as the structural shift in the US trajectory around 2000 and the sustained decline in China's position post-2000—though unlike moving averages, interpolation maintains the timing of critical inflection points like the mid-1980s dip in US receipts.

Sources: IMF Balance of Payments Statistics Yearbook and data files.

### China's Indigenous AI Stack Push

The Chinese state (at national and local levels) heavily subsidizes AI companies and research. Sometimes this is done via public-private contests and funds—for example, the Ministry of Science and Technology and industry groups sponsor AI competitions where winners receive large R&D grants (Luong and Konaev, 2023). Beyond contests, local governments offer cash rewards, rent subsidies, and tax breaks to AI firms setting up in their jurisdictions. Chinese government spending on AI-related development reached tens of billions of dollars annually by the late 2010s (Zhang, 2024; Luong and Konaev, 2023). Key enterprises in AI (from giants like Baidu, Alibaba, Tencent to startups) often receive state grants to pursue specific projects (e.g. autonomous vehicles or medical AI applications), blurring the line between public and private sector efforts. This has meant accelerating growth in targeted subfields like facial recognition,

where companies like SenseTime and Megvii thrived with government contracts and funding (Lee, 2018).

The Chinese government also subsidizes the infrastructure that AI development needs. This includes building AI research parks and innovation zones in cities like Beijing, Shanghai, Shenzhen, and Hefei (Zhang, 2024). Such parks often come with incentives for companies and shared facilities (like supercomputing centers or cloud platforms) that firms can use at low cost. Additionally, China's relatively lax data regulation (until recently) and huge population have provided AI developers with abundant data—effectively a resource the state allowed them to exploit for AI training (Wei, 2021). (China's new Personal Information Protection Law may tighten data use, but historically, Chinese AI firms had more leeway to gather data than their European counterparts constrained by GDPR.)

The Chinese state has been an enthusiastic early adopter of AI, particularly for surveillance, public security, and administrative purposes. By deploying AI solutions (like computer vision systems for city surveillance or AI-powered Covid tracking apps), the government has given domestic AI companies large contracts and valuable real-world experience (Luong and Konaev, 2023; Wei, 2021). This domestic market nurtured firms that are now globally competitive in certain niches (e.g., Chinese companies supply over half the world's facial recognition tech). The close relationship between industry and a government willing to spend on AI creates a virtuous cycle of revenue and iterative improvement for Chinese AI developers.

### *Infrastructure*

Chinese data centers consumed 140 billion kWh in 2024, and this figure was projected to quadruple by 2035, reflecting the country's ability to support energy-intensive AI workloads (Du et al. 2025). To cover the country in 5G, Beijing spent hundreds of billions of dollars on base stations, new cell towers, and other infrastructure; it allocated significant chunks of radio spectrum that mixes fast speeds with moderate transmission distances to three state owned telecommunication companies. It directed national regulators and provincial and local governments to coordinate the nationwide rollout of 5G, using its muscle over land rights (Woo 2019). It also awarded Huawei lucrative contracts to provide equipment to the network. And for decades, Chinese government procurement centered on computers, telecommunication infrastructure, office equipment, software, renewable energy, and energy efficiency.

### *Semiconductors*

During the early 2020s, China pursued semiconductor self-sufficiency, with ambitious plans to construct dozens of major chip factories (Fuller et al., 2020; Luong and Konaev, 2023). National champion Huawei played a crucial role in fostering collaboration within China's chipmaking ecosystem through co-investment and staff exchanges with chip foundries, while making progress in EDA software development with Empyrean (Zhang, 2024).

To help bankroll the refashioning of China's semiconductor supply chain, Beijing substantially increased subsidies and launched the "Information Innovation" project (xinchuang). Moreover,

the Chinese government maintained significant stakes in key players like SMIC, AMEC, and Empyrean (Branstetter et al., 2023; Zhang, 2024).

By 2025, these efforts met with considerable success. Huawei demonstrated progress by producing a smartphone with a seven-nanometer chip and was developing 5nm chips with SMIC, while its Ascend chips were being used by companies like Baidu for AI applications (Luong and Konaev, 2023). Despite these advances, China still requires substantial support to achieve complete independence in chip production, having already invested approximately \$150 billion in subsidies over the past decade (Branstetter et al., 2023; Fuller et al., 2020).

China achieved rapid improvements in its domestic semiconductor design and fabrication capabilities as it went about building out an entire semiconductor supply chain, beginning after US export controls on high-end chips (Hilson 2024; Mims 2025c). As US policymakers redoubled their efforts, banning EUV-class lithography tools, advanced CPU/GPU designs, and certain chemicals involved in chipmaking, Chinese firms redoubled their efforts.<sup>175</sup> During the early 2020s, Chinese firms like Cambricon and Huawei introduced competitive alternatives that challenged Nvidia's hardware because they had no choice. This included Cambricon's A100-like processor and Huawei's CloudMatrix cluster of Ascend AI chips.

Despite years on a US trade blacklist and Washington's sanctions on both advanced manufacturing equipment and memory components, Huawei shipped around 800,000 units of its Ascend 910B and 910C models in 2025 to major state carriers—in 2025, Chinese state data centers reported that most of their new deployments used Chinese-made chips—and private developers such as ByteDance; these customers ramped up orders of the 910C after Nvidia's H20 chips were added to America's export-restricted list by the Trump administration in 2025 after the success of efficient Chinese models like DeepSeek that demonstrated competitive AI performance with less compute by dint of exploiting these inference accelerator chips.<sup>176</sup>

---

<sup>175</sup> Chinese companies also employed various strategies to circumvent US export restrictions on high-end AI chips. Huawei relied on indirect procurement from overseas foundries via intermediaries (TrendForce 2025). Chinese cloud and datacenter operators leased foreign-made Nvidia servers through third countries. There was also large-scale smuggling of restricted AI chips into China, with estimates suggesting that 10–50% of China's AI training capacity in 2024 relied on such illicitly obtained chips (Grunewald 2024). Also, Chinese AI developers physically carried terabytes of training data on hard drives to overseas data centers. There, they rented servers equipped with restricted chips, ran their training jobs, and then flew the resulting model parameters back to China. This illicit activity, in turn, spurred nearly 2,000 MW of new data-center capacity across Singapore, Malaysia, Thailand, and Indonesia (see Huang and Lin 2025). US Enforcement of these controls was further hampered by a severe staffing shortage at the Commerce Department's Bureau of Industry and Security and by the inherent difficulty of tracing chips once they leave Nvidia's official channels (see The Economist 2025a).

<sup>176</sup> In the fiscal year ending January 2025, Nvidia sold about \$12 billion worth of H20 inference chips to China: Anticipating restrictions, major Chinese firms—including Alibaba, Tencent, and ByteDance—placed massive orders for H20-based servers and modules (Sanghvi 2025). However, while the Trump administration at first banned shipments of Nvidia's H20 and AMD's

Perhaps most importantly, along with other Chinese technology firms, Huawei laid the groundwork for an indigenous chip supply chain (Mims 2025c; Sbeglia Nin 2025; Lee and Li 2025). Huawei created the Ascend 910D AI processor with the explicit goal of matching or surpassing Nvidia's H100 chip.<sup>177</sup> And in late 2025, Huawei disclosed a three-year roadmap for its Ascend AI chips as a direct challenge to Nvidia's market dominance (Liu 2025). The plan, which includes four new generations of AI accelerators through 2028, contained a critical announcement: the 2026 Ascend 950 series would feature Huawei's own high-bandwidth memory (HBM). This move was a major step toward complete self-reliance, as it would eliminate a key bottleneck and end the company's dependence on U.S. and South Korean HBM suppliers (ibid.). Meanwhile, while CXMT made significant advances in high-bandwidth memory, equipment makers AMEC and Naura unveiled etching and layer-deposition tools once exclusively produced by Western giants (The Economist 2025b). Concurrently, firms such as Cambricon and DeepSeek successfully designed their own AI accelerators (Mims 2025c).

To compensate for its lack of cutting-edge chips, China pursued a "swarms beat the titan" strategy (Chin and Huang 2025). This involves building AI supercomputers that bundle together thousands of less powerful, domestically produced accelerators (like Huawei's Ascend chips) to match the performance of a smaller number of Nvidia's "titan" GPUs. And Beijing strongarmed local governments to shower data centers with deep electricity subsidies if they agreed to use these domestic chips to address the fact that this "swarm" approach consumes significantly more power (ibid).

Most strikingly perhaps, by 2025, SMIC, China's largest foundry, had become the world's third-largest chip manufacturer by revenue, behind only TSMC and Samsung (TrendForce 2025; Zhu 2025). Despite relying on older process nodes, SMIC managed to produce 7nm-class chips for Huawei and was actively working toward 5nm production.<sup>178</sup> Chinese foundries adopted local equipment, with companies like AMEC and Naura gaining ground against foreign competitors (Luong and Konaev, 2023).<sup>179</sup>

---

MI308 chips to China, it partially reversed this decision after industry and diplomatic pushback, allowing exports to resume under certain license conditions (Sherman 2025).

<sup>177</sup> Huawei's chip-making arm leaned on "chip-stacking" techniques and system-level designs to close the performance gap created by its exclusion from TSMC's leading-edge foundries and SMIC's inability to acquire the latest lithography tools. Its Ascend 910C AI chips delivered approximately 60% of the inference performance of Nvidia's H100 in practical applications. While the Ascend 910C was more power-efficient at 310W compared to the H100's 700W, Huawei compensated for lower per-chip performance by deploying more chips in parallel. This strategy allowed its CloudMatrix 384 system—linking 384 Ascend 910C chips—to exceed the throughput of Nvidia's 72-GPU Blackwell rack under certain workloads, achieving competitive raw computing power and memory capacity, albeit with higher total energy consumption (Qureshi 2025; The Economist 2025b).

<sup>178</sup> On SMIC's capabilities, challenges, and potential see TrendForce 2025; Mims 2025c; Hilson 2024; Averroes AI 2025; Zhu 2025.

<sup>179</sup> Yet, circa 2025, critical dependencies remained. China still lacked a domestic counterpart to ASML's extreme-ultraviolet lithography machines and in addition relied on related semiconductor equipment imports from Japan (Feldgoise et al., 2024). Many "indigenous" chips

Notwithstanding these substantial gains in mature chip production, however, by 2025 China's state-led model continued to struggle to close the gap on the cutting-edge frontier. Even as SMIC clocked impressive improvements, TSMC remained the unmatched industry leader in semiconductor fabrication, widening its lead in both capital expenditure and technological sophistication (Loo 2025). And while there is no denying that by underwriting a self-reliant supply chain for older chips China's massive subsidies were helping it close the gap, the state-directed firms within this nascent semiconductor ecosystem fell significantly short of the scale and R&D velocity required to match the technological frontier at the 3nm-and-below nodes powering advanced AI during this period (ibid.).

### *AI Models*

In terms of AI models, Beijing recognized the strategic value of an open-source playbook and had backed its own domestic variants—such as the MOSS and InternLM models—through partnerships between government labs and local cloud providers.

Along these lines, Beijing designated DeepSeek as a "national treasure" in China in 2025. Since its breakthrough, DeepSeek's CEO, Liang Wenfeng, has been invited to high-level gatherings with Chinese leaders, including President Xi Jinping and Premier Li Qiang, a rare occurrence for a relatively young, privately-funded startup. It implemented travel restrictions on some employees involved in AI model research and development, requiring them to hand in their passports due to concerns about confidential information, potentially including trade and state secrets, being leaked. Furthermore, the local government of Zhejiang Province, where DeepSeek's parent company is headquartered, has begun screening potential investors before allowing them to meet with company leaders. This involves registering investment inquiries with the provincial Communist Party committee. Additionally, headhunters have reportedly been contacted by government officials and asked not to poach talent from DeepSeek.<sup>180</sup>

Beijing also sought to aggressively diffuse open source models across China. One way to do this was to mandate data-localization policies that ensured wide internal uptake. DeepSeek's open-sourced models were being rapidly integrated into the IT infrastructure and workflows of local governments and state-owned companies across China.

### *Robots*

Finally, China also mobilized massive state support—some \$23 billion in local government funds plus a \$138 billion central venture pool—to accelerate the mass production of robots and build a domestically sourced supply chain for critical hardware like motors, sensors, and custom drives (Casado and Neuberger 2025). This state-led push yielded dramatic results. By 2025,

---

incorporated Western components. On the software side, Nvidia's CUDA ecosystem retained a near-monopoly on AI programming; Huawei's CANN framework for its Ascend chips were far less mature and suffered from bugs, slowing developer adoption (see The Economist 2025b).

<sup>180</sup> On all these points see Osawa and Liu (2025).

while the U.S. celebrated its lead in AI software, China was quietly winning the race for "embodied AI." China now has more robots in production than the rest of the world combined, operating "dark factories" staffed entirely by robots, and has surpassed the U.S., Germany, and Japan in robot density (ibid.).

This boom was not just in quantity but in quality. Chinese firms like Unitree produced advanced humanoid robots for as little as \$5,900, demonstrating they were not cheap knockoffs but sophisticated machines (ibid.; Robotics Tomorrow 2025). This created a profound strategic risk for the U.S., which lacks a comparable robotics ecosystem. As China builds its domestic component supply chain—or acquires foreign leaders like the German giant KUKA—American firms may become reliant on Chinese suppliers for the physical hardware of the AI revolution (Casado and Neuberger 2025).

What's more: China's advantage stems not only from state subsidies and its vast manufacturing base (which provides diverse real-world scenarios for training), but from a rapid innovation cycle and unique market structure and economies of scale. This cycle boasts not only extraordinarily velocity, but nimble feedback mechanisms that drive quality improvement. By 2025, leading companies were iterating new robot generations every few months, rapidly integrating cost-saving design changes and manufacturing advances (Asian Robotics Review 2024). By controlling every stage from chipmaking to final assembly, firms compressed supply chains, immediately translating R&D advances into large-scale production (Han et al. 2025). Meanwhile, China's vast, diverse market enabled real-world field deployment of new robotic systems at enormous scale, providing rapid performance data and user feedback—allowing for a virtuous cycle of rapid refinement (Asian Robotics Review 2025).

### **China's Strategy Over Standards**

In the early 2020s, Beijing moved to shape AI standards through a multipronged strategy. This included integrating the pursuit of standards dominance through the aggressive industrial policy outlined above; centralizing national standard-setting practices; and imposing binding platform rules. Also, like the US, it sought to capitalize on the growing commercial importance of leading AI models that strove to become market leaders. By 2025, Greater Beijing and Shanghai ranked among the world's top hyperscale markets, facilitating the Chinese government's efforts to operationalize its preferred formats, interfaces, and testing practices domestically, if not internationally (Synergy Research Group 2025).

In 2025, Chinese tech firms began to win a key part of the AI race: the battle for open-source dominance. While American companies were widely seen as offering the most powerful proprietary models, Chinese firms decisively outcompeted their U.S. rivals in "open-weight" AI—technology that anyone can freely download and build upon. This "open-weight" approach, which makes the model's learned parameters available for modification, was championed by Chinese firms with high-performing models like DeepSeek R1 and Alibaba's Qwen3 (Huang 2025).

This marked a significant reversal from the previous year, when U.S.-based Meta offered the world's best open model. By October 2025, the top-ranked open models on LMArena, a site that ranks models based on blind user preference tests, were all from Chinese companies like

DeepSeek, Z.ai, and Alibaba (Fu 2025). This popularity was confirmed on Hugging Face, the leading platform for AI developers: In early November of 2025, DeepSeek-R1's model card showed about 12.8k likes—more than twice Meta's most-liked recent Llama-3.x instruct release, which garnered about 4.9k likes (Hugging Face 2025a; Hugging Face 2025b).

Because these Chinese models excel at local-language tasks and can be run in-house, they rapidly became a default choice for enterprises prioritizing data control. This adoption included multinational banks like HSBC and Standard Chartered, industrial giants like Saudi Aramco, and even major U.S. cloud vendors like AWS, Microsoft, and Google (Lin et al. 2025).

This trend, which directly challenged the Trump administration's AI strategy, left the U.S. scrambling to respond with new initiatives like the American Truly Open Models (ATOM) Project (ATOM Project Substack 2025). In a sign of a potential strategic retreat, Meta's Mark Zuckerberg signaled in a new essay that the company would need to be "careful about what we choose to open source," suggesting Meta might keep its next frontier model closed (Nolan 2025). And OpenAI unveiled its first open-weight model family, gpt-oss, releasing the model weights under a permissive Apache-2.0 license in a strategic shift widely interpreted as an attempt to counter the momentum of China's flexible and accessible AI ecosystem (Willison 2025).

This battle for open-source dominance was being fought most intensely in the lucrative market for AI coding assistants. In 2025 Chinese startups like MiniMax, Zhipu, and Moonshot AI were releasing a wave of new models that are not only high performing but radically cheaper than U.S. alternatives (Osawa 2025). Consider MiniMax's M2 model, which ranked highly on developer leaderboards, and was priced at just 8% of the cost of Anthropic's competing Claude Sonnet model (ibid.). This aggressive strategy was driven by the need to find paying customers outside of China, where fierce domestic price wars make it difficult to generate revenue. The approach was proving successful: in late 2025, U.S. cloud company Vercel had partnered with Zhipu and celebrated venture capitalist Chamath Palihapitiya had migrated workloads to Moonshot AI's Kimi K2 (ibid.).

To complement this global commercial success, Chinese ministries committed to developing over 50 new national and industry AI standards by 2026. This top-down effort includes creating committees to harmonize technical specifications for LLMs and conduct risk assessments across different vendors. In parallel, in 2025 industry alliances were working to standardize the technical connections—the APIs and data formats—between AI models and domestically produced accelerator chips to ensure interoperability within China's tech ecosystem (Reuters 2024a). Finally, Beijing used interconnection standards to create a unified national grid for computing power to allow surplus capacity in one Chinese region to be reallocated elsewhere (MFA PRC 2025).

Where the United States leaned heavily on NIST's voluntary frameworks, China anchored standardization in binding, pre-market obligations administered by the Cyberspace Administration of China (CAC) and a dense lattice of national/industry standards (CESI 2020). China's standards propagate through market entry and platform licensing: to serve the public in

China, model and service providers must file algorithms, meet content-identification and traceability rules, pass security evaluations, and align with national implementation guides.

Three cornerstone rule sets created de-facto technical baselines. First, the Algorithmic Recommendation Provisions, which require large platforms to file algorithms in a national registry, disclose basic operating principles, curb manipulative ranking, and maintain a “features library” for harmful content screening. Second, the Deep Synthesis (deepfake) Provisions, which mandate content labeling/watermarking, identity verification, and periodic reviews of synthesis models. Third, the Interim Measures for Generative AI Services, which require security assessments, data-quality controls, user-rights handling, and clear API terms.

To ensure compliant AI systems could be deployed at a vast scale, Beijing expanded its national compute and cloud infrastructure. Its most notable initiative on this front, the “Eastern Data, Western Computing” program, established a network of eight national computing hubs and ten data-center clusters (NDRC 2022). This project created a powerful, geographically distributed foundation for processing and storing the immense amounts of data required for a national AI ecosystem.

Finally, China’s active presence in helping shape the international AI standards set forth by ISO/IEC SC 42 gave those domestic baselines pathways into international committees where they can be negotiated, adapted, or referenced by other jurisdictions (see ISO/IEC 2025). Plus, Beijing’s July 2025 Global AI Governance Action Plan calls for accelerating technical standards in security, ethics, and industrial application, dialogue among national standards bodies, unified computing-power standards, and even mutual-recognition platforms for safety assessment—explicit pathways to socialize China’s testing and documentation templates abroad (MFA PRC, 2025). Indeed, China went so far as to formalizing mirror structures to international AI standards bodies and tasked them with domain-specific working groups (e.g., AI for steel, energy) to translate policy goals into technical specs.

## EUROPE’S INDUSTRIAL POLICY

Before 2025, Europe’s proto-industrial policy was twofold, focusing on strengthening its traditional manufacturing base while simultaneously promoting digital interoperability (Mattoo and Staiger 2020; Veugelers 2018). This strategy translated into targeted R&D investments, with initiatives like the Chips Joint Undertaking aimed at bolstering the semiconductor supply chain (European Commission 2021), and programs like “Industry 4.0” designed to foster data-sharing among established firms (European Parliament 2020). In parallel with these internal efforts, the EU’s posture toward China hardened; labeling Beijing a “systemic rival” (European Commission 2019), it began to follow the U.S. playbook: implementing increased regulations on foreign investments and imposing new tariffs on critical Chinese imports, such as electric vehicles (Crawford 2021).

During the early 2020s, European tech firms experienced some successes. These included significant investments in European AI startups like Aleph Alpha (half a billion dollars in 2023)

and Mistral AI (over \$600 million since early 2023). The European Union also worked to foster its own AI industry with initiatives like the GAIA-X cloud project and investments in R&D.<sup>181</sup>

By and large, however, Europe's industrial and innovation strategy struggled to keep pace with the U.S. and China, as reflected in several stark metrics and structural challenges. Over the last half-century up to 2025, the U.S. produced 241 new \$10-billion-plus public companies from scratch, versus only 14 in Europe, while EU labor productivity per hour slipped from roughly 95% of U.S. levels in the late 1990s to under 80% by 2025 (Fairless and Luhnnow 2025). Europe also suffered from missed commercialization opportunities: seminal European inventions such as the World Wide Web (born at CERN) and MP3 (developed at Fraunhofer) were ultimately scaled and monetized disproportionately by firms outside Europe (CERN n.d.; Brandenburg n.d.). Consequently, EU members demonstrated a relatively small tech footprint. As of early 2025, there were about 100 European "unicorns" (about one-third of a trillion dollars in combined value), far behind the U.S. and China in both counts and valuations (CB Insights 2025d; Fairless and Luhnnow 2025). By 2025, only four of the world's top 50 publicly traded tech companies were European, despite the EU representing a sizable share of global output (about 14% of world GDP on a PPP basis) (Draghi 2024).

By 2025, Europe's semiconductor industry revealed a particularly troubling decline from past strength. While the region (broadly construed) still hosted critical champions like ASML and ARM, Europe's share of global manufacturing capacity had eroded (Saint-Martin 2024; the EU itself set a goal to double its global chip share to 20% by 2030—an implicit acknowledgment that it hovered near about 10%—and Europe had no prominent pure-play foundry competing with TSMC or Samsung at advanced nodes (European Court of Auditors 2025). Investment in domestic chipmaking also lagged the scale of rival programs, with the U.S. CHIPS and Science Act appropriating \$52.7 billion, while estimates of China's support exceed \$150 billion over 2014–2030 (Blevins et al. 2023).

The AI and advanced-tech gaps further illustrated Europe's competitive disadvantage. In 2025, none of the top ten global investors in quantum computing were European, and only a handful of firms on the continent pursued frontier-scale AI models (Fairless and Luhnnow 2025). Germany's Aleph Alpha, for example, pivoted away from the frontier-model race toward enterprise solutions and consulting-style deployments, while many promising teams sought U.S. capital markets, looser regulation, and stock-option incentives more conducive to rapid scaling (Lomas 2024; Fairless and Luhnnow 2025).

### **Contextualizing Europe's Technological Backwardness**

Capital markets presented a significant weakness for European innovation. In 2025, European venture-tech investment remained roughly one-fifth the U.S. level, as much financing still originated from banks or large pension funds, leaving a "missing middle" of dynamic, growth-oriented capital. weak venture capital environment and a shallow entrepreneurial ecosystem and culture. Historically, European capital markets—especially in Germany—provided less venture capital compared to the U.S. and even China, partly due to more conservative banking structures,

---

<sup>181</sup> On all these points see Black 2024.

fewer tax incentives for entrepreneurship, and differences in equity culture (Veugelers 2018). This sluggish venture capital environment hampered the growth of promising European startups and limited their ability to scale globally. While the U.S. fostered high-growth startups across the Third and Fourth Industrial Revolutions, including AI powerhouses such as Anthropic and OpenAI, European innovators struggled with early-stage funding (Schneider 2024).

Regulatory and cultural barriers compounded Europe’s challenges around innovation. Europe’s complex multi-jurisdictional market and stricter labor frameworks made cross-border scaling and hiring/firing harder, while historically less-friendly equity rules dampened stock-option use (Baroudy et al. 2020). In several countries, stock options were even taxed at grant rather than at exercise or sale—further reducing their attractiveness (Palmer 2019). Companies reported spending about 40% of IT budgets on compliance with EU regulations, and roughly two-thirds of firms said they still struggled to understand their obligations under the EU AI Act (Duffer et al. 2025).

And while U.S. government agencies like DARPA, the NSF, and NIH long underwrote high-risk, high-reward research—leading to breakthroughs in AI—the EU did much less of this (Mazzucato 2013).

### **Europe’s Response to its Precarious Position**

In response to this precarious position, the EU and national governments began to pivot, making significant public investments and providing systematic incentives characteristic of industrial policy.<sup>182</sup> For example, in 2024 the German government promised to provide €9.9 billion in subsidies for Intel’s semiconductor factory in Magdeburg, which represents a significant portion of Intel’s €33 billion investment. The factory was expected to create 3,000 direct jobs at Intel and potentially up to 20,100 jobs in total, including indirect employment. The Chancellery and Economy Ministry emphasized that the factory could serve as an innovation hub, contributing to the development of the semiconductor industry and fostering economic growth in the region.<sup>183</sup> Moreover, the project aligned with the EU’s strategic goal of increasing its share of global chip production from 10% to 20% by 2030, thereby reducing dependence on Asian suppliers.

This turn toward state subsidies for hardware was complemented by a growing strategic desire for “digital sovereignty,” driven by anxieties over Europe’s deep reliance on American cloud providers (Johnson and Hay 2025). By 2025, there were clear signs that European companies and governments were souring on their dependence on U.S. “hyperscalers” like Amazon Web Services, Microsoft Azure, and Google Cloud; Denmark’s two largest municipalities, for example, moved to end use of Microsoft systems amid sovereignty concerns (Desmarais 2025). These anxieties were amplified by the return of the Trump administration, renewed uncertainty over transatlantic data-sharing arrangements, and fears tied to U.S. surveillance authorities such as the CLOUD Act and FISA §702 (Johnson and Hay 2025). Switzerland, meanwhile, pressed

---

<sup>182</sup> This paragraph closely draws on Gerresheim and Krahe. 2024.

<sup>183</sup> Circa 2025, these commitments were complicated by Intel’s market challenges, as it had fallen significantly behind chip fabricators such as TSMC and Samsung and had recorded heavy losses throughout 2024 and suffered sizable drops in its market capitalization after stock selloffs.

for stronger data-residency options, with major investments framed to keep more workloads within Swiss borders (Revill 2025).

This push for sovereignty was intellectually framed by “EuroStack,” a programmatic blueprint commissioned by European institutions (Bria et al. 2025). The EuroStack report outlines a strategy for sovereign digital infrastructure anchored in interoperability, sustainability, and treating “data as a common good,” and calls for tools like a “Buy European” act, a close to €300 billion sovereign tech investment program, and a “sovereign AI cloud” marrying public supercomputers with decentralized systems (ibid).

By 2025, the EuroStack agenda was mirrored in multi-billion-euro EU programs aimed at achieving “sovereign compute”—a domestic, EU-controlled infrastructure free from reliance on foreign platforms. This strategy was spearheaded by the EuroHPC (European High-Performance Computing Joint Undertaking), an agency that began upgrading its traditional scientific supercomputers with thousands of GPUs to create “AI factories.”<sup>184</sup> This was complemented by the EU Commission’s “InvestAI” plan, a €200 billion initiative designed to fund these “AI gigafactories” and provide European startups with the domestic capacity needed to build large-scale models (Kyosovska and Renda 2025).

### **Europe’s AI Standards Strategy**

Unlike the reliance on market scale driving de facto standards or a reliance on voluntary standard setting, the EU dictates AI technical standards. The European Commission, the EU’s executive body, issues standardization requests to European standards bodies—primarily CEN-CENELEC—to develop detailed, harmonized European Standards.<sup>185</sup> These standards act as a common rulebook for the entire market, specifying reference architectures, which are like standardized blueprints for designing trustworthy AI systems; data-quality criteria to ensure the information used to train AI is accurate and free from bias; common test and assurance methods to verify an AI’s performance and safety; and interface profiles, which act like universal connection points to ensure different AI tools can work together seamlessly. This gives suppliers a predictable “single-build” target for the European market, encouraging uniform implementations, modular component swapping, and cross-border interoperability (European Commission 2025; CEN-CENELEC 2025; European Union 2024).

To understand what this means on the ground, consider a vendor selling a computer-vision module, a type of AI that visually inspects products for defects on an assembly line, to several automotive plants within the EU. Instead of creating different versions for each country, the vendor follows a single technical rulebook known as a harmonized European Standard. This standard, developed by CEN-CENELEC, details all the technical requirements, which include

---

<sup>184</sup> The strategic goal of this push was to break Europe’s critical dependence on American cloud providers like Amazon Web Services, Google Cloud, and Microsoft Azure for the high-performance computing required for AI model training.

<sup>185</sup> CEN-CENELEC is the joint standardization platform that brings together CEN—the European Committee for Standardization—and CENELEC, the European Committee for Electrotechnical Standardization.

common data interfaces so the AI can connect to different factory systems; specific validation and robustness tests to prove its accuracy and reliability; and clear rules for lifecycle documentation, which is a record of how the AI was built, trained, and tested. The standard also specifies the need for logging hooks, which are connection points for monitoring the AI's decisions, and change-control procedures that dictate a safe process for updating the software.

To prove it has followed these rules, the company compiles a technical file, an "evidence binder" containing test reports, dataset descriptions, and monitoring plans that map directly to each requirement in the standard. This enables it to roll out synchronized updates across all its European sites, preserving the system's interoperability, which is the ability to work with other factory equipment, and its modularity, the ability to be updated in parts without breaking the entire system.

Returning to EU standard setting in general, AI systems are routed through a structured "conformity-assessment" regime (European Union 2024). For less risky AI, this can be a self-assessment, but for higher-risk systems, it must be reviewed by a Notified Body, an independent, EU-approved auditor. Once this process is complete, the final integrator grants the CE marking, a product "passport" allowing the exact same software package (the binary code, APIs, and documentation) to be deployed across Europe without having to create costly, country-specific software versions.

Because many non-EU suppliers adapt products and internal processes to these harmonized standards to retain EU market access, the framework often projects outward, shaping practices well beyond Europe—a "Brussels Effect" that encourages firms to adopt the "world's toughest standard" as a default to avoid region-specific variants (Engler 2023). In anticipation, American and Chinese firms have expanded AI governance, conducted risk assessments, trained staff on data-handling and model-use policies, and staged deployments until they were confident of alignment with these standards before scaling (Engler 2023; Crawford 2021).

Workday, a U.S.-based enterprise software company headquartered in Pleasanton, California, provides a concrete example of how an American firm proactively adapted to EU standards to maintain market access and avoid compliance pitfalls.<sup>186</sup> As a provider of AI-enhanced HR and finance platforms used by global clients (including many in the EU), Workday expanded its AI governance framework by integrating ISO/IEC 42001 for AI management systems, which emphasizes structured risk management and documentation; this involved creating detailed model cards for transparency, embedding bias-mitigation protocols in development pipelines, and implementing watermarking for generated content to comply with limited-risk transparency rules. Workday also trained over 1,000 staff members on EU-specific data-handling policies, ethical AI use, and incident reporting through mandatory certification programs, while staging deployments of new features—such as AI-driven talent analytics—in regulatory sandboxes to test alignment with conformity assessments, ensuring no region-specific variants disrupted global scalability.

---

<sup>186</sup> This paragraph draws on Foeth 2025a; Workday, Inc. 2023; and Workday, Inc. 2024.

## EMERGING REGULATORY ISSUES AROUND A.I.

Like previous industrial revolutions, the increasing deployment of AI throughout the early 2020s sparked a fierce political struggle over privacy, labor, and corporate capture, with sharp swings between extremes. Take education. Early in the AI boom, districts in New York and Los Angeles banned tools like ChatGPT, fearing a "cheating epidemic." By 2025, however, the strategy had shifted from prohibition to negotiation. The turning point came in July of that year, when the American Federation of Teachers (AFT)—representing 1.7 million educators—struck a landmark partnership with OpenAI and Microsoft to launch the National Academy for AI Instruction (Chalkbeat 2025).

This represented a notable shift in how organized labor engaged with AI. Instead of resisting automation, the AFT moved to shape its deployment. The union's strategy was clear: if AI was inevitable, teachers would be its pilots, not its victims. The AFT secured guarantees that AI would be deployed as a "co-pilot" to reduce administrative drudgery—lesson planning, grading, and IEP generation—rather than as a substitute for instruction. As AFT President Randi Weingarten noted, the goal was to "harness" the technology to humanize the profession, buying back the hours teachers lost to bureaucracy (AFT 2025).

Yet the labor question was only one front in a broader political struggle. By late 2025, states like Ohio and Tennessee had begun passing mandatory AI disclosure laws, forcing districts to declare exactly how student data was being used and whether it was shared with third-party model providers (Agile Education 2025). The resulting tension—between the learning gains of personalized AI and the surveillance risks of algorithmic schooling—ensured that education remained one of the most contested frontiers of the AI economy.

Education was not the only sector where opposition to AI emerged. The breakneck pace of data center construction documented in the book generated local resistance. In Memphis, xAI's use of gas turbines to power its Colossus facility sparked complaints about air pollution, prompting local candidates critical of the project to enter 2026 elections. As opposition mounted, xAI shifted subsequent infrastructure across the state line into Mississippi—a pattern of jurisdictional arbitrage that would become increasingly common as communities grappled with the environmental footprint of AI's physical layer (Wayt 2025).

This resistance was emblematic of a broader pattern: across domains, the commercialization of AI outpaced the regulatory frameworks designed to govern it. Given that AI-powered applications and systems thrive on data, it was only natural that issues around data ownership, informed consent, anonymization techniques, data security protocols, misinformation, errors, and transparency rose to the surface. And as AI became integral to everything from creative content generation to critical decision-making, the lines between original data, algorithmic processing, and the resulting outputs became increasingly blurred, creating significant legal challenges. Several regulatory blind spots affected liability rules, intellectual property, privacy, and data security regimes (Buiten 2019; Schaake 2021).

Moreover, governments grappled with how traditional legal frameworks could address harm caused by AI products and services across diverse applications such as customer service,

healthcare, legal services, and autonomous vehicles. All the while, antitrust issues about Big Tech firms' size, digital ecosystems, and market power also loomed in the background.<sup>187</sup> While several implemented A.I. governance mechanisms, others took the drastic step of outright restricting or banning generative A.I. tools due to data-privacy and security concerns (Naidu et al. 2023). Furthermore, the cross-border flow of data is essential for many AI applications, which may require international cooperation and, potentially, the harmonization of data privacy regulations (Aaronson and Leblond 2018).

Consider liability for the moment and the “black box problem”—the combination of technical non-transparency and legal proof difficulties that arise when AI systems are powerful but not meaningfully explainable (Hubbard 2023). Consider that, as explained in the previous section of this appendix, many modern AI systems, especially complex machine-learning and deep-learning models, make decisions in ways that are opaque even to their creators. Users can usually see the inputs and outputs, but they cannot trace the internal steps the system takes to reach a particular decision, such as why an algorithm flagged one applicant as high-risk or caused a vehicle to swerve rather than brake (Burrell 2016). This opacity creates serious accountability and evidentiary challenges: it is hard to diagnose errors, detect bias, or prove that a specific defect in the system was the proximate cause of a harm (ibid). In liability contexts, that means injured parties face an uphill battle to show “defect” and “causation” when the relevant technical evidence is locked inside an inscrutable model controlled by the producer (Hubbard 2023).

By 2025, the stakes around the outstanding regulatory issues outlined above were high. Fears and challenges around liability and IP, for example, represented potentially prohibitive costs that may deter risk-averse organizations from widely adopting AI going forward. Specifically, without clear safe harbors or mature insurance markets, the threat of “ruinous” litigation—whether for copyright infringement or algorithmic discrimination—disproportionately burdened smaller innovators who lack the massive legal war chests and diversified revenue streams of the tech giants. Unless new laws and regulations fill this breach after 2025, the resulting dynamic may threaten to create a two-tier AI economy where only entities with sufficient resources to absorb liability risks can fully leverage AI capabilities, leaving smaller organizations and individuals in the lurch.

While traditionally in the U.S. Congress would step in to level this playing field, the tools to address these problems changed drastically between 2016 and 2025. The technocratic, evidence-driven, expertise-anchored model of rulemaking that characterized the late-20th and early-21st centuries—especially in economically and technologically dynamic domains—was increasingly displaced by an emergency/statutory carve-out model, reactive rulemaking, and “vertical consolidation” of executive power. This shift was not isolated to single domains (trade, AI, pandemic, environment), but was a general pattern, and was accelerated by crisis and turbocharged during the second Trump administration as the US shifted towards a more powerful

---

<sup>187</sup> There is also the issue of bias and fairness, as A.I. algorithms can perpetuate and amplify existing societal biases, leading to discriminatory outcomes. Regulators have focused on addressing algorithmic bias in areas like hiring, lending, and criminal justice by elaborating standards for bias testing and fairness audits. Exploring these issues is beyond the scope of this chapter, however.

and patrimonial executive branch. One important example is the resurgence of populist antitrust thought, which advocates for stricter regulation of tech giants.

During this period, other polities were differently positioned to address these problems. If the U.S. model devolved from the "Creative Destruction Paradigm" into a messy, transactional form of patrimonialism—where outcomes depend on proximity to political power—China consolidated a unique brand of "State-Centric Managerialism." Unlike the American approach, which remained largely reactive in 2025 and relied on ex post litigation to address harms, Beijing's model was fundamentally proactive and preventative, treating AI not as a neutral market good but as a political instrument that must be aligned with state goals before it reaches the public. And where the U.S. struggles with a chaotic patchwork of state laws and federal inaction, China built a comprehensive, vertically integrated regulatory stack—from data labeling standards to algorithm filings—that enforced ideological conformity ("core socialist values") as a prerequisite for market entry. Meanwhile, the European approach consolidated into a "Regulatory-First" model. Driven by the "Brussels Effect," the EU sought to export its values of safety, privacy, and fundamental rights through comprehensive, precautionary legislation like the AI Act, betting that high trust would eventually yield high quality—even if it meant sacrificing the speed of initial commercialization.

As this section of the appendix shows, these diverging regulatory philosophies are not merely academic differences; they are actively reshaping the technology itself, forcing global companies to fragment their products into distinct, region-specific versions—a "GDPR AI" for Europe, a "State-Approved AI" for China, and a "Liability-Shielded AI" for the United States (Jin and Wagman Zhong 2025).

## US APPROACH TO REGULATION AND AI REGULATION

The start of the AI era coincided with a move away from the Creative Destruction Paradigm (CDP) era of evidence-based regulation. In Chapter 3 of the book I reviewed the legislative, executive, and judicial basis behind the CDP in terms of evidence-based policy that pursued innovation or was at least not unfriendly to it.

To remind readers, beginning with the APA's requirements for notice-and-comment rulemaking and a reasoned administrative record, the groundwork was laid for systematic analysis of regulatory impacts. Until the 2020s, cost-benefit analysis in federal rulemaking was largely an internal discipline imposed by the White House rather than a binding legal test. Under Executive Order 12866 (see Clinton 1993) and its companion, OMB Circular A-4 (see OMB 2003), agencies had to prepare a "RIA" for any economically significant rule—quantifying expected benefits (e.g., lives saved, pollution avoided) and costs (e.g., compliance expenditures, paperwork burdens) and demonstrating that benefits justify the intervention. Although courts reviewed whether an agency had followed the APA's procedures and offered a reasoned explanation, they generally deferred to the agency's statutory interpretation under *Chevron USA Inc.* (1984).<sup>188</sup> That meant agencies could defend rules by showing only that their interpretation

---

<sup>188</sup> As I will explain ahead, after *Loper Bright Enterprises* (2024) the Supreme Court overruled *Chevron* deference.

was “reasonable,” even if the underlying statutory authority was ambiguous, and courts seldom second-guessed the merits of an agency’s CBA so long as it wasn’t arbitrary or capricious.

## **The End of the CDP?**

The second Trump administration pursued several regulatory approaches and policies that may portend a shift from a Weberian public-service bureaucracy toward a patrimonial administration, in which loyalty to the President—rather than impartial expertise—becomes the primary qualification for office. During the first year of Trump’s second term, career constraints and independent norms were stripped away and decision-making centralized in the Oval Office. Specialized agency expertise and policy continuity faded too. The cost-benefit approach to policymaking and regulation also seemed endangered.

### *How Things Got to this Point*

The first Trump Administration (2017–2021) implemented several major trade and immigration measures that bypassed traditional OIRA/Circular A-4 cost-benefit analysis by using “emergency” and statutory carve-outs. For example, in March 2018, President Trump imposed 25 percent steel and 10 percent aluminum tariffs under Section 232 of the Trade Expansion Act of 1962, citing national security threats (United States 2018a; United States 2018b). This provision authorized unilateral presidential action without a formal Regulatory Impact Analysis. Similarly, beginning in July 2018, the Administration levied broad duties on Chinese imports under Section 301 of the Trade Act of 1974 (Office of the U.S. Trade Representative 2018), which requires a USTR investigation but contains no mandate for a comprehensive RIA. Rather than negotiating reciprocal concessions, the administration levied duties by presidential proclamation, claiming these statutory delegations empowered the President to bypass the longstanding GATT/WTO principle of reciprocity and the TPA process (Casey 2024; Murrill 2018). In each case, the administration argued that these national-security or “unfair practice” statutes provided their own decision frameworks, making a full cost-benefit analysis legally unnecessary (Harrell 2019; White House 2025).

Similarly, during a wave of financial deregulation beginning in 2017, agencies increasingly abandoned comprehensive cost–benefit analysis in favor of administrative flexibility.<sup>189</sup> Key financial regulatory rollbacks—such as revisions to stress test thresholds, tailoring rules, and Volcker Rule modifications—were often enacted via interim final rules or agency guidance, justified by assertions of market instability or international competitiveness (GAO 2024).

The first Trump Administration also invoked narrow public health and agricultural exemptions. This occurred most notably in March 2020, when the CDC issued an order under the Public

---

<sup>189</sup> Compare this to the aftermath of the 2008 financial crisis: the Dodd-Frank Act introduced “living wills” and periodic systemic risk stress tests, requiring major financial institutions to conduct detailed self-assessments and resolution planning through formal notice-and-comment rulemaking and rigorous economic modeling. These measures imposed significant informational and oversight obligations, compelling banks to prepare systematic plans for rapid and orderly resolution under bankruptcy, aimed at mitigating systemic risk and reducing moral hazard.

Health Service Act to authorize border expulsions to control COVID-19 (Centers for Disease Control and Prevention 2020).

The Biden Administration (2021-2025) continued many of these approaches. It maintained both the Section 232 steel/aluminum and Section 301 China duties without fresh OIRA analysis and extended the Title 42 expulsions through most of 2022, again outside the OIRA framework. The closest this administration came to a traditional analysis for an emergency rule was in late 2021, when the Occupational Safety and Health Administration (OSHA) issued its Vaccine-or-Test Emergency Temporary Standard (ETS). While this was accompanied by an economic analysis of compliance costs and benefits, OSHA nonetheless invoked its emergency authority rather than undertaking a full Circular A-4 RIA (see Occupational Safety and Health Administration 2021).

Indeed, during the COVID-19 pandemic, virtually no public health regulations underwent the comprehensive cost-benefit analysis typically required under Circular A-4, with only a handful of tentative techno-economic snapshots as exceptions (Coghlan 2025). Practically, pandemic timelines were incompatible with Circular A-4's hallmark requirements—multi-year benefit projections, sensitivity analyses across discount rates, and side-by-side comparisons to baselines—making full analyses infeasible in fast-moving conditions (ibid 2023). In short, emergency legal authorities and urgent public-health needs displaced the detailed Circular A-4 process that accompanies normal federal rulemaking.

The primary legal justification for bypassing traditional regulatory analysis stemmed from emergency exemptions in federal administrative law: the APA's "good cause" provision allowing agencies to forgo notice-and-comment during urgent threats to public health or safety (Coghlan 2025); and Executive Order 12866's carve-out for emergencies addressing "serious and unforeseen risks to public health," which permits proceeding without the usual OIRA review (see Clinton 1993). Since most COVID interventions qualified for these emergency paths, agencies did not prepare full, Circular A-4-style Regulatory Impact Analyses (OMB 2023).

Federal agencies further sidestepped traditional analysis by relying on non-regulatory instruments—guidance documents, recommendations, and emergency orders—which are generally exempt from the APA's notice-and-comment process (see Legal Information Institute 2025).<sup>190</sup> This enabled rapid deployment (e.g., travel-masking orders) without months of modeling (OMB 2023). However, the Biden administration's reasoning for circumventing cost-benefit analysis helped cement a momentous movement away from the CDP.

## **The Second Trump Administration**

During 2025, the second Trump Administration doubled down on using statutory "emergency" authorities to bypass traditional OIRA review. It issued new, economically significant Section 232 tariff proclamations for aluminum and steel, invoking the Trade Expansion Act's national-security authority (see US Congress 1962; United States 2025a; United States 2025b). It also launched a formal inclusions process and related national-security reviews to sweep additional

---

<sup>190</sup> It is nonetheless the case that courts sometimes reclassify "guidance" as a substantive rule if it has binding effect.

derivative products and critical minerals into 232 coverage (BIS 2025). Separately, on April 2 (“Liberation Day”), it unveiled the sweeping “reciprocal tariffs,” invoking the International Emergency Economic Powers Act (IEEPA).<sup>191</sup> Both sets of actions were implemented via presidential proclamations rather than agency rulemakings, allowing them to proceed outside the cost-benefit requirements of EO 12866 (Burkhart and Hammond 2025).

Moving beyond trade policy, during its first year in office, the Trump Administration unleashed an unprecedented drive to bring every corner of the executive branch firmly under presidential control. Indeed, while some legal scholars euphemistically referred to President Trump’s power grab as “vertical” or hyper-presidentialism (see Ginsburg 2015), these machinations resulted in a radical reimagining of the executive branch and how regulation is done in the US. Table 14.1 summarizes the various measures the administration pursued under that umbrella: It outlines ten key mechanisms by which presidential authority was expanded at the expense of Congress, the courts, and independent agencies during its first year in office. Below I discuss a selection of these and some others.

First, the second Trump administration recast the role of career civil servants and purged several of them. In a January 20<sup>th</sup> EO reinstating "Schedule F," it reclassified many long-serving employees as “policy/career” staff—subject to removal if they failed to “faithfully implement” White House priorities (Balsler et al. 2025). Thousands of provisional hires were summarily fired, permanent staff were lured out with buyouts, and several inspectors general—Congress’s own embedded watchdogs—were dismissed without the statutorily required thirty-day notice. These moves bypassed traditional merit-system safeguards and eroded the neutral competence of the bureaucracy (US Congress 1946).

Second, the White House effectively nullified the DOJ’s Office of Legal Counsel (OLC) as an internal check. Historically, OLC opinions had carried binding weight across the executive branch, subject only to review by the Attorney General or President. Instead, the administration sidelined OLC memoranda in favor of ad hoc White House legal memoranda, operating under the slogan, “If the President wants it, it’s lawful.” This collapse of independent legal review deepened the unitary-executive dynamic noted in the EO requiring all federal agencies to follow White House or Department of Justice interpretations of their own statutes and magnified the rhetoric-driven erosion of judicial oversight (see Seila Law LLC 2020).

A third, pivotal shift occurred when longstanding Supreme Court precedents safeguarding independent-agency tenure were effectively suspended. Although Humphrey’s Executor (1935) had insulated multi-member commissions—like the National Labor Relations Board (NLRB) and the Merit Systems Protection Board (MSPB)—from at-will removal, the administration fired Democratic appointees to both bodies. Lower courts initially blocked those dismissals under Humphrey’s rule, but a divided D.C. Circuit panel reinterpreted the NLRB and MSPB as wielding “considerable executive power,” and thus subject to presidential removal. When the

---

<sup>191</sup> The reciprocal tariffs were a sweeping set of duties targeting over 90 nations and included a new 10% baseline on all U.S. imports, plus additional, higher duties on approximately 60 countries, including 34% for China, 46% for Vietnam, and 49% for Cambodia. These moves were promptly met with retaliatory counter-tariffs from China, the EU, and Canada.

Supreme Court, in *Trump* (2025), issued an emergency stay permitting the dismissals to take effect, it sent an unsigned signal that Humphrey’s protections were exceptions to rather than the rule of Article II removal power—echoing and extending the logic of *Seila Law* (see *Seila Law LLC* 2020; *Humphrey’s Executor* 1935). And in a later emergency order, the court permitted removal of three Consumer Product Safety Commission members as well (Kruzel 2025). If the Court upholds the D.C. Circuit’s reasoning, it will enshrine near-absolute presidential authority over bureaucratic appointments.

Finally, beyond personnel and legal theory, the administration wielded its budgetary and enforcement powers like blunt instruments (see Table 14.1). It impounded congressionally appropriated funds in violation of the Impoundment Control Act of 1974, threatened to abolish the Consumer Financial Protection Bureau (CFPB), the United States Institute of Peace (USIP), and the National Endowment for the Humanities (NEH) without new legislation, and consolidated discretionary spending in the Oval Office (Impoundment Control Act 1974). Simultaneously, it directed federal agencies—from the Pipeline and Hazardous Materials Safety Administration (PHMSA) to the Office of Federal Contract Compliance Programs (OFCCP)—to drastically curtail or suspend enforcement actions: pausing pipeline-safety cases; stripping pollution-control initiatives at the Environmental Protection Agency (EPA) of funding; and ordering the DOE to stand down on water-efficiency standards for household appliances.

**Table 14.1 The Trump Administration’s Efforts to Consolidate Executive Power**

<b>Mechanism</b>	<b>Example Action</b>	<b>Legal/Institutional Constraint Circumvented</b>	<b>Effect on Checks &amp; Balances</b>
<b>Coercion of Private Actors</b>	Pressuring law firms to pledge pro bono work under threat of barring them from government buildings	Violates free-association and imposes extra-legal sanctions	Divides elite professions and chills independent representation
<b>Political Pressure on Universities</b>	Forced resignations (e.g., UVA President) and shifting compliance demands after campus protests	Undermines institutional autonomy	Weakens independent scrutiny and campus free expression
<b>Undermining “For-Cause” Protections</b>	Summary firing of NLRB Chair Gwynne Wilcox despite Humphrey’s <i>Executor</i> precedent requiring “for-cause” removal	Contradicts Humphrey’s <i>Executor v. U.S.</i> (1935)	Threatens insulation of independent agencies from political whims
<b>Unitary-Executive Orders</b>	EO directing agencies to follow White House/AG legal interpretations of their organic statutes	Erodes Chevron deference and agency expertise	Centralizes rule-making authority in the President
<b>Impoundment of Appropriations</b>	Withholding funds Congress has appropriated (violates the Impoundment Control Act of 1974)	Violates Impoundment Control Act	Shifts budget power from Legislature to President

<b>Mechanism</b>	<b>Example Action</b>	<b>Legal/Institutional Constraint Circumvented</b>	<b>Effect on Checks &amp; Balances</b>
<b>Dismantling Independent Entities</b>	Attempts to abolish CFPB, U.S. Institute of Peace, NEH without new legislation	Circumvents congressional “take-care” and appropriation functions	Weakens legislative oversight of federal programs
<b>Removal of Inspectors General</b>	Firing multiple IGs charged with agency oversight	Violates statutes protecting IG independence	Eliminates internal watchdogs and reduces transparency
<b>Tariffs via Emergency Declarations</b>	Unilateral imposition of steep tariffs on allies citing “national emergency”	Circumvents Trade Act constraints and congressional oversight	Concentrates foreign-policy & trade powers in the executive
<b>Rhetorical Undermining of Courts</b>	Public attacks on judicial orders and slow compliance, coupled with elimination of universal injunctions	Erodes judicial ability to protect broad classes of litigants	Weakens the judiciary’s practical check on executive action
<b>Foot-Dragging on Court Orders</b>	Delayed or partial compliance with injunctions (e.g., deportation reversal of Kilmer Ábrego García)	Runs against district/circuit rulings; exploits narrowing of injunctions	Undermines timely judicial enforcement and remedies

Notes: Each row in the table corresponds to a discrete mechanism by which the second Trump Administration “vertically consolidated” power—expanding executive control over personnel, rule-making, budgeting, enforcement, and judicial review. The selection criteria included only those actions that involved a clear executive act (executive order, firing, directive, or litigation strategy), bypassed or challenged an existing statutory or constitutional check, and affected either internal agency autonomy or external private actors. The mechanisms are grouped conceptually across several categories: personnel control (rows 1–3), legal-jurisprudential control (rows 4, 9–10), budgetary authority (row 5), agency structure (row 6), watchdog removals (row 7), trade and emergency powers (row 8), and enforcement roll-backs (row 10 extension). The evidence base for each entry draws on contemporaneous executive orders, federal statutes, major Supreme Court and D.C. Circuit decisions, and documented agency memos or reports. Wherever possible, primary legal sources were used, while enforcement-drop metrics rely on internal agency dashboards and credible press accounts.

Sources: Bravin and Timiraos (2025); Chevron U.S.A. Inc. (1984); Free Enterprise Fund (2010); Ginsburg (2025); Humphrey’s Executor (1935); Impoundment Control Act (1974); Loper Bright Enterprises (2024); Office of Management and Budget (2003); Seila Law LLC (2020); US Congress (1946); US Congress (1962); EPA (2022).

The second Trump administration also adopted a broad “don’t enforce” strategy across multiple departments. PHMSA initiated only five safety-violation cases—a 92 percent drop from the same period in Trump’s first term—after a new legal directive required all enforcement actions to be routed through a single official (Cantwell 2025). EPA enforcement also declined significantly, with watchdogs reporting a sharp drop in new civil and criminal cases initiated by

the agency (Na 2025). At the Department of Labor’s Office of Federal Contract Compliance Programs (OFCCP), staff were ordered to “cease and desist” enforcement under EO 11246 (U.S. Department of Labor 2025), despite binding civil-rights mandates, following Secretary’s Order 03-2025 (Mitchell and Trotta 2025). The DOE was instructed to halt implementation of longstanding water- and energy-conservation standards for household appliances, including tankless water heaters and air conditioners (NAHB 2025). Internal directives empowered political appointees to override or reprimand career enforcement officers deemed to have “over-enforced” regulations (Cantwell 2025). Critics argued these directives violated the APA and amounted to an unprecedented assertion of unilateral executive power, chilling enforcement and leaving vital protections dormant (Callahan et al. 2025).

### **But there were Offsetting Pressures in Favor of CBA**

There is an argument to be made that there are real limits to the second Trump administration’s attempts to weaken independent agencies, disregard CBA, and succeed in its attempt to circumvent orthodox administrative procedures. As Table 14.1 shows, in 2024 the Supreme Court overruled the four-decade-old Chevron deference doctrine (see *Chevron U.S.A. Inc.* 1984), holding that, instead, judges must independently interpret ambiguous statutes rather than defer to an agency’s “reasonable” reading of the law (Loper Bright Enterprises 2024). Specifically, this ruling empowers courts to decide for themselves whether an agency’s interpretation of ambiguous statutory text aligns with Congress’s intent—so agencies can no longer rely on a “reasonable” interpretation safe harbor. Moreover, under the major-questions doctrine affirmed in *West Virginia* (2024), any rule with vast economic or political significance now demands “clear congressional authorization” before it can proceed.<sup>192</sup>

Together, these two watershed changes may ironically end up strengthening the cost benefit approach to policymaking in some ways. Future CBAs may be more firmly grounded as they will have to not only quantify benefits and costs, but also to map those quantitative findings back to unmistakable statutory language. Agencies will have to build a dual record—showing that (1) the economic trade-offs favor regulation and (2) the statute unambiguously empowers them to act on those trade-offs (Barczewski 2025). Judges, in turn, will have to scrutinize the CBA as a piece of evidence about the rule’s net societal effect and to evaluate the agency’s fidelity to Congress’s design. The upshot is that cost–benefit analysis may become a central battleground in judicial review rather than a mere internal checkpoint (Sunstein 2024).

Along these lines, consider that the “major questions” doctrine (limiting agencies’ ability to assert sweeping power without clear congressional authorization) and the rollback of Chevron deference (no longer forcing courts to defer to agency interpretations of ambiguous statutes) were wielded against the second Trump administration early in his second term.<sup>193</sup> Lower courts used the major questions doctrine to block Trump’s freezes on spending for Biden’s Inflation

---

<sup>192</sup> In *EPA* (2022), the Court formally invoked the “major-questions” doctrine for the first time—requiring agencies to show clear congressional authorization before regulating matters of vast economic and political significance—and struck down the Clean Power Plan on the ground that Congress had not plainly granted the EPA such sweeping power.

<sup>193</sup> This paragraph draws closely on Bravin and Wolfe 2025.

Reduction Act and other signature programs, as well as to enjoin White House orders targeting “woke” initiatives and foreign-aid funding (New York 2025). In May of 2025, a three-judge panel of the U.S. Court of International Trade unanimously struck down Trump’s global tariffs, imposed under the International Emergency Economic Powers Act (IEEPA). The court held that IEEPA authorizes targeted sanctions to address specific foreign threats, not the broad “blunderbuss” use of tariffs to pressure countries on assorted policy goals. It invoked both the major questions doctrine and *Loper Bright v. Raimondo* (2024), which curtailed Chevron deference (see *Chevron U.S.A. Inc.* 1984). Portions of that ruling were later stayed on appeal and the Supreme Court heard consolidated oral argument on November 5, 2025 (see Knauth and Wiessner 2025; U.S. Supreme Court 2025).

### **The Return of Populist Antitrust**

In Chapter 7 of the book, I laid out the political economy and history of the Rule of Reason approach to antitrust. I discussed how changes to U.S. antitrust law catalyzed the rise of digital platforms and AI. There, I traced its evolution from early 20th Century populist approaches, which sought to curb corporate dominance through per se prohibitions of various business strategies, to the adoption of the consumer welfare standard via the Rule of Reason approach (Bork 1978). Contrary to the conventional narrative attributing this shift to the Chicago School, however, the chapter argued that it was primarily driven by courts and policymakers responding pragmatically to the complexities of real-world cases, decades before intellectuals such as Bork, Stigler, and Posner warned against capricious government intervention. Antitrust decision-makers, faced with concrete challenges, increasingly drew on insights from price theory, industrial organization, game theory, and Schumpeterian economics, leading to a more nuanced understanding of market power, efficiency, and innovation. As antitrust decision-makers worked to refine market definitions and analytical tools to address the distinct dynamics of network effects and data-driven business models, these insights were then applied to multisided platforms like Google, Facebook, and Amazon.

The neo-Brandeisian antitrust movement, which perhaps reached its apogee during the Biden administration, seeks to replace the traditional consumer welfare standard with a broader mandate that focuses on protecting competition itself, emphasizing the dangers of “bigness” in both economic and political terms. Proponents argue that the original intent of antitrust laws was to limit the concentration of corporate power, a focus they believe has been lost. Lina Khan, a prominent figure in this movement and the FTC Chairperson under President Biden, argued in her influential article, “Amazon’s Antitrust Paradox,” that the narrow focus on consumer prices has allowed monopolistic practices to flourish (see Khan 2017). This view, shared by other neo-populists like Tim Wu, suggests that the sheer size of tech companies can pose a threat to the balance of power between government and private entities, regardless of their immediate effect on prices (see Wu 2018).

While anti-Big Tech sentiment may reflect legitimate concerns about accountability and power—such as the potential for algorithmic bias, the concentration of data in the hands of a few companies, and the impact of misinformation on democracy—some antitrust enforcers have recast themselves as populist champions of “the little guy,” demonizing scale per se rather than focusing on measurable harm to consumer welfare. This shift is embodied in Wu’s (2018) view

that certain forms of monopolistic behavior, vertical restraints, exclusive dealing, and specific instances of vertical integration should be treated as inherently anticompetitive and harmful to democracy, and should therefore be considered illegal per se.

This is far from an academic concern, as the resurgence of this populist approach garnered bipartisan support across the first Trump administration and Biden administration. The House Judiciary Committee’s 2020 report on Big Tech called for a sweeping overhaul of antitrust law, and politicians from Democratic Senator Elizabeth Warren to Republican Senator Josh Hawley criticized the economic and cultural power these companies wield (see US House Judiciary Committee 2020). This emerging consensus helped legitimize a slew of antitrust actions and was embodied in the 2023 Merger Guidelines, which place far greater scrutiny on mergers that could increase market concentration, create an “entrenchment” of a dominant firm, or cause non-price harms to labor markets.

Neo-populism of this sort does violence against the Rule of Reason approach to evaluating competitive effects on a case-by-case basis, even if it does so when adjudging large corporations with considerable market power. It marks a clear departure from the more technocratic, evidence-driven era of antitrust enforcement that dominated from 1979 through 2019. Indeed, in 2023, a bipartisan coalition of 17 former FTC and DOJ chief economists publicly denounced the Biden administration’s merger guidelines for flouting “consensus economic understanding,” noting that the rules do not even cite the still-binding consumer-welfare test used by the courts (Baker et al. 2023).

Practically-speaking, this ideological battle resulted in two co-existing—and sometimes colliding—enforcement paradigms. On one hand is the traditional, evidence-driven “technocratic” school, which zeroes in on measurable consumer-welfare harm and advocates for the Rule of Reason (see Chapter 7 of the book). On the other is the neo-Brandeisian or “populist” school, which treats corporate size and power as inherently suspect and often pushes for structural remedies like breaking up firms. Many of the recent cases against Big Tech firms represent a “hybrid” approach, blending detailed economic proof with broader concerns over platform power. Table 14.1 compares antitrust lawsuits brought against digital platform-owning technology companies during the late 2010s and early 2020s and places them in these different categories; I discuss the major cases below.

### *Major Lawsuits Against Big Tech Firms*

The FTC’s investigation into Amazon is a prime example of the populist approach. The commission’s inquiries, which culminated in a lawsuit filed in September 2023, focused on the dangers of platform power, such as Amazon’s alleged use of third-party seller data to launch competing products and its acquisition of MGM to entrench its dominance (see FTC 2020a; Amazon.com, Inc. 2023: 3). These actions were targeted as presumptively anticompetitive, reflecting a focus on Amazon’s size and vertical integration, rather than a need to quantify specific price harms to consumers. Separately, in September 2025 Amazon agreed to a \$2.5

billion settlement with the FTC, consisting of both a civil penalty and consumer refunds) over Prime sign-up “dark patterns” (FTC 2025b).<sup>194</sup>

Another example along these lines is the FTC’s December 2022 lawsuit to block Microsoft’s acquisition of Activision Blizzard (see Microsoft, Corp and Activision Blizzard, Inc. 2023). The FTC argued that the merger would harm competition in the gaming industry by giving Microsoft an unfair advantage in its Xbox console and cloud gaming services. However, the court ultimately ruled against the FTC in July 2023, allowing the merger to proceed.

Similarly, the FTC’s December 2020 lawsuit against Facebook (Meta) advances a populist claim that its buying Instagram (2012) and WhatsApp (2014) were “killer acquisitions” meant to neutralize nascent rivals and maintain its monopoly in social networking. The remedy sought—a full unwind of the deals—echoes neo-Brandeisian calls to break up bigness itself, even without direct evidence of price increases or output reduction for consumers (see Facebook, Inc. 2020).

The DOJ’s multi-front assault on Google showcases a more hybrid approach.

The first lawsuit, filed in 2020, accused the company of unlawfully maintaining its dominance in search through exclusionary agreements, most notably its multi-billion-dollar deal with Apple to be the default search engine on the iPhone. While the case leaned on technocratic evidence, the DOJ’s framing invoked populist fears, highlighting “the dangers of allowing a single company to control a critical gateway to information” (Google LLC 2024: 12). In late 2024, a court found Google liable; in September 2025, the court issued its remedies decision, forbidding Google from making agreements that force partners to give Google the same or better terms than any other search provider gets, and restricting deals that make Google the default search engine on devices (Mehta 2025).

A second case targeted Google’s advertising technology business, and in April 2025, a court found Google had violated the Sherman Act and the case proceeded to a remedies phase through fall 2025 (see Google LLC 2025). Anticipating a potential forced breakup, by late 2025 Google had already begun restructuring its ad tech unit to operate more independently and was pushing into new markets like streaming-video ad sales (Perloff 2025).

A third effort, to force Google to divest its Chrome browser and Android operating system, the courts rejected structural breakup, opting for conduct remedies instead (see Swain 2025). In a notable divergence from the U.S. approach, however, European regulators took a much harder and earlier line against these practices. While rooted in traditional theories of monopoly leveraging, the EU’s crackdown on Android strongly aligned with the goals of the emerging neo-populist movement rather than the consumer-welfare standard alone. In a landmark 2018 decision, the European Commission ruled that Google’s mandatory bundling of its apps was

---

<sup>194</sup> “Dark patterns” are user-interface design choices crafted to coerce or trick consumers into making decisions that benefit the company at the user’s expense. In its complaint against Amazon, the FTC alleged the company used such manipulative designs to enroll consumers in Prime without their consent (known as “non-consensual enrollment”) and to make the cancellation process intentionally difficult.

illegal, finding that the company had abused its market power to unfairly stifle competition from rival search engines and browsers. In its decision, the European Commission levied a record €4.34 billion fine against Google. The Commission found that Google had illegally tied its Chrome browser and Google Search app to the license for its essential Google Play Store. As a remedy, Google was forced to unbundle its services in Europe, allowing manufacturers to license its app store without being forced to pre-install and give default status to Google's other products (see European Commission 2018).<sup>195</sup>

Apple faced antitrust scrutiny for its App Store practices, most notably in a lawsuit brought by Epic Games, which also represents a hybrid approach. The suit relied on consumer-welfare economics to challenge a powerful tech gatekeeper—a core concern of the neo-populist movement, which argues that such platform owners unfairly disadvantage competitors and limit consumer choice (Nicas and Wakabayashi 2020). Epic accused Apple of monopolistic practices by forcing developers to use its payment system and charging a 30% commission.<sup>196</sup> After years of litigation, in April–May 2025, the court enforced the injunction requiring Apple to allow external payment links and ruled that Apple remained in violation until it fully complied.<sup>197</sup>

In contrast to these broad, populist challenges, the traditional “technocratic” approach to antitrust also persisted, focusing on classic economic evidence to prove direct, measurable harm to consumer welfare. The Supreme Court’s decision in *Apple Inc. (2019b)* is a prime example. The case did not hinge on Apple's size or its gatekeeper power, but on a highly technical interpretation of antitrust standing and overcharge theory, ultimately allowing a class-action lawsuit from consumers to proceed on narrow economic grounds. Similarly, the DOJ’s case against “Google’s ad-tech” business was a deeply technocratic affair. The government’s arguments relied on sophisticated market-share data and complex economic models to prove that Google’s practices had foreclosed competition and harmed publishers and advertisers (Google 2025c). In both instances, the legal arguments were rooted in measurable economic effects, standing apart from the broader political concerns about corporate power that animate the populist movement.

---

<sup>195</sup> The General Court largely upheld the Android decision in Sept. 2022, and in June 2025 the Advocate General recommended that the Court of Justice dismiss Google’s appeal (see European Commission 2022 and Kokott 2025).

<sup>196</sup> The DOJ also asked a court to void Apple’s multi-billion-dollar agreement with Google—under which Google pays to remain Safari’s default search engine—on competition grounds. If the DOJ succeeds, Apple could lose nearly 20% of its operating profit, according to Bank of America estimates, forcing the company to rethink one of its most lucrative streams of “rental income” from its tightly controlled ecosystem (Winkler and Rattner 2025).

<sup>197</sup> Within days of the ruling, multiple apps added external-payment options. However, most consumers continued to use Apple’s in-app payments, since redirecting to an external website required leaving the App Store and entering payment details. And while some developers incentivized outside payments with discounts, others remained cautious, fearing Apple might retaliate by downgrading App Store search rankings or withholding features from apps that avoid its commissions; moreover, Apple appealed the court decision, though a federal appeals court refused a stay of the order (on all these points see Tilley 2025b).

Finally, consider the government’s complaints against Meta. The FTC’s non-merger case against the company began in December 2020, alleging that Facebook illegally maintained monopoly power in “personal social networking” by acquiring nascent rivals Instagram (2012) and WhatsApp (2014); after an initial dismissal, the FTC refiled an amended complaint in 2021 and the case proceeded to trial in 2025 (see FTC 2020b). However, in a major setback for breakup-oriented enforcement, a federal court rejected the FTC’s bid to force Meta to divest Instagram and WhatsApp in November of that year, holding that the agency failed to prove Meta currently holds monopoly power in a properly defined market and noting competition from platforms like TikTok and YouTube. In other words, the ruling left Meta’s 2012 and 2014 acquisitions intact; it also underscored that despite the executive branch’s overt flirtation with populism during the early 2020s, courts continued to prefer “conduct remedies” or narrow relief over structural divestiture in fast-moving digital markets (Meta Platforms Inc. 2025c).

**Table 14.1 Big Tech Antitrust Cases and Hybridity**

<b>Case</b>	<b>Date / Court</b>	<b>Core Allegation</b>	<b>Evidence &amp; Theory</b>	<b>Remedy Sought</b>	<b>Antitrust Philosophy</b>
<b>Epic Games v. Apple</b>	2020 (N.D. Cal. trial; 2021 9th Cir. appeal)	Apple’s 30% in-app-purchase “tax” and gate-keeping of App Store distribution constitute a monopoly on iOS app commerce.	App Store policies; developer/deployment data; internal Apple memos on “walled garden”; Rule of Reason showing foreclosure of rival payment systems.	Injunctions forcing Apple to allow alternative in-app payment and third-party app stores on iOS.	<b>Hybrid:</b> relies on consumer-welfare economics (harm to developers & consumers) but also taps populist “big tech gatekeeper” rhetoric.
<b>Apple v. Pepper</b>	2019 (U.S. Supreme Court)	iPhone owners, as “direct purchasers,” can sue Apple for monopolistic over-pricing of apps.	Apple’s commission structure; pass-through pricing data; statutory “direct purchaser” standing analysis.	Class-action damages against Apple if monopoly found.	<b>Technocratic:</b> focuses narrowly on classical antitrust standing and overcharge theory under the Clayton Act.
<b>DOJ v. Google (Ad-Tech)</b>	Aug 2024, D.C. District Court	Monopoly in ad-tech: exclusionary agreements making Google the “default” ad-seller	Detailed market-share data, price-impact studies, internal documents on exclusivity deals; Rule-of-Reason analysis of harm to publishers and advertisers	Divestiture of key ad-tech units; injunctions on default-setting practices	<b>Technocratic:</b> consumer-welfare, evidence-driven; judge applied classic economic metrics to define market and assess foreclosure effects

<b>DOJ v. Google (Search)</b>	Filed Oct 2020; hearings mid-2025, D.C.	Monopoly in search services & search ads via default-search deals (Chrome, Apple iOS, Android)	Browser-usage stats; consumer-survey evidence; documents on partner payments; dynamic-competition analysis under Rule of Reason, “major questions” concerns	Divestiture of Chrome browser; limits on default-search accords	<p><b>Hybrid:</b> formally technocratic (price, output, data); DOJ’s framing (“gateway to information”) and talk of “critical infrastructure” invoke populist fear of concentration</p> <p><b>Populist:</b> centers on bigness and threat to competition itself; seeks structural breakup even absent proven price rises or output reduction (per se-style remedy)</p>
<b>FTC v. Meta</b>	Filed Dec 2020, pending in D.C. Circuit	Anticompetitive acquisitions (Instagram 2012, WhatsApp 2014) to neutralize nascent rivals	Internal strategy memos; counterfactual “would-have-grown” analyses; consumer-welfare projections; testimony on barriers to entry	Unwind Instagram & WhatsApp deals; injunctions on future M&A	

Notes: this table provides a comparative analysis of major antitrust actions against digital platforms; it is organized to highlight each case's procedural status, legal arguments, evidence, and proposed remedies. The columns present: the case identification with date and court; a summary of the alleged monopolistic conduct; the key evidence and economic theory supporting the claims; the specific remedies requested to restore competition; and a classification of the case's underlying antitrust philosophy as “Technocratic” (consumer welfare-focused), “Populist” (concerned with size and power), or “Hybrid.” Readers can use this chart to understand not only the specific allegations against each platform, but also to discern whether enforcers are emphasizing detailed economic proof or broader structural arguments about market power, reflecting either traditional consumer-harm frameworks or neo-populist approaches focused on curtailing corporate size and influence.

Sources: see the text.

### *Patrimonial Antitrust Enforcement*

If neo-Brandeisian antitrust re-centers ideas about bigness, early in the second Trump administration a patrimonial mode of enforcement surfaced, in which outcomes appear to hinge less on neutral rules and more on proximity to political power and "national champion" status. While the formal docket still moved forward in terms of Section 2 cases and merger challenges, there was an unmistakable rise in lobbying efforts by well-connected defendants to curry favorable settlement terms or selective forbearance.

This dynamic was epitomized by the administration's pivot on Google. While the DOJ's case continued, President Trump publicly questioned the wisdom of a structural breakup, arguing that destroying such a "powerhouse" would only benefit China. "China is afraid of Google," he noted, suggesting that the firm's scale was a strategic national asset that should be disciplined for "fairness" rather than dismantled (see Weiss and Tangalakis-Lippert 2024). This marked a shift from consumer welfare to mercantilist welfare, where monopolies are tolerated if they serve national (or political) interests. Similarly, the administration's reversal on the TikTok ban, following engagement with major donor Jeff Yass, underscored that regulatory threats could be dissolved through transactional alliances (see Kim and Ibsa 2024).

The Meta litigation discussed in the last row of Table 14.1 further crystallizes this dynamic. On the one hand, the company pressed a classic competitive-constraints defense—framing TikTok and YouTube as vigorous rivals. On the other, it undertook highly visible gestures to placate the new political equilibrium, paring back diversity initiatives and relaxing content moderation. This dual strategy bore fruit on November 19, 2025, when a federal court formally rejected the FTC's bid to break up the company, ruling the agency failed to prove a monopoly in "personal social networking" (Meta Platforms, Inc. 2025c). The ruling left Meta's empire intact, reinforcing the perception that in the patrimonial era, the most effective antitrust defense may be a combination of legal delay and political alignment.

### **U.S. Regulation of A.I.**

As this section of the appendix has made clear, right out of the gate, the second Trump Administration emphasized maintaining and extending American dominance in A.I., with the goal of building "the most powerful AI systems" on U.S. soil (Trump 2025). The administration viewed stringent rules as potential innovation strangling, criticizing the EU's Digital Services Act (Shapero 2025). To secure this deregulated environment, the administration supported Congressional efforts to preempt state-level AI laws (Quinlan 2025). Following the failure of a 10-year moratorium proposal in July 2025, House Republican leaders actively considered attaching broad preemption language to the National Defense Authorization Act (NDAA), arguing that a patchwork of state regulations threatened national security and American competitiveness (Mondeaux 2025). This push culminated in late November 2025, when House Republicans formally circulated—and the White House quietly advanced—a draft executive

order and NDAA amendment package to preempt state AI laws, including an “AI Litigation Task Force” and funding levers to pressure non-compliant states (Alder 2025).<sup>198</sup>

As outlined above, the US federal government’s light approach to AI also fell short of the strictures associated with the CDP. Given the fact that, since 2017, AI regulation by federal agencies has leaned heavily on nonbinding guidance—like NIST’s AI Risk Management Framework or DHS’s voluntary standards for facial recognition (NIST 2023)—these measures were accompanied by little formal OIRA review or economic impact analysis despite potentially significant effects on industry practices and sectoral norms. White House AI executive orders and emergency technology controls (e.g., export controls on AI chips, Section 1758 authority) have been justified on national security grounds, thus bypassing traditional open regulatory processes (BIS 2023).

It is in these policy interstices where high-profile litigation unrelated to AI-specific statutes, but that deeply implicates AI, has bubbled to the surface. An example is the pioneering class-action lawsuit Mobley (2025), where a job applicant alleged that Workday’s AI-driven hiring tool systematically rejected his applications for over eighty positions based on race, age, and disability. What makes this lawsuit particularly significant is its legal foundation. First, rather than proceeding as a tort claim, it constitutes a federal civil-rights action under established statutes including the Age Discrimination in Employment Act, Title VII, and the Americans with Disabilities Act. Second, the court held that Mobley had plausibly alleged that Workday, as an AI vendor, could be treated as an “employer” under these statutes via agency theory, because its tools allegedly performed traditional hiring functions such as screening and rejecting applicants—even though Workday did not directly employ the affected individuals (Mobley 2025; Tyman 2024). The implication is that, even as comprehensive AI legislation stalls at the federal level, courts can still constrain AI applications by applying existing civil-rights statutes and agency doctrines to vendors whose tools perform core decision-making functions.

The absence of a comprehensive federal law regulating AI also created a vacuum that states rushed to fill, resulting in a fragmented regulatory landscape. By late 2025, 47 states had enacted laws specifically targeting deepfakes, focusing on election interference and nonconsensual pornography (Bonatesta 2025). Beyond these commonalities, however, local political priorities drove a diverse array of specific mandates. In Texas, the legislature enacted the "Responsible AI Governance Act," which explicitly bans government agencies from using AI for "social scoring"—a direct response to populist concerns about algorithmic bureaucracy (Nahra et al. 2025). Utah and Tennessee focused on consumer deception and creative rights; Utah’s "AI

---

<sup>198</sup> By late 2025, however, this federal push for deregulation faced a surprising political backlash from within President Trump’s own coalition. A growing rift emerged between the administration’s “innovation-first” alliance with tech billionaires and populist Republicans who viewed unchecked AI as a threat to workers and families (De Vynck 2025). Prominent GOP leaders, including Governors Ron DeSantis and Sarah Huckabee Sanders and Senator Josh Hawley, actively opposed the White House’s preemption efforts, arguing that blocking state-level AI regulation would surrender sovereignty to “massive companies that have hugely concentrated power” (ibid.). This “MAGA rift” created a strange bedfellows dynamic where populist conservatives aligned with progressive Democrats to demand stricter guardrails.

Policy Act" mandates clear disclosure when a human is interacting with an AI chatbot (Levi et al. 2024), while Tennessee's "ELVIS Act" created the nation's first protections against unauthorized AI cloning of a musician's voice (Damle et al. 2024). In California, SB 53 established robust whistleblower protections for employees at frontier AI labs (Tene et al. 2025). In Ohio, Republican lawmakers introduced a novel bill to legally define AI as "non-sentient," explicitly banning humans from marrying AI systems to forestall any future claims of algorithmic personhood (Landymore 2025).

### *Privacy and Data Security*

In this spirit, data privacy in the U.S. is primarily addressed at the state level.<sup>199</sup> This has resulted in a complex patchwork of regulations that range from strict to lax.<sup>200</sup> State laws include provisions for data access, deletion, and opt-out rights regarding the sale or sharing of personal information (Solove and Schwartz 2021). Most US states have limited or sector-specific privacy laws, such as laws focused on education and employment. Many states still lack comprehensive data privacy legislation, leaving residents with fewer protections. For its part, California boasts comprehensive consumer privacy protections granting residents significant control over their personal data.

In this spirit, by 2025, the most "successful" example of a standalone state level attempt to regulate AI was perhaps in the realm of biometrics: Illinois' Biometric Information Privacy Act (BIPA) is an atypically stern state law that requires companies to obtain written consent before

---

<sup>199</sup> To be sure, at the federal level there are specific laws about privacy focusing on areas like healthcare (HIPAA) or financial information (GLBA) and data security that matter for AI. In healthcare settings, for example, AI impacts patient care documentation by transcribing and summarizing doctor-patient interactions for electronic health records, with potential expansion into preliminary diagnoses. This application faces stringent regulatory requirements under HIPAA and similar federal laws, necessitating careful attention to patient consent, data encryption, and sharing limitations. Privacy around medical issues at the federal level remains the exception to the rule, however.

<sup>200</sup> States with consumer privacy laws in 2025 include California: California Consumer Privacy Act (CCPA), Virginia: Virginia Consumer Data Protection Act (VCDPA), Colorado: Colorado Privacy Act (CPA), Connecticut: Connecticut Data Privacy Act (CTDPA), Utah: Utah Consumer Privacy Act (UCPA), Iowa: Iowa Consumer Data Protection Act (ICDPA), Indiana: Indiana Consumer Data Protection Act, Tennessee: Tennessee Information Protection Act (TIPA), Texas: Texas Data Privacy and Security Act (TDPSA), Florida: Florida Digital Bill of Rights (FDBR), Montana: Montana Consumer Data Privacy Act (MCDPA), Oregon: Oregon Consumer Data Privacy Act (OCPDA), Delaware: Delaware Personal Data Privacy Act (DPDPA), New Hampshire: New Hampshire Privacy Act (NHPA), New Jersey: New Jersey's Act Concerning Online Services, Consumers, and Personal Data (NJOPA), Kentucky: Kentucky Consumer Data Protection Act (KCDPA), Nebraska: Nebraska Data Privacy Act (NDPA), Rhode Island: Rhode Island Data Transparency and Privacy Protection Act (RIDPA), Maryland: Maryland Online Data Privacy Act (MODPA), and Minnesota: Minnesota Consumer Data Privacy Act (MNDPA).

collecting biometric data like fingerprints or face scans.<sup>201</sup> Unlike many privacy laws, BIPA allows individuals to sue for statutory damages of \$1,000 to \$5,000 per violation (see Illinois General Assembly 2008). In 2023, the Illinois Supreme Court interpreted "per violation" to mean every single time a person's data was scanned (e.g., every time an employee clocked in) (see Cothron 2023). Because this created the potential for "ruinous" liability—the fast-food chain White Castle, for instance, faced an estimated \$17 billion judgment for a class of just 9,500 employees (Brown 2023)—a 2024 amendment fixed this by clarifying that repeated scans of the same person count as a single violation, drastically lowering the liability cap while preserving the core privacy mandate (Nahra et al. 2025). Nonetheless, even with this reform, BIPA remains the most potent legal check on the deployment of commercial AI in the United States (Hartzog 2020). This has significantly influenced companies to adopt strict consent protocols to avoid the risk of non-compliance and potential legal consequences; for example, Clearview AI ceased its contracts with non-law enforcement entities in Illinois specifically to avoid penalties under the Act (ibid).

Yet, while these targeted laws proliferated, broad "comprehensive" AI regulation faced stiff headwinds beyond the federal government's attempt to muzzle state level initiatives. In California, Governor Newsom vetoed the controversial SB 1047, which would have mandated "kill switches" for frontier models, arguing it would stifle the state's innovation economy (Montgomery and Bhuiyan 2024). Similarly, Colorado passed the nation's first comprehensive AI Act but, following intense industry pushback, was forced to delay its implementation until mid-2026 (Schwartz et al. 2025). This dynamic left the U.S. with a "bifurcated" reality: aggressive state-level policing of specific harms like deepfakes and voice cloning, but a stalled attempt to regulate the fundamental technology itself.

### *Liability Issues Around AI*

During the Third Industrial Revolution, the U.S. exhibited considerable leniency when regulating new digital technologies, often relying on existing law and after-the-fact litigation to handle harms. As already explained in this book manuscript, internet platforms enjoy broad immunity due to Section 230, which, since 1996, "immunized online intermediaries from liability for content posted by others" (US Congress 1996), enabling online speech and commerce to flourish without constant lawsuits. This hands-off approach treated tech intermediaries not as the speaker, but as neutral platforms.

As I explained in Chapter 8 of the book, while Section 230 of the U.S. Communications Decency Act historically provided safe harbor for online platforms, this protection does not extend to A.I.-

---

<sup>201</sup> While BIPA does not explicitly use the term "Artificial Intelligence," it is functionally one of the most consequential AI regulations in the country. This is because modern biometric systems—such as facial recognition, iris scanning, and voiceprint analysis—are powered almost exclusively by computer vision and machine learning algorithms (GovFacts 2025). These AI models work by ingesting physical characteristics, converting them into mathematical vectors (templates), and matching them against databases (ibid). By strictly regulating the collection of the necessary input data (the biometrics), BIPA effectively regulates the deployment of the AI systems that rely on that data to function (Bryan et al. 2022).

generated content, as it represents new material created by the A.I. system rather than third-party content merely hosted on a platform. Indeed, it may not necessarily extend to all algorithmically amplified content produced by third parties. Specifically, some rulings leading up to 2025 signaled a potential shift as they seemed to narrow Section 230's protections and exposed platforms to greater accountability for algorithmic harms.

In *Malwarebytes (2020)*, Malwarebytes, a cybersecurity firm, flagged Enigma's software as "malicious," effectively blocking it from users' computers. Enigma, which is also involved in cybersecurity, sued, arguing this was not a good-faith effort to protect users but an anticompetitive move to sabotage a business rival. In its defense, Malwarebytes invoked the specific part of Section 230 designed to protect services that filter "objectionable" content. After a series of appeals, courts ultimately decided that this protection is not absolute. They ruled that the law's "good faith" requirement has real teeth: if a platform's decision to block content is motivated by anticompetitive malice rather than user safety, it could lose its immunity and be sued for making a false statement.

The case drew national attention when Supreme Court Justice Clarence Thomas, in a rare statement accompanying the Court's decision not to hear the appeal, signaled his own skepticism about the broad immunity platforms enjoy, suggesting courts had interpreted the law too generously. Between the ultimate decision and Thomas's unusual unsolicited decision to weigh in on Section 230 jurisprudence generally, the Malwarebytes saga established that a platform's motives for moderation matter, creating a precedent that its legal shield can be challenged if it's used as a pretext to harm competitors.

Similarly, in *Anderson (2024)*, the Third Circuit Court of Appeals departed from typical Section 230 jurisprudence by reversing a lower court's dismissal when a 10-year-old died attempting TikTok's algorithmically promoted "Blackout Challenge." The court reasoned that TikTok's algorithm wasn't merely a "neutral conduit" but actively pushed deadly content to a vulnerable minor—conduct the court likened to TikTok "telling" the user to attempt the challenge, drawing on First Amendment precedents to classify recommendations as the platform's own expressive activity (not third-party content protected by Section 230). As of 2025, the ruling stood, but was under petition for further review, potentially signaling a broader erosion of algorithmic immunity.

Finally, a 2021 case involving then-President Trump's Twitter account, *Biden (2021)*, upped the Section 230 skepticism ante when the US's highest court weighed in on whether some platforms were public squares that should abide by rules governing public access. While the Supreme Court declared that the specific dispute in the case—whether President Trump's blocking of critics on his Twitter account was an unconstitutional violation of their First Amendment rights—was moot because Trump was out of office, Justice Thomas used the occasion to write a separate concurring opinion, where he warned that an "unprecedented" amount of control over public speech was now concentrated in the hands of a few dominant digital platforms. He then outlined two long-standing legal frameworks that could be used to regulate these platforms. He suggested they could be treated as "common carriers," like telephone companies, which are required to provide service to all without discrimination. Alternatively, he proposed they could be seen as "public accommodations," like hotels or restaurants, which are legally barred from

refusing service to customers. Thomas argued that applying these models could limit a platform's right to exclude users and called on courts and legislatures to reconsider the broad immunity granted by Section 230.

Therefore, it is not a given that Section 230 will cover AI liability issues. Justice Gorsuch's comments in *Gonzalez* (2023) signaled that AI content is potentially outside Section 230's shield. This suggested the Court might be open to allowing AI-related suits to proceed without that immunity. While as of 2025 these AI cases had not reached the Supreme Court, during the latter part of the early 2020s, there was a growing bipartisan push to reconsider Section 230's scope in relation to generative AI (Perault 2023). Senators Richard Blumenthal (D-CT) and Josh Hawley (R-MO) introduced a bipartisan bill (S. 1993) to explicitly eliminate Section 230 immunity for generative AI (Blumenthal and Hawley 2023). Both senators floated a "comprehensive regulatory framework" for AI, including expanded liability for AI developers, a new AI-specific federal agency, licensing requirements for advanced AI systems, and transparency mandates (ibid). At the same time, critics of these initiatives warn that regulatory overreaction could stifle innovation and point to the internet's vibrant commercialization as a reason to allow Section 230 to serve as the legal foundation for AI applications around speech and related content (Kosseff 2019).

In 2025, California proposed updated regulations to allow testing and deployment of heavy-duty AVs, moving to end the de facto ban (California DMV 2025). At the federal level, DOT streamlined exemptions so domestic manufacturers can deploy limited volumes of vehicles without traditional manual controls (like steering wheels) where safety is shown (NHSTA 2025). Additionally, other embodied AI systems such as commercial delivery drones and warehouse robots are also regulated through existing product-safety and aviation frameworks. However, in 2025 legal experts debated how liability and compliance may evolve as these technologies become more autonomous (McKirahan 2025).

Separately, during the early 2020s, regulators like the FTC and DOJ also flexed existing laws and approaches. The FTC repeatedly warned firms about deceptive AI claims and accuracy, including about AI safety or capabilities (Chander 2023). The DOJ and EEOC (Equal Employment Opportunity Commission) likewise stated that discrimination by algorithms is still discrimination—and companies using AI hiring tools that unfairly screened out protected groups confronted enforcement actions. In one case, the EEOC even argued that an AI vendor could be directly liable under civil rights laws if their tool caused biased hiring outcomes (Ball 2025). Additionally, the U.S. AI Safety Institute Consortium (AISIC), established by the Department of Commerce and managed by NIST, brought together hundreds of organizations from industry, academia, and government to develop science-based guidelines for testing, evaluating, and ensuring trustworthy AI systems (NIST 2024).

Returning to AI and biased outcomes, binding state regulations were also ascendant during this period. In 2025, California implemented a new rule clarifying that state anti-discrimination laws apply fully to "automated-decision systems" used in hiring, performance management, and termination. This regulation places a significant compliance burden on employers, requiring them to maintain personnel records for four years and audit third-party tools for bias. Similar

regulatory frameworks were adopted in New York, Illinois, and Colorado, creating a growing state-level liability shield for workers against algorithmic bias (Abril 2025).

In the absence of clear and preeminent federal legislation, therefore, many U.S. companies have undertaken proactive compliance measures—including bias audits, rigorous vendor oversight, and thorough documentation practices—to mitigate exposure to legal risk. Take key players in the financial sector. Mastercard examined generative AI’s implications for banking. In its 2024 industry report, it noted that ensuring data privacy, accuracy, and bias mitigation in AI is critical and that banks’ adoption of AI will depend on meeting global regulatory requirements (Mastercard 2024). Similarly, Goldman Sachs rolled out a proprietary generative AI assistant for its staff in a controlled environment using internal data, with oversight to ensure outputs met compliance standards such as refraining from offering unauthorized financial advice (Son 2025). Big banks and insurance companies also formed internal AI governance councils to track regulatory developments. And Nationwide Insurance’s approach involved a dedicated AI ethics committee and a “red team/blue team” process to continually test the AI’s outcomes for risks (Carrel 2025).

### *Copyright Law and AI*

Meanwhile, the courts clarified the application of copyright law to AI training during this period. As already outlined in Chapter 8 of the book, in two landmark 2025 rulings (see Bartz et al. 2025; Kadrey et al. 2025), federal judges affirmed that training AI models on copyrighted books constitutes “fair use.” However, they drew a hard line on sourcing, ruling that downloading pirated datasets to build those training corpora constituted copyright infringement. A court-approved \$1.5 billion settlement by Anthropic to the plaintiffs in Bartz et al. (2025) subsequently resolved claims tied to pirated datasets, accelerating a broader “clean data” pivot while protecting the underlying act of model training.

However, several cases that were ongoing in 2025 promised to further complicate the copyright safe harbor currently protecting companies that train AI models. Outstanding issues include whether outputs that mimic an author’s style constitute derivative works (Madigan 2025). The New York Times (2025) litigation highlighted the distinct risk of “regurgitation”—where an LLM outputs verbatim or near-verbatim excerpts of paywalled content—complicating the defendants’ reliance on the “fair use” defense, which assumes outputs are sufficiently transformative (Mass Law Blog 2025).

### CHINA’S REGULATION OF AI

In the early days of the internet and e-commerce, China pursued a relatively “hands-off” growth. Chinese tech giants like Alibaba and Tencent thrived in a period of light regulation like Silicon Valley’s, with the government stepping in later to impose rules once the sectors matured (see Webster 2019).

However, when it comes to content and data, China has always exercised tight control. Liability for online content in China is essentially the inverse of Section 230: internet platforms *are* held responsible for user-posted content and must censor prohibited material or face penalties (see

Shao 2012). Laws like the Cybersecurity Law and various content regulations require providers to proactively monitor and remove banned speech, or the government can sanction the company (including fines or shutdown). Consequently, Chinese tech companies have long acted as content gatekeepers, liable if they facilitate dissemination of, say, anti-government messages or pornography. So, whereas Western companies worried about being sued by users, Chinese companies historically worried about punishment by authorities for failing to police users (Jinhe and Yang 2022). This backdrop forms the foundation for how China approaches AI liability: the developer or operator will be held accountable for the AI's behavior, especially its output content.

## **China's AI-Specific Regulations**

Leading into 2025, China had rolled out some of the world's earliest AI-specific regulations. The Provisions on the Administration of Deep Synthesis Internet Information Services, which went into effect in 2023, regulate "deep synthesis" services (AI-generated or manipulated audio-visual content) by requiring providers to label synthetic media and prevent its use to spread illegal or misleading information (CAC 2022). Also in 2023, the Cyberspace Administration of China and six other ministries issued the Interim Measures for the Management of Generative Artificial Intelligence Services (IMG AIS), which govern public-facing generative-AI services by requiring security assessments, algorithm filings, content controls, and multi-agency supervision for systems offered to users in China (CAC 2023).

Although the IMG AIS framework initially appeared highly restrictive, the final version eased several obligations and explicitly framed generative AI as a strategic growth sector, narrowing mandatory security reviews to systems with "public opinion attributes or social mobilization capabilities" and toning down some of the strict performance guarantees—an adjustment widely read as an attempt to balance control with innovation (Sun and Zeng 2025). Even so, providers remained responsible for model outputs and had to complete filings and security assessments before launch, so major platforms such as Baidu and Alibaba were only allowed to deploy ChatGPT-style services after demonstrating robust censorship of "sensitive topics" and safety filters (White & Case LLP 2025).

First, the generative-AI regulations codified stringent content requirements. Providers must ensure that model outputs adhere to China's censorship regime and "core socialist values," filtering or preventing content that is politically sensitive, violent, sexual, false, or otherwise illegal under Chinese law; if a model generates prohibited material, the service provider—not the user—is held responsible and must correct the problem (CAC 2023). This effectively inverts the U.S. Section 230 approach: Chinese generative-AI platforms are treated as accountable publishers rather than neutral intermediaries.

Second, the rules introduced security assessments and algorithm-filing obligations. Early drafts would have required security reviews for virtually all models, but the final IMG AIS text focuses filings and formal assessments on models with public-opinion or mobilization effects—such as large chatbots and content platforms—while still allowing regulators to conduct *ex post* inspections of smaller or enterprise systems (see Sun and Zeng 2025). Providers must disclose basic information about model architecture, training-data categories, and intended use cases,

giving authorities leverage to intervene if outputs are deemed unsafe or politically destabilizing (White & Case LLP 2025).

Third, transparency and labeling duties were strengthened and then generalized. Building on the 2022–23 deep-synthesis rules, which mandated labels for AI-generated audio-visual material, China adopted dedicated Measures for the Administration of Labels for Artificial Intelligence-Generated Content in March 2025, requiring AI-generated text, images, video, and audio that could be mistaken for real content to carry clear labels (Yan 2025). To enable platforms and regulators to authenticate content even if visible labels are removed, the accompanying national standard on AI-content labeling specifies a dual system of explicit indicators (such as on-screen marks, watermarks, or audio tones) and implicit technical markers embedded in metadata or digital signatures (White & Case LLP 2025).

Fourth, China moved to harden AI systems across the full data lifecycle. In April 2025, the National Information Security Standardization Technical Committee issued three national cybersecurity standards addressing (1) security of data annotation for generative AI, (2) protection and governance of pre-training and fine-tuning datasets, and (3) baseline security requirements for generative-AI services (ibid). Together, they call for stricter vetting and logging of annotators, classification and provenance tracking of training data, privacy-preserving techniques for sensitive datasets, secure storage and deletion protocols, and requirements for input validation, output filtering, incident-response plans, and ongoing monitoring of deployed models. Many of these provisions may be formally “recommended,” but for firms operating in regulated sectors they function as de facto mandatory technical benchmarks.

Finally, the IMG AIS measures called for industry-specific regulation. The government directed various sector regulators in China to create sector-specific AI guidelines for their industries. For example, China’s financial regulator issues rules for AI in finance and its healthcare regulator for medical AI.

### *Judicial Developments: Defining Liability through the Courts*

Just as U.S. courts wrestled with Section 230, Chinese courts in the mid-2020s began developing their own case law on AI liability, sometimes moving faster than national legislation. In February 2024, the Guangzhou Internet Court issued the so-called Ultraman decision, widely described as the first effective judgment on generative-AI copyright infringement. The court held an AI-image platform liable because its system generated pictures of the famous Ultraman character that were substantially like the copyright holder’s works and emphasized that providers have a “duty of care” to implement keyword filtering, monitoring, and complaint mechanisms—effectively articulating a negligence standard for AI services (Song and Wang 2024).<sup>202</sup> The

---

<sup>202</sup> Ultraman is a Japanese superhero created in 1966 who transforms into a giant alien to battle massive monsters (kaiju). In East Asia, his cultural ubiquity is comparable to that of Superman in the West; he is an instantly recognizable, multigenerational icon of justice whose likeness appears on everything from toys to government campaigns. The character is particularly massive in China, where he has been a dominant children's franchise for decades, making the unauthorized generation of his image a high-profile test case for copyright law.

court’s reasoning underscored that, unlike in the U.S., Chinese providers cannot invoke broad immunity when their models output infringing content.

Chinese judges also moved more aggressively than their U.S. counterparts on authorship. In a separate Beijing Internet Court case, a user who generated an image with a text-to-image system was granted copyright protection after demonstrating that their prompts and iterative adjustments reflected sufficient “intellectual inputs” and personal choices to qualify as creative contribution (Wang and Zhang 2024). Rather than treating AI outputs as categorically uncopyrightable, the court focused on the human user’s control over the process, suggesting a hybrid model in which AI is merely a tool, but authorship still attaches to a human who can show meaningful creative effort.

Over the early 2020s, enforcement of AI-specific obligations heavily relied on China’s broader digital-governance framework. The deep-synthesis rules and IMG AIS both state that violations will be punished under long-standing instruments such as the Cybersecurity Law, Data Security Law, and Personal Information Protection Law, which already authorize fines, business suspensions, and—in severe cases—criminal liability for executives (Yu and Li 2024). Authorities thus retain wide latitude to shut down services, block access through the Great Firewall, or impose administrative sanctions on providers whose systems disseminate illegal content, leak data, or otherwise endanger “network security” (Su and Zhang 2024).

As AI cases multiplied, Beijing also moved to centralize control over AI jurisprudence. In October 2025, the Supreme People’s Court issued new Provisions on the Jurisdiction of Internet Courts that removed Internet Courts in Beijing, Hangzhou, and Guangzhou from hearing certain online copyright disputes, including key categories of AI-related copyright cases, shifting them to higher-level or more general courts (Supreme People’s Court 2025). Legal commentators interpreted this as an effort to ensure that precedents on AI and copyright are aligned with national policy priorities rather than emerging piecemeal from specialized local tribunals (see Wininger 2025).

Finally, regulation in China has traditionally been accompanied by forward-looking guidance aimed at shaping norms and technical practice and, true to form, national AI governance principles and the AI Safety Governance Framework—updated to version 2.0 in 2025—stress themes such as controllability, safety, fairness, and protection of user rights (Sheehan 2023), while emphasizing that AI must remain “under effective human supervision” (Geopolitechs 2025). In December 2024, the Ministry of Industry and Information Technology announced a 41-member AI standardization technical committee bringing together major tech firms like Baidu and leading universities to develop standards on LLMs and AI risk assessment, signaling Beijing’s intent to be a global rule-maker rather than merely a rule-taker in AI governance (Reuters 2024).

## EUROPE’S REGULATORY APPROACH

Neither has the EU’s regulatory environment helped Europe narrow its innovation gap. Europeans have prioritized an aggressive approach to preventing misinformation, harmful content, and anti-competitive behavior, resulting in strict privacy regulations, strong data security

protocols, and stringent antitrust enforcement (Renda 2021; Crawford 2021). In turn, while the EU—along with important European nation states—espoused a uniquely “European brand” of A.I. focused on safety, open collaboration, and trust, Europe fell considerably behind the United States and China in commercial AI applications (Veugelers 2018). Indeed, by 2025, the gap between Europe’s innovation performance and that of these other polities only widened over time (Spirlet 2025).

Consider the privacy and data security law known as the General Data Protection Regulation (GDPR), which grants consumers broad data rights and places strict obligations on firms (European Union 2016). Although designed to bolster user trust and reinforce data security, during the early 2020s it often led to higher compliance costs for firms and delayed the commercialization of data-intensive services where AI was central, including targeted algorithms or personalized advertising (European Parliamentary Research Service 2021). Moreover, the European Union’s Digital Services Act and Digital Markets Act impose obligations on large technology companies to police illegal or harmful content and to encourage competition (European Commission 2022). These twin regulations represent a fundamental shift from ex-post enforcement to ex-ante compliance, requiring “gatekeeper” firms to fundamentally redesign their algorithms and data practices before they even reach the market.

However, it is perhaps antitrust where Europe has shown the most divergence from the US model, which as argued in Chapter 7 of the book, has been based on a CDP inspired quest to bolster innovation, despite a more recent turn, at least rhetorically and symbolically by the executive branch, as documented earlier in this section of the appendix.

### **European Antitrust**

Historically, European antitrust authorities—though roughly guided by a consumer welfare standard, like the United States—were more vigilant about preventing market dominance and less inclined to promote innovation for innovation’s sake (Geradin 2020; Schneider 2024). In contrast to the more permissive U.S. approach, Europe’s emphasis on curbing market concentration discouraged “winner-takes-most” dynamics, which may have limited the explosive platform growth seen in America. Smaller user bases and more fragmented data pools ultimately hampered cutting-edge AI model training in Europe (Veugelers 2018).

Europe has coupled classic ex-post abuse cases with an ex-ante rulebook—the Digital Markets Act (DMA)—that treats the largest platforms as “gatekeepers” subject to hard obligations on steering, self-preferencing, and access. In March 2024 the Commission opened the first non-compliance investigations against Alphabet, Apple, and Meta; by April 2025 it formally found Apple (anti-steering) and Meta in breach, and in 2025 pressed Google to revise Play Store terms on fees and external offers. The DMA sits atop a decade of headline abuse decisions (e.g., Android), where EU courts have largely upheld multibillion-euro fines—signaling that Brussels will mix ex-ante conduct rules with ex-post structural and behavioral remedies when needed (Car 2025).

*The Enforcement Crackdown (2024-2025)*

This signal was confirmed by a series of definitive enforcement actions that marked the end of the transition period. September 2024 proved to be a watershed month for the "old" antitrust regime. The European Court of Justice (ECJ) upheld a €2.4 billion fine against Google for self-preferencing its Shopping service and, in a separate landmark ruling, ordered Apple to pay €13 billion in back taxes to Ireland, overturning a lower court victory for the company. While the General Court annulled a separate €1.5 billion fine against Google's AdSense business that same month, the broader message was clear: the era of endless appeals for Big Tech was closing (Court of Justice of the European Union 2024).

By 2025, the focus shifted to the new Digital Markets Act (DMA), which proved to have immediate bite. In April 2025, the Commission issued its first non-compliance decisions, fining Apple €500 million for anti-steering violations in its App Store and Meta €200 million for its "pay or consent" data model, underscoring that DMA obligations would be backed by real sanctions (Car 2025). Perhaps the most significant victory, however, came not from a fine but from a structural concession. Under sustained EU pressure, Microsoft agreed to offer Office and Microsoft 365 suites without Teams at reduced prices and to improve interoperability and data portability; in September 2025, the Commission accepted legally binding commitments and closed its Teams bundling probe, allowing Microsoft to avoid a potentially massive fine while still forcing a global product and pricing reset (Reuters 2025). This move demonstrated the immense reach of the "Brussels Effect," as an EU regulatory action successfully forced a change in the product strategy of a U.S. tech titan.

### **The European Approach to AI**

In the mid-2020s, the European Union's strategy for AI fused a comprehensive legal framework with detailed technical standards and major infrastructure investments (European Union, 2024). Following the European Parliament's adoption of its negotiating position in June 2023, EU institutions engaged in "trilogue" negotiations, reaching a provisional agreement on December 8, 2023 (European Parliament 2023). The result was the EU Artificial Intelligence Act (EU AI Act), whose primary goals are to protect fundamental rights, foster trust and innovation, and ensure safety through defined standards.

The act establishes a comprehensive "risk-based" regime that bans certain unacceptable uses of AI, places strict pre-approval controls on "high-risk" systems used in areas like healthcare or critical infrastructure and imposes transparency duties on lower-risk applications like chatbots (European Union 2024). Specifically, the act categorizing AI applications into four tiers: unacceptable risk (banned applications like social scoring systems), high risk (systems affecting health, safety, or fundamental rights), limited risk (requiring transparency), and minimal risk (few or no requirements) (European Parliament 2023; Renda 2021). The Act introduced a certification regime for high-risk AI systems, with a tiered approach to foundation models and stricter obligations on "high-impact" models.

Key provisions included specific bans on certain AI uses—such as social scoring by public authorities and real-time biometric identification in public spaces, with limited exceptions—and imposed transparency and disclosure rules for chatbots, generative AI, and deepfakes. Non-

compliance can result in fines up to €30 million, or 6% of global annual turnover, whichever is higher (European Parliament 2023).

A new European AI Office was created within the European Commission to enforce these rules, particularly for powerful general-purpose AI models, or GPAI (European Commission 2023).<sup>203</sup> To guide companies, a voluntary Code of Practice was developed with over a thousand stakeholders, offering a path for developers to demonstrate compliance with key requirements like model documentation, security controls, and risk evaluation before formal standards are finalized (European Commission 2025).<sup>204</sup>

During the early 2020s, Europe also took the lead in addressing other regulatory issues around AI. As of 2025, the EU was drafting specific rules associated with liability. For example, by making it easier for people to sue for AI harms and updating product liability for digital products and pre-emptively clarifying responsibility (Duffourc and Gerke 2024). Its proposed AI Liability Directive eased victims' burden of proof. It also mandated disclosure requirements and safe harbor provisions tied to compliance standards (European Union 2024: Art. 50).

### *Beyond the AI Act: Liability, Copyright, and Labor*

While the AI Act established broad rules governing AI systems across all sectors of the economy, the EU simultaneously updated legal frameworks addressing specific types of harm. While the AI Act set the “horizontal rules,” the EU simultaneously updated its more “vertical” legal frameworks to address specific AI issues around liability, IP and labor markets, often diverging sharply from the U.S. approach.

In the realm of liability, Brussels moved to tackle the “black box” problem primarily through legislation rather than litigation. The revised Product Liability Directive, adopted in late 2024, explicitly brings standalone software and AI systems within the definition of a “product,” extends damage to include data loss, and broadens the range of potentially liable economic operators (European Union 2024a). To ease the burden on injured parties confronted with opaque or technically complex systems, the directive introduces rebuttable presumptions of defect and causation in situations where product malfunction and scientific complexity make it excessively difficult for claimants to prove fault, shifting much of the practical explanatory burden back onto manufacturers (ibid). By contrast, the separate AI Liability Directive—which would have added an additional presumption of causality tailored to AI under national fault-based regimes—was formally withdrawn in 2025 after lawmakers failed to reach agreement on its final shape (Brachmann 2025).

---

<sup>203</sup> The Act is reinforced by adjacent digital laws that create a broad compliance perimeter for tech companies (European Commission 2023). The Digital Services Act, for instance, requires the largest online platforms to assess and mitigate systemic risks from their algorithms, while the Cyber Resilience Act adds baseline cybersecurity requirements for all products with digital elements (ibid).

<sup>204</sup> To be sure, by 2025 several unresolved issues remained. This included how to define “high-risk” AI, addressing compliance complexity for large generative AI models, and ensuring adequate enforcement resources across EU member states (see Schneider 2024).

In copyright, the EU clung to a legislative compromise that came under increasing strain by 2025. Rather than U.S.-style “fair use,” the European framework hinges on the Text and Data Mining (TDM) exceptions in the 2019 DSM Directive, which allow AI developers to mine protected works unless rights holders opt out via machine-readable means (European Union 2019). As generative AI spread, cultural organizations and legislators warned that this opt-out regime had become a “devastating” loophole, enabling large-scale scraping by major AI firms that the original directive never anticipated (Ranking 2025). That critique was reinforced by a 2025 European Parliament study on generative AI and copyright, which concluded that the current TDM exceptions sit uneasily with large-scale commercial training practices and called for a re-examination of the framework (Lucchi and Hunter 2025). By late 2025, the conflict crystallized in a case where, after a German music rights collecting society called GEMA sued OpenAI, the Munich I Regional Court held that training and operating ChatGPT on song lyrics without a license infringed German copyright law and rejected reliance on TDM in situations where models memorize and regurgitate lyrics—marking the first European judgment to find an AI developer directly liable for training-data use and putting real pressure on the existing TDM compromise (Soppe and Schubert 2025).

Finally, Europe moved ahead of most jurisdictions in regulating the AI influenced labor market. The EU Platform Work Directive introduces a presumption of employment for platform workers and establishes information and consultation rights regarding automated monitoring and decision-making systems (European Union 2024b). Among other things, it requires platforms to provide transparency about how algorithms allocate work and evaluate performance and to ensure meaningful human oversight of significant automated decisions, curbing purely automated dismissals and other high-impact actions (Foeth 2025b). By late 2025, the European Parliament’s Employment Committee was already calling for a new directive on algorithmic management that would extend similar safeguards beyond the gig economy to the broader workforce, aiming to make platform-style transparency and human-in-the-loop requirements the default scenario across sectors (European Parliament 2025).

### **The Tension Between Regulation and Competitiveness**

After its formal adoption in 2024, the Act began to set a worldwide standard for AI regulation, not unlike the GDPR’s impact on data protection (European Union 2016). This phenomenon, often called the “Brussels Effect,” meant that multinational firms began adopting the EU’s stringent rules as their global baseline to ensure market access and avoid the cost of maintaining separate compliance systems for different regions (Bradford 2023).

Despite the legislative milestone represented by the EU AI Act, however, the EU’s regulatory ambition increasingly collided with its economic anxieties. Supporters argued it would create a level playing field for well-audited AI systems and endow Europe with a competitive advantage as global concerns over privacy, security, and ethical AI practices proliferated (Renda 2021). Critics warned of potential burdens on AI startups, particularly regarding documentation and risk assessment requirements (Schneider 2024).

Indeed, as the Act moved toward finalization, a deep rift emerged between the EU “regulators” in Brussels and the “industrialists” in key member states. Leaders like French President Emmanuel

Macron and German Chancellor Olaf Scholz actively pushed back against the Act's strictest provisions for "general-purpose" AI models (Henshall 2023). Their fear was that excessive red tape would strangle Europe's own nascent champions—such as France's Mistral AI and Germany's Aleph Alpha—before they could compete with U.S. giants (Davies 2023). This political tug-of-war mirrored the American tension between safety and innovation forged in the crucible of the AI race between the US and China. Surprisingly, by 2025 this highlighted a growing consensus across the Atlantic that strategic competitiveness might require tempering the impulse to regulate (World Economic Forum 2025).

While the final Act did include concessions to aid innovation, such as requiring every member state to create regulatory sandboxes for testing new AI, the EU recognized that legal safe harbors alone were insufficient. To truly compete, it needed to back its rules with hard power, which meant making major investment commitments to bolstering computing power and data sharing (EuroHPC Joint Undertaking 2024). Brussels was no longer content to merely regulate the AI revolution; it intended to engineer a European alternative to U.S. and Chinese dominance. By mid-2025, scattered initiatives were consolidated into a unified industrial strategy known as the "AI Continent Action Plan." Launched in April, this roadmap explicitly aimed to transform Europe into a global AI leader by integrating infrastructure investment with regulatory certainty and tying AI compute "gigafactories" and AI Factories to a new InvestAI financing facility (European Commission 2025). It was followed in October by the "Apply AI Strategy," a targeted push to fast-track AI adoption across Europe's legacy industrial base—from automotive manufacturing to healthcare—addressing the continent's persistent productivity gap (Spirlet 2025). Yet, despite these efforts, Europe's innovation comeback had "barely scratched the surface" relative to the U.S. and China (Spirlet 2025).

## POTENTIAL FUTURE TRAJECTORIES

Looking beyond the early 2020s, the dream of a harmonized global AI governance regime appears increasingly remote. Instead, the distinct regulatory philosophies of the three major powers suggest a future of regulatory bifurcation and strategic fragmentation.

In the United States, the "innovation-first" approach, tempered by a new patrimonial drift, suggests a future defined by litigation and sectoral fragmentation. Rather than a single "AI Law," we can expect a continued patchwork where liability is hammered out in the courts (via defamation and copyright suits) and specific rules are set by executive-level agencies like the FDA or SEC, creating a flexible but unpredictable environment for developers.

In China, the "State-Centric" model points toward an increasingly closed loop. As Beijing tightens its "chokehold" on data and algorithms to ensure ideological conformity and the pursuit of political-economic objectives, its AI ecosystem will likely diverge further from the West, creating a distinct "intranet of AI" optimized for state control and industrial policy rather than open commercial exchange.

In the European Union, the tension between the "regulatory instinct" and industrial anxiety will likely persist. While the AI Act aims to set a global standard via the "Brussels Effect," the rise of

"sovereign AI" policies in member states suggests that even Europe may eventually prioritize national champions over pure compliance.

In response to this complex landscape, large tech companies are likely to adapt by maintaining distinct product versions for each major region. While this approach allows for localized compliance, it also increases development costs. Similarly, these companies will also need to work closely with local insurers to develop specialized policies addressing algorithmic risk in each jurisdiction.

However, there is also the possibility that governments may take bold initiatives and attempt to solve several outstanding problems that emerged during the early 2020s. Below, I outline what challenges remain and sketch out potential regulatory solutions without predicting how likely any of these attempts are likely to arise. However, because the European and Chinese models rely on comprehensive statutes to preempt legal ambiguity, the following analysis focuses primarily on the United States and similar common-law jurisdictions, where the absence of top-down legislation leaves these complex questions to be resolved through litigation, insurance markets, and evolving property rights.

## **Legal Liability and AI**

In the US and several other countries legal liability is governed by traditional tort law doctrines like negligence and product liability—and, for several reasons, extrapolating this framework to AI is not straightforward, as extant legal frameworks were designed for clear chains of human responsibility (Smith et al. 2024). While libel and defamation laws govern harmful speech, existing product-safety frameworks are centered on physical injury and death. The overarching question is whether a software-based AI service should be treated like a product (with potential strict liability for defects) or like a publisher of speech (with different standards) or as a platform (potentially immune to liability in countries such as the US due to Section 230). Regulators face difficult decisions about whether liability should attach to the model’s developers, the data providers, or the distributors of the model, potentially requiring new approaches to liability that can account for the complex interactions between human actors and autonomous systems (see European Union 2024). In other words, if for example a medical diagnosis AI or an autonomous legal-advice chatbot makes a grave mistake, it’s unclear whether liability falls on the tool’s manufacturer, the professional using it, or the AI itself—which, as property, lacks legal personhood.

### *Opacity and the Assignment of Blame*

By 2025, courts struggled to apply negligence standards to AI because of the complex “AI supply chain” (from developers to deployers) that blurs who is at fault. As explained already, the opaque, “black box” nature of advanced AI exacerbates this challenge: existing liability frameworks are “not fully equipped” to handle AI’s lack of transparency, leaving consumers, providers, and developers unsure who is responsible for AI-caused injury.

Consider medical liability: both in the U.S. and elsewhere, this issue is particularly complex when AI systems make errors or omissions. Indeed, AI systems in healthcare can misdiagnose

patients or recommend incorrect treatments that cause appreciable harm. However, if a clinician relies on a flawed AI recommendation, determining liability is tricky. The hospital or doctor might be sued for malpractice, but they could in turn claim the AI software was defective. In product liability terms, an AI diagnostic tool might be considered a “product”; if it’s unreasonably unsafe or flawed in design, the manufacturer could be strictly liable (Smith et al. 2024).

Regulatory oversight might mirror FDA medical device approval processes, especially for AI tools that advance beyond basic transcription to diagnostic recommendations. This may potentially call on new co-liability models between healthcare providers and software developers (Adams et al. 2024). Under such a framework, the software developer would retain strict liability for technical failures—such as coding errors or data biases that deviate from the product’s specifications—while the healthcare provider would remain liable for the “human in the loop” decision of whether to accept or reject the AI’s recommendation in a specific clinical context (see Custers et al. 2025; Mello and Guha 2024).

### *Harm Associated with Speech*

While legal systems typically treat speech harms (like libel) differently than product harms, generative AI blurs this line by automating speech creation (see Smith et al. 2024). For example, defamation law requires showing that the defendant—whether a publisher, platform, or AI provider—was at least negligent in publishing the false statement (or, for public figures, acted with “actual malice,” meaning knowledge of falsity or reckless disregard for the truth). But an AI has no intent or knowledge, so plaintiffs must attribute negligence or malice to the company’s development, deployment, or monitoring of the model in scenarios where AI provides incorrect advice or information that causes reputational harm (see Frasher 2025; Goldman 2023). For example, Air Canada lost a tribunal case over misinformation from its website chatbot and was held liable for the bot’s statements—an early legal marker that heightened caution about fully automated front-line service (American Bar Association 2025).

Legal services that employ AI tools for case law research, contract drafting, and standardized legal advice pose a similar challenge in terms of how to handle liability for AI-generated legal errors. Courts and bar associations have emphasized that lawyers remain professionally responsible for vetting AI output, and reliance on erroneous AI-generated content without adequate review could expose practitioners to malpractice liability (see Menaldo 2025).

Practically speaking, companies like OpenAI prominently warn users about ChatGPT’s accuracy limitations and have argued that, together with their efforts to reduce erroneous output, these warnings demonstrate they did not act negligently—a position that received early support in Walters (2025), a state level case in Georgia where the court dismissed a defamation claim in part due to OpenAI’s safeguards and disclaimers. However, this remains a relatively untested legal area, with only a handful of defamation suits against AI companies filed as of late 2025, and courts are just beginning to address how traditional defamation standards apply when there is no human author (Scarcella 2025).

In the future, potential regulatory frameworks might include laying out liability rules for new categories of speech. Returning to the legal services example outlined above, this may include “certified legal AI” with attorney oversight, or safe harbor provisions for routine tasks, building on emerging bar guidance that lawyers must supervise and remain responsible for any AI they use in practice (Huffstetler 2024). The preservation of attorney-client privilege in cloud-based AI services presents another crucial challenge, which may require robust security measures and clear guidelines for data handling, given the risk that sharing confidential material with third-party AI tools could waive privilege (Eichen and Kumarasamy 2025).

### *Physical Harm*

The emergence of autonomous systems and AI has complicated traditional approaches to tort law and the assignment of blame. Self-driving cars, drones, and robots raise complex questions about strict liability, negligence, and product liability. When an autonomous vehicle causes harm, the legal system must determine whether responsibility lies with the manufacturer, the software developer, the hardware supplier, or the owner.

Self-driving cars exemplify AI acting in the physical world, where malfunctions can cause injury or death. Under existing law, vehicle accidents have well-established liability regimes. Many experts believe product liability law (holding manufacturers responsible for defects) will be the main mechanism for AV-related harm (Anderson et al 2016; Villasenor 2017). In practice, this could mean if an autonomous car’s AI causes a crash, the carmaker (and its software suppliers) would face lawsuits instead of the human occupant. This creates strong incentives for companies to make their AI as safe as possible (ibid). Indeed, early AV legislation efforts in the U.S. avoided broad immunity for manufacturers, on the assumption that tort law can handle most scenarios

While novel questions have arisen (e.g. if aftermarket self-driving software is installed, who is liable?), legal scholars suggest the “very strong set of incentives” under decades of product liability precedent generally apply to AI drivers as well (Villasenor 2014). Indeed, early driverless rideshare pilots in U.S. cities such as San Francisco and Phoenix proceeded under permit-and-insurance regimes that allocate responsibilities among manufacturers/permit holders and set reporting obligations—an incremental pathway that allows scaling while preserving recourse for harms (see California Public Utilities Commission 2020).

Further regulation may create an even more seamless, streamlined process. This is because, unlike generative AI which produces infinite and unpredictable outputs, an autonomous vehicle is a discrete physical object with clear performance metrics. To reduce uncertainty about liability in the event of a crash, regulation could focus on rigorous safety testing and certification processes for autonomous vehicles, for a wide range of conditions to ensure the safe operation of autonomous vehicles before deployment—analogous in spirit to regimes for drugs, food, or aircraft (see Anderson et al. 2016; Gianclaudio and Pasquale 2024).

### *Insuring Against AI Centric Risks*

As industries increasingly adopt A.I. at larger scales, the insurance sector faces new challenges in developing appropriate risk allocation mechanisms. In response, it may provide specialized policies covering like algorithmic malpractice and new types of data breaches. This may call on insurers to demand that companies seeking insurance coverage to hedge against these hazards to prove that they adhere to recognized A.I. safety or bias-mitigation standards.<sup>205</sup> This may entail that third parties administer rigorous audit standards. In this way, the private insurance market may effectively take on the role of “privatized regulator,” refusing to underwrite AI systems that do not meet specific benchmarks for safety and robustness and thereby enforcing discipline if the government fails to fill the vacuum (see OECD 2021).

Even if this happens, however, there is still some room for the government to play a role. If the insurance industry develops ways to address the new AI related risks outlined above, governments may mandate minimum liability insurance for some AI applications to militate against adverse selection and broaden the coverage pool—akin to how it compels automobile insurance (OECD 2019). Similarly, because insurance coverage raises concerns about moral hazards in A.I. development and deployment, as it may encourage developers and users to take more risks or distort use in ways that create negative externalities (see Shavell 1987), governments may seek to proactively encourage responsible development practices and provide robust safety measures as a complement for insurance coverage. This may include mandating mandatory third-party algorithmic audits, rigorous pre-deployment "red-teaming" to stress-test models, and strict adherence to technical standards like the NIST Risk Management Framework (NIST 2023).

## **IP Issues Around AI**

As AI chatbots began displacing traditional search engines and publishers witnessed declining referral traffic and ad revenues, the media industry began to question the open web paradigm that had been pivotal to the training of foundational AI models I outlined in the previous section of this appendix.<sup>206</sup> Around 2015, platforms started providing direct answers without requiring

---

<sup>205</sup> Testing methodologies and evaluation metrics are crucial components in assessing AI system performance and reliability. These methodologies need to examine how the AI performs across diverse and representative data sets to ensure effectiveness for all intended populations. A systematic approach following structured and repeatable processes ensures consistent and objective evaluation while reducing the risk of ad-hoc or biased testing. Testing tries to simulate real-world conditions as closely as possible, considering factors like noisy data, unexpected inputs, and potential adversarial attacks. Practitioners try to promote transparency by documenting all the tests they conduct and their results to allow for review, replication, and continued improvement over time. The evaluation metrics used to measure AI system performance and characteristics usually meet several key criteria: relevant to the specific goals and intended use of the AI system; quantifiable, providing numerical values that can be objectively compared and analyzed, enabling progress tracking and identification of areas needing improvement; statistically sound, based on appropriate methods that allow for meaningful comparisons while avoiding misleading conclusions. On all these points see NIST 2023.

<sup>206</sup> This paragraph draws on Simonetti and McMillan (2025).

users to click through to the original source, through features like Google’s “featured snippets” (short excerpts from web pages shown at the top of search results) and “AMP cached articles” (faster-loading versions of web pages stored directly on Google’s servers). This trend was later extended by “generative AI interfaces” like ChatGPT, Bard/Gemini, and Copilot, which provided direct summaries of information, further reducing “referral clicks” (clicks back to the original website). As website traffic and advertising revenues weakened, publishers’ business models—heavily reliant on predictable pageviews—came under growing strain.

Publishers pursued multiple strategies to reclaim control over their content. Major outlets including The Atlantic, Dotdash Meredith, and News Corp titles negotiated licensing deals with AI companies like OpenAI, establishing a “pay-or-play” dynamic (David 2024). Additionally, publishers implemented technical barriers by using services to throttle or block unwanted automated data collectors; Cloudflare, for example, shifted defaults and tooling so that new sites block AI crawlers by default, require “bot verification,” and let publishers state whether access is for training, inference, or search (Prince 2025). On July 1, 2025, Cloudflare made this default change and reported that over one million customers had already enabled its one-click AI-block, while rolling out programs (e.g., “Pay Per Crawl”) and authentication enhancements to move toward permission-based access (ibid). Even nonprofits took defensive action: iFixit publicly blocked an Anthropic crawler after it hit the site about 1 million times in a single day (Koebler 2024), and Wikipedia urged AI companies to stop scraping and use its paid API with clearer access controls (Perez 2025).

Predictably, as attested to by multiple high-profile lawsuits, the legal landscape was another arena where acrimony between publishers and web scrapers scavenging for AI model training data intensified.<sup>207</sup> For example, while Reddit sued Anthropic for alleged continued scraping despite promises to stop, The New York Times pursued copyright claims against Microsoft and OpenAI, and News Corp subsidiaries filed suit against Perplexity. In turn, these escalating conflicts raised concerns about a partitioned web that could shut out smaller publishers and legitimate researchers while threatening the economic incentives for creating original content that feed both AI and journalism ecosystems.

In turn, these escalating conflicts raised concerns about a partitioned web that could shut out smaller publishers and legitimate researchers while threatening the economic incentives for creating original content that feeds both AI and journalism ecosystems. Perhaps recorded music embodied this new dynamic best. In 2025, settlements between startups like Suno and Udio and major record labels signaled a shift away from the “open” training model; under a November 2025 deal, for example, Suno agreed to introduce licensed AI models that will replace its current versions in 2026, with Udio similarly resolving disputes with Warner Music and Universal Music, raising the bar for entrants in high-value creative media (Singh 2025).

Crucially, indemnity obligations increasingly shift liability from developers to downstream users (Madigan 2025). Unlike many traditional software licensing arrangements, where developers often indemnify users against intellectual property claims, major AI companies have in some cases structured their service agreements to narrow or reallocate copyright liability—requiring

---

<sup>207</sup> This paragraph draws on Simonetti and McMillan (2025).

users, rather than model developers, to bear much of the legal responsibility for any infringement in AI-generated outputs (Rastogi 2025). This evolving liability architecture can turn enterprise AI deployments into potential litigation risks, necessitating new categories of spending on legal reserves, specialized insurance products, and compliance systems.

Considering this increased balkanization, it is conceivable that different industries and policymakers could establish standardized data-licensing frameworks—analogue to collective-rights agencies in music—which aggregate royalties for AI training on third-party content (Lemley and Casey 2021). New statutory provisions or patent-office guidelines might also define how ‘AI-generated’ inventions are claimed; for instance, by requiring disclosure of significant training data sources or characteristics and, more speculatively, by exploring mechanisms to recognize the interests of data owners whose material was integral to the invention (WIPO 2024a).

### *IP Rights to AI Generated Content*

While India, Ireland, New Zealand, South Africa, the United Kingdom, and Ukraine extend copyright protection to “computer-generated works,” in the US and several other countries fully AI-generated assets effectively enter the public domain upon creation (WIPO 2024b). Repeated rulings by the U.S. Copyright Office and federal courts, most notably the D.C. Circuit’s decision in *Thaler* (2025) firmly established that works generated solely by AI systems cannot, by themselves, be copyrighted (see Carroll et al. 2025). This baseline insecurity about IP may make it less likely that organizations will fully invest in creative processes that involve AI.

This situation may call on the American government—and possibly others—to create a new, *sui generis* forms of intellectual property protection specifically for AI-generated assets. It could draw on distinct rights introduced for database protection or semiconductor mask works and grant AI users a limited exclusivity period to monetize their creations without requiring the “human authorship” standard of traditional copyright (Senftleben and Buijtelaar 2020). Alternatively, it could look more like patent protection for software, where eligibility hinges not on expressive creativity but on the novelty and utility of the underlying technical solution, effectively protecting the functional outcome of the AI process rather than just its artistic form (WIPO 2024a).

Such an AI centric regime would likely shift the focus from the final artistic output to the specific technical process used to generate it. Rather than individual examiners vetting millions, if not billions, of individual images or texts—which is logistically impossible—the system could allow developers to register their unique “prompt chains,” fine-tuning workflows, and parameter settings as a protected “functional recipe.” Creators would securely lodge their workflows in an encrypted registry, with protection attaching automatically. Infringement would then be defined not by aesthetic similarity, but by technical misappropriation—meaning a competitor would be liable only if they demonstrably reverse-engineered or copied the specific algorithmic steps and prompt sequences to replicate the work, thereby protecting investments in the method behind creation without locking up the general style or idea in the public domain (Senftleben and Buijtelaar 2020; WIPO 2024a).

## Standards

As discussed in Chapter 10 of the book, technical interoperability standards help A.I. systems integrate safely across industries and borders. These cover everything from data formats to algorithmic auditing protocols to provenance standards like C2PA (Coalition for Content Provenance and Authenticity), which evaluates the origin and history of digital media (C2PA 2023). By embedding cryptographically secure metadata into files, these standards allow different AI systems to verify the origin of a piece of media—confirming whether it was created by a human, a specific AI model, or a camera—thereby enabling a trusted ecosystem of interoperable applications.

While voluntary standard-setting efforts by organizations like the IEEE and ISO may guide best practices and foster a common language for A.I. developers (see IEEE 2021), going forward there might also be a need for governments to step in more assertively. They may have a critical role to play in coordinating standard-setting efforts, ensuring that standards are developed through open and inclusive processes, and potentially mandating compliance with certain key standards to ensure safety and interoperability (see Abbott 2012; Harris 2025).

Work on this front had already begun in 2025. In the US, government agencies like NIST, alongside international SSOs like ISO and IEEE, were actively developing crucial testing methodologies and benchmarks. For example, the U.S. AI Safety Institute, established under NIST, created specialized taskforces to collaborate on research and testing of advanced AI models for national security and public safety. Similarly, in 2025, the European Union's Horizon Europe program began funding initiatives like ELLIOT, a major R&D project focused on creating open, trustworthy, general-purpose AI models that can robustly generalize across various data types.

Beyond interoperability, a new class of standards focused on "model evaluations" (or "evals") was quickly becoming a central pillar of AI governance. Because AI capabilities evolve faster than legislation, regulators were increasingly relying on standardized technical benchmarks—akin to the "SafetyPerf" suite discussed earlier, or specific "red-teaming" protocols to look for chemical-weapon knowledge or cyber-offensive capabilities—to define the threshold of acceptable risk (Narajala and Narayan 2025). Establishing these common "yardsticks" for intelligence and safety is arguably an urgent regulatory challenge: they may ultimately determine which models are deemed safe for release and which trigger strict oversight (NIST 2023).

As the previous section of this appendix explained, many AI systems, particularly those based on deep learning, operate as "black boxes," making it difficult to understand how they reach their conclusions (Doshi-Velez and Kim 2017). While this lack of transparency raises concerns about accountability, trust, and fairness, there are technical challenges to improving it: ham-fisted attempts to increase transparency may compromise the performance of AI systems. Moreover, companies that develop closed-source proprietary AI algorithms often want at least some of the ways in which outputs are generated to remain secret for purposes of IP protection (see Rudin 2019).

The push for explainable AI (XAI) has sought to make AI decision-making processes more transparent and understandable. During the early 2020s, regulators explored ways to require developers to provide explanations for AI outputs, particularly in high-stakes domains like healthcare and finance. Since revealing the raw code or mathematical weights of a deep learning model offers little insight to a layperson, regulators increasingly demanded "counterfactual explanations." Instead of explaining how the model thought, these systems explain why a decision was made by showing what would have changed the outcome (e.g., "You were denied the loan because of your debt-to-income ratio; if that ratio had been 5% lower, you would have been approved") (Wachter et al. 2017). This focuses on actionable intelligibility for the user rather than technical transparency for the engineer.

Looking ahead, governments may move beyond these individual explanations toward rigorous, standardized algorithmic auditing regimes. Rather than relying on company self-assessments, future regulations may empower independent oversight bodies—modeled after financial auditors—to inspect the internal logic and training data of high-risk models under strict confidentiality (NIST 2023). This "qualified transparency" would solve the trade secret dilemma by granting regulators "white box" access to verify safety and fairness, while providing the public with "black box" summaries, ensuring accountability without exposing proprietary code (Malgeri and Pascale 2024).

Finally, standards may be required to correct market failures in the economic pricing of AI. At least as late as 2025, the industry suffered from a "pricing fog" where inconsistent metering—by tokens, time, or compute type—created high transaction costs and prevented businesses from easily switching providers (see OpenAI 2025h). Future regulation may therefore need to define a standardized, abstract unit of "AI work" (analogous to the kilowatt-hour) to create a fungible metric for consumption. This would enable businesses to budget reliably and create a transparent, liquid market for compute where providers compete on clear cost-per-unit metrics.

## WHERE AI GOES FROM HERE

While quantum and evolutionary methods offer powerful optimization tools, some researchers argue that the most comprehensive blueprint for intelligent systems comes from biology itself. A growing body of work proposes moving beyond brain-inspired metaphors toward treating neuroscience findings as literal design specifications (Clark 2013; Shipp 2024; Aston-Jones and Cohen 2005; Yu and Dayan 2005; Foster and Wilson 2006; Mattar and Daw 2018; Miller 2025).

The argument proceeds as follows. Current AI systems, despite their impressive performance on narrow benchmarks, lack capabilities that biological brains achieve routinely: learning reliably from just a few examples, maintaining performance when conditions change, and planning multiple steps ahead while consuming minimal power (Miller 2025). Proponents of the neuro-inspired approach contend that these limitations stem from architectural choices that diverge fundamentally from how primate cognition actually works. They point to specific biological mechanisms—hierarchical prediction systems that enable rapid error-correction, sparse neural circuits that support working memory, chemical signaling that manages confidence and uncertainty, and memory consolidation processes that enable skill transfer—as templates for a different kind of AI (Halassa and Kastner 2017; Marcus 2025c; Marcus 2025d).

These proposals remain largely at the conceptual stage rather than demonstrated at scale. But their advocates envision three interconnected research directions. The first involves translating biological circuits into computational modules that could be integrated with existing systems—essentially allowing current AI models to delegate high-stakes decisions to more reliable, biologically-grounded algorithms. The second emphasizes training protocols that test for genuine understanding through interactive environments with verifiable outcomes, using evolutionary algorithms to discover effective architectures rather than hand-coding them (Back et al. 2018; Mattar and Daw 2018). The third prioritizes reliability, pairing every reasoning chain with verification mechanisms such as symbolic proofs and constraint solvers (Marcus 2025c).

Complementing these architectural shifts is an urgent focus on energy efficiency. The human brain performs remarkable cognitive work on roughly twenty watts of power—a fraction of what current AI systems consume. Researchers have proposed that energy-to-solution metrics should become as important as accuracy benchmarks, with hardware designed around brain-inspired computing patterns: event-driven processing, compute-in-memory architectures, and near-sensor inference to reduce data movement (Merolla et al. 2014; Sebastian et al. 2020; Rasch et al. 2024). Community benchmarks like MLPerf Power have begun measuring energy consumption as a primary outcome and could expand to include reasoning and planning tasks (Tschand et al. 2024).

Along these lines, CMOS (Complementary Metal–Oxide–Semiconductor)—the primary technology used today to build integrated circuits—may give way to entirely new computing paradigms that are much more energy efficient. For example, optical computing, which uses light, or photons, instead of electrons to process information, potentially offers much faster speeds and lower power consumption, and neuromorphic computing, which designs chips to mimic the structure and function of the human brain, aiming for more efficient AI processing (Buntz 2024). One can also anticipate further refinements in process node technology (the ongoing effort to shrink transistor sizes on chips) into the late 2020s: even smaller features (e.g., sub-3nm nodes) should increase transistor density and improve performance and power efficiency (Chang 2024).<sup>208</sup>

### *A Call to Action*

---

<sup>208</sup> This will require continued advancements in EUV lithography and the development of new lithography techniques (Intel Newsroom 2024). Advanced packaging technologies should also play a critical role, including wider adoption of 3D stacking and chiplet architectures (breaking down a large, complex chip into smaller, specialized "chiplets" that are then combined in a single package, allowing for more flexibility and higher yields), allowing for the integration of more complex systems combining different types of components, like logic, memory, and sensors, often from different manufacturing processes, within a single package (ibid). This will likely involve further development of high-bandwidth interconnects (extremely fast connections that allow large amounts of data to move quickly between different parts of a chip or between chiplets) and advancements in thermal management, techniques and systems for dissipating the increased heat generated by these densely packed chips (McKinsey & Company 2024a).

Some advocates for a new AI and computing paradigm have called for a coordinated national effort—a "Phase Two" translation program—to bring these elements together. Such an initiative would expand neuroscience research aimed at identifying reusable cognitive patterns, maintain open libraries of standardized modules, and establish cross-laboratory benchmarks measuring not just capability but also robustness and energy efficiency (Miller 2025; Marcus 2025c; Marcus 2025d).

For the political economist, these proposals raise familiar questions about industrial organization and public investment. If neuro-inspired approaches prove viable, they could shift competitive advantage toward actors with access to neuroscience research infrastructure—potentially universities and government laboratories rather than the hyperscale cloud providers that currently dominate. The emphasis on energy efficiency, meanwhile, could alter the geographic distribution of AI development by reducing dependence on massive power supplies. Furthermore, a reduced reliance on massive datasets would erode the "data moats" currently guarded by tech giants, potentially democratizing capabilities that today require the entire internet's worth of text to train. Whether these possibilities materialize depends on technical progress that lies beyond this book's scope to evaluate. What can be said is that the trajectory of AI development remains contested, with significant implications for the market structures and power relations examined in this book.

## CONCLUSION

This section of the appendix has charted the rise of a new "AI Cold War," a geostrategic competition that has coincided with the end of the era of hyper-globalization after the end of the actual Cold War. Instead, in the new age of technological mercantilism, three major economic blocs diverged, each forging its own distinct industrial policy to try to win this race—or in the case of Europe, at least try to keep from falling further behind.

The United States adopted a hybrid model, blending its market-driven "Creative Destruction Paradigm" with a new, state-directed approach of subsidies, export controls, and "America First" dealmaking. Circa 2025, the US blended market mechanisms, robust capital markets, and strong IP protection to drive rapid commercial innovation with a more state directed industrial strategy that called for significant investments in strategic sectors like AI and semiconductors.

China doubled down on its state-centric model, leveraging massive subsidies, "national champions" like Huawei, and a successful open-source "swarms beat the titan" strategy to build a self-reliant AI stack. China wielded powerful tools of central planning, resource mobilization, and a massive domestic market to achieve rapid technological advancement in targeted areas. While this state-directed approach yielded impressive gains, it faced potential long-term sustainability challenges stemming from structural economic issues, demographic pressures, and increasingly centralized governance under Xi Jinping. In a departure from overt state centralization, it also strengthened its IP regime to facilitate the acquisition of technology from the global technological frontier and foster nascent indigenous innovation.

The European Union, meanwhile, pursued a "regulatory-first" path, attempting to build "digital sovereignty" through programs like EuroHPC while seeking to export its values globally via the

"Brussels Effect"—a topic I will take up again in the following section . In doing so, it mirrored the US's increased use of export controls, investment screening, and sanctions. It also increased public investment in high-tech fields such as semiconductors.

While I did not look outside of these big three players in this section of the appendix, I also note that, in this fragmenting landscape, smaller, but technologically advanced, economies like South Korea, the Netherlands, and Israel also emerged as crucial players; they wielded outsized influence in critical niches such as semiconductor equipment and photonics. Their innovation trajectories and AI policies, despite individual quirks and challenges, carried global significance. For example, the entire global AI industry—including chipmakers in the U.S., South Korea, and Taiwan—is completely dependent on a single Dutch firm, ASML, which holds a 100% monopoly on the extreme ultraviolet (EUV) lithography machines required to manufacture all advanced AI chips.

Moreover, other countries mimicked the industrial policies pursued by the three big polities that were the focus of this section of the appendix—especially around semiconductors. South Korea's K-Chips Act provided tax credits for domestic semiconductor investments (Park 2024). Taiwan implemented its own "Chips Act," offering tax credits for R&D costs to encourage domestic innovation and maintain technological leadership (Kao 2025). Japan allocated substantial funds to support its semiconductor industry, including generous subsidies for TSMC's Kumamoto plant (Kelly 2024). It also launched the Rapidus initiative, a venture focused on developing cutting-edge chips in collaboration with firms like IBM (Kelly and Lee 2022).

During the late 2020s, we can expect deeper political involvement in steering technological progress towards specific goals around AI and adjacent innovations such as quantum computing. We can also anticipate potential economic conflicts over critical resources—from rare earth minerals to semiconductors, data, and intellectual property. While this heightened geopolitical competition may spur advanced R&D reminiscent of the Sputnik effect, it also risks resource misallocation as states fiercely compete to pick winners in the technological lottery.

Yet, the focus on a "winner-take-all" race for a hypothetical AGI may be misguided. The more immediate and consequential contest is what Alhassani and Bak (2025) call the "AI implementation" race: the practical challenge of embedding today's AI into every factory, hospital, school, and military unit to unlock real-world productivity gains. This perspective is consistent with the view put forth in this book about industrial revolutions. For the Third Industrial Revolution, centered on personal computers and the internet, thorough and economic transformation deployment depended on the CDP.

Indeed, for all the Herculean efforts that US firms and the federal government have poured into winning the AI race with China, it may very well be the case that their focus has been wrongheaded. It has been all about going big or going home, if you will. But what if China's "swarms beat the titan" strategy is better tailored to what it takes to consolidate an industrial revolution? Returning to Alhassani and Bak (2025): even if America maintains an 18-month lead in raw model performance by restricting chips and other inputs, foreign competitors can quickly catch up at far lower cost thanks to the steep returns on modest hardware investments and continual efficiency gains in both algorithms and semiconductor design. Therefore, perhaps the

“AI implementation” race, integrating the best available models into every sector of the economy and public service, is what all countries may eventually settle on, especially if the promise of AGI is left unkept.

This section of the appendix also analyzed the divergent AI regulation approaches of the US, China, and the EU in the early 2020s. The US blended elements of the CDP with permissive AI regulation. China pursued a state-directed, application-focused strategy with early AI-specific regulations. The EU prioritized consumer protection and comprehensive regulations, including the groundbreaking AI Act. Examining key regulatory challenges (liability, standards, privacy, transparency), this section of the appendix reveals how each polity’s approach reflected its institutional heritage and strategic priorities. It also anticipates how future AI governance trajectories may evolve in a fragmented policy landscape.

## BIBLIOGRAPHY

- 3GPP. 2019. "Third Generation Partnership Project: Working Procedures." 3GPP Organizational Partners. [https://www.3gpp.org/ftp/information/working\\_procedures/3gpp\\_wp.htm](https://www.3gpp.org/ftp/information/working_procedures/3gpp_wp.htm)
- 3GPP. 2021b. Technical Specification Group working methods (TR 21.900). Version 17.0.0. Valbonne, France: 3GPP Organizational Partners.  
<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=555>
- 3GPP. 2024. "About 3GPP." <https://www.3gpp.org/about-us>
- A&P 1949. *United States v. New York Great A&P Tea Co.*, Civ. No. 52-139. US District Court, S.D. New York (filed Sept. 15). <https://law.justia.com/cases/federal/district-courts/FSupp/67/626/2311543/>
- AAR. 2025. Staggers Rail Act of 1980. AAR. <https://www.aar.org/issue/staggers-act-of-1980/>
- Aaronson, Susan Ariel, and Patrick Leblond. 2018. "Another Digital Divide: The Rise of Data Realms and Its Implications for the WTO." *Journal of International Economic Law* 21(2): 245-272.
- ABA 1969. Commission to Study the Federal Trade Commission. 1969. Report of the ABA Commission to Study the Federal Trade Commission. <https://www.uschamber.com/assets/documents/Primer-on-The-FTC's-New-Section-5-Guidance.pdf>
- ABA. 1956. "Report and Recommendations Presented to House of Delegates by Section of Antitrust Law at Midwinter Meeting, Chicago, Illinois, February 20–21, 1956." <https://www.justice.gov/atr/merger-guidelines/overview>
- Abbate, Janet. 1999. *Inventing the Internet*. Cambridge, MA: The MIT Press.
- Abelson, Hal 1994. Summary of In re Alappat. Prepared by counsel for Tektronix, Inc. Hosted at MIT CSAIL. <https://groups.csail.mit.edu/mac/classes/6.805/assorted-short-pieces/alappat-summary.html>
- Abernathy, Frederick, John Dunlop, Janice Hammond, and David Weil. 1999. *A Stitch in Time*. New York: Oxford University Press.
- Accredited Standards Committee X9 (ASC X9). 2022. *Financial Services Data Security Standards*. ASC X9.

Acemoglu, Daron, and David Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." In *Handbook of Labor Economics*, Volume 4. Edited by Orley Ashenfelter and David Card. *Handbook of Labor Economics*. Amsterdam, Holland: Elsevier, 1043–1171.

Acemoglu, Daron. 2002. Directed Technical Change. *Review of Economic Studies* 69(4): 781–809.

Acemoglu, Daron. 2009. *Introduction to Modern Economic Growth*. Princeton: Princeton University Press.

Acemoglu, Daron. 2024. "The Simple Macroeconomics of AI." MIT Economics Working Paper / Economic Policy. <https://economics.mit.edu/sites/default/files/2024-04/The%20Simple%20Macroeconomics%20of%20AI.pdf>

ACT. 2021. Digital Markets Act Position Paper. The App Association. <https://actonline.org/wp-content/uploads/ACT-The-App-Association-DMA-Position-Paper-March-.pdf>

Action. 2025. Overcoming Manufacturing's Biggest Data Challenges: From Silos to Trust. White Paper. <https://www.action.com/wp-content/uploads/2025/02/overcoming-manufacturings-biggest-data-challenges-from-silos-to-trust.pdf>

Addyston Pipe. 1898. *United States v. Addyston Pipe and Steel Co.*, No. 498. US Court of Appeals, Sixth Circuit, February 8. <https://law.justia.com/cases/federal/appellate-courts/F2/85/271>

Adeboye, Daniel. 2025. "How much does an NVIDIA H100 GPU cost?" Northflank Blog. <https://northflank.com/blog/how-much-does-an-nvidia-h100-gpu-cost>

Adelman, Morris Albert 1959. *A&P: A Study in Price-Cost Behavior and Public Policy*. Cambridge, MA: Harvard University Press.

ADQ. 2025. ADQ and Energy Capital Partners to establish a USD 25 billion US-based investment partnership. ADQ. <https://www.adq.ae/newsroom/adq-and-energy-capital-partners-to-establish-a-usd-25-billion-us-based-investment-partnership-focused-on-developing-new-power-generation-to-serve-the-growing-electricity-needs-of-data-centers/>

Affinius Capital. 2025. "From Cloud to Capital." *PREA Quarterly Global* (Spring 2025), Exhibit 2.

Affinius Capital. [https://affiniuscapital.com/app/uploads/Spring\\_2025\\_PREA\\_Quarterly\\_Global.pdf](https://affiniuscapital.com/app/uploads/Spring_2025_PREA_Quarterly_Global.pdf)

Aggarwal, Sanjay, and Betsy Mulé. 2025. State of Robotics. F-Prime Capital.  
[https://docs.google.com/presentation/d/1jMNIjt3rU2lQ\\_WNIKSy2UUQkE\\_Dr9FJhMecAnjCQnw8/edit?slide=id.p#slide=id.p](https://docs.google.com/presentation/d/1jMNIjt3rU2lQ_WNIKSy2UUQkE_Dr9FJhMecAnjCQnw8/edit?slide=id.p#slide=id.p)

Aghion, Philippe, and Peter Howitt. 1992. "A Model of Growth Through Creative Destruction." *Econometrica* 60(2): 323-351.

Aghion, Philippe, and Peter Howitt. 1998. "On the Macroeconomic Effects of Major Technological Change." In *General Purpose Technologies and Economic Growth*, edited by Elhanan Helpman. Cambridge, MA: MIT Press, 121–144.

Aghion, Philippe, Ufuk Akcigit, and Peter Howitt. 2014. What Do We Learn from Schumpeterian Growth Theory? In *Handbook of Economic Growth*, Vol. 2, edited by Philippe Aghion and Steven N. Durlauf. Amsterdam, Holland: Elsevier, 515–563.

AGNCSAL. 1955. Report of the Attorney General's National Committee to Study the Antitrust Laws. Washington, DC: U.S. Government Printing Office.

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb. 2023. Similarities and Differences in the Adoption of General Purpose Technologies. NBER Working Paper No. 30976. Cambridge, MA: National Bureau of Economic Research.

Ahmed, Nur and Muntasir Wahed. 2020. "The De-democratization of AI: Deep Learning and the Compute Divide in Artificial Intelligence Research." arXiv preprint arXiv:2010.15581.  
<https://arxiv.org/abs/2010.15581>

AInvest. 2025. "Intel's Layoffs: A Necessary Reset or a Last Gasp?" AInvest.  
<https://www.ainvest.com/news/intel-layoffs-reset-gasp-2506/>

Akhtar, Shayerah Ilias. 2013. Trade reorganization: Overview and issues for Congress. CRS Report R42555. <https://www.everycrsreport.com/reports/R42555.html>

Alappat. 1994. In re Alappat. No. 92-1381. US Federal Circuit Court, July 29.  
<https://law.justia.com/cases/federal/appellate-courts/F3/33/1526/513542/>

Albergotti, Reed 2024. Amazon announces new "Rainier" AI compute cluster with Anthropic. Semafor. <https://www.semafor.com/article/12/03/2024/amazon-announces-new-rainier-ai-compute-cluster-with-anthropic>

Albertus, Michael, and Victor Menaldo. 2018. *Authoritarianism and the Elite Origins of Democracy*. Cambridge: Cambridge University Press.

Alcatel-Lucent, Ericsson, NEC, NextG Networks, Nokia, Nokia Siemens Networks, Nortel, and Sony Ericsson. 2008. *Joint Proposal on 4G Mobile Broadband: Standard Essential Patents*. ETSI.

Alcoa. 1945. *United States v. Aluminum Co. of America*, No. 144. US Court of Appeals, Second Circuit, March 12. <https://law.justia.com/cases/federal/appellate-courts/F2/148/416/1503668/>

Alhassani, Mehdi, and Anthony Bak. 2025. America Is Winning the Wrong AI Race. *The Wall Street Journal*. <https://www.wsj.com/articles/america-is-winning-the-wrong-ai-race->

Alice Corp. 2014. *Alice Corp. v. CLS Bank International*. No. 13-298. US Supreme Court, June 19. <https://supreme.justia.com/cases/federal/us/573/208>

Alicke, Knut, Tacy Foster, Katharina Hauck, and Vera Trautwein. 2023. "Tech and regionalization bolster supply chains, but complacency looms." *McKinsey & Company*.

Aljunaid, Saif Khalifa, Almheiri, Saif Jasim, Dawood, Hussain, and Khan, Muhammad Adnan. (2025). Secure and transparent banking: Explainable AI-driven federated learning model for financial fraud detection. *Journal of Risk and Financial Management* 18(4): 179.

Allen, Robert. 2009. *The British Industrial Revolution in Global Perspective*. Cambridge: Cambridge University Press.

Allied Orthopedic Appliances. 2010. *Allied Orthopedic Appliances Inc. v. Tyco Health Care Group LP*, No. 08-56314, 9th Cir., Jan. 6. <https://law.justia.com/cases/federal/appellate-courts/F3/592/991>

Allison, John, and Emerson Tiller. 2003. "The Business Method Patent Myth." *Berkeley Technology Law Journal* 18: 987-1084.

Alok, Yashashwy 2024. Top 10 Open-Source AI Libraries for Developers. <https://www.analyticsvidhya.com/blog/2024/12/open-source-ai-libraries>

Alphabet Inc. 2020–2025. Annual Reports and SEC Filings. SEC.

Alphabet Inc. 2025. “2024 Q4 Earnings Call” (transcript). Alphabet Investor Relations.  
<https://abc.xyz/investor/events/event-details/2025/2024-Q4-Earnings-Call/>

Alsop, Thomas. 2025. Qualcomm R&D expenditures worldwide 2016–2024. Statista.  
<https://www.statista.com/statistics/980315/research-and-development-expenditures-of-qualcomm/>

Alston, Eric, Lee Alston, Bernardo Mueller, and Tomas Nonnenmacher. 2018. Institutional and Organizational Analysis. Cambridge: Cambridge University Press.

Altman, Sam. 2025. “Reflections.” January 5, 2025. <https://blog.samaltman.com/reflections>

Alvarez, Robert. 2025. "How to Accelerate Training for Large AI Models with Pure Storage FlashBlade and PyTorch Checkpointing." Purely Technical (blog), Pure Storage. <https://blog.purestorage.com/purely-technical/pure-storage-flashblade-and-pytorch-asynchronous-checkpointing/>

Amazon Inc. 2024. "AWS Trainium2 Instances Now Generally Available." US Press Center, December 03. <https://press.aboutamazon.com/2024/12/aws-trainium2-instances-now-generally-available>

Amazon Inc. 2025. “Amazon.com (AMZN) Q4 2024 Earnings Call Transcript.” The Motley Fool.  
<https://www.fool.com/earnings/call-transcripts/2025/02/06/amazoncom-amzn-q4-2024-earnings-call-transcript/>

Amazon Staff. 2024. Amazon and Anthropic deepen their shared commitment to advancing generative AI. Amazon News. <https://www.aboutamazon.com/news/company-news/amazon-anthropic-ai-investment>

AMD Inc. 2025. Current Report (Form 8-K), “Entry into a Material Definitive Agreement” (Warrant to OpenAI OpCo, LLC). SEC. <https://ir.amd.com/financial-information/sec-filings/content/0001193125-25-230895/0001193125-25-230895.pdf>

American Can. 1949. United States v. American Can Co., No. 26345-H. US District Court, N.D. California, S.D., November 10. <https://law.justia.com/cases/federal/district-courts/FSupp/87/18/>

American Express. 2018. Ohio v. American Express Co., 585 U.S. \_\_\_\_ (138 S. Ct. 2274).  
<https://supreme.justia.com/cases/federal/us/585/16-1454>

AIPLA. 2019. Report of the Economic Survey. AIPLA. <https://www.aipla.org/detail/journal-issue/2019-report-of-the-economic-survey>

American Nuclear Society 2025. Talen and Amazon expand their partnership for Pennsylvania nuclear. American Nuclear Society Newswire. <https://www.ans.org/news/2025-06-16/article-7113/talen-and-amazon-expand-their-partnership-for-pennsylvania/>

Aminabadi, Reza Yazdani, Samyam Rajbhandari, Minjia Zhang, Ammar Ahmad Awan, Cheng Li, Du Li, Elton Zheng, Jeff Rasley, Shaden Smith, Olatunji Ruwase, and Yuxiong He. 2022. DeepSpeed Inference: Enabling Efficient Inference of Transformer Models at Unprecedented Scale. arXiv preprint arXiv: 2207.00032. <https://arxiv.org/abs/2207.00032>

Anderson. 2024. Anderson v. TikTok, Inc. No. 22-3061. US Court of Appeals, Third Circuit, August 27. <https://law.justia.com/cases/federal/appellate-courts/ca3/22-3061/22-3061-2024-08-27.html>

Andreessen, Marc. 2011. "Why Software Is Eating the World." The Wall Street Journal. <https://www.wsj.com/articles/SB10001424053111903480904576512250915629460>

Anthropic. 2024. Introducing the Model Context Protocol. Anthropic. <https://www.anthropic.com/news/model-context-protocol>

Anthropic. 2025a. Piloting Claude for Chrome. Anthropic. [https://www.anthropic.com/news/claude-for-chrome?\\_bhlid=bd9fae9ffe489cd5db411d4be7bf4c89adf41d9c](https://www.anthropic.com/news/claude-for-chrome?_bhlid=bd9fae9ffe489cd5db411d4be7bf4c89adf41d9c)

Anthropic. 2025c. "Expanding our use of Google Cloud TPUs and Services." Anthropic News. <https://www.anthropic.com/news/expanding-our-use-of-google-cloud-tpus-and-services>

Anthropic. 2025d. "Anthropic invests \$50 billion in American AI infrastructure." Anthropic News. <https://www.anthropic.com/news/anthropic-invests-50-billion-in-american-ai-infrastructure>

Anthropic. 2025e. Anthropic Economic Index: Insights from Claude 3.7 Sonnet. <https://www.anthropic.com/news/anthropic-economic-index-insights-from-claude-sonnet-3-7>

Appen. 2021. Overcoming AI Deployment Challenges/5 Tips for Handling AI Obstacles. Appen. <https://www.appen.com/blog/overcoming-ai-deployment-challenges>

Apple Computer. 1994. Apple Computer, Inc. v. Microsoft Corp. Nos. 93-16867, 93-16869 and 93-16883.

US Court of Appeals, Ninth Circuit, September 19. <https://law.justia.com/cases/federal/appellate-courts/F3/35/1435/605245/>

Apple Developer. n.d.a. “Get Started — iOS.” Apple Developer Documentation.

<https://developer.apple.com/ios/get-started>

Apple Developer. n.d.b. "Core Location." Apple Developer Documentation.

<https://developer.apple.com/documentation/corelocation/>

Apple Developer. n.d.c. “Configuring your app to use location services.” Apple Developer

Documentation. <https://developer.apple.com/documentation/corelocation/configuring-your-app-to-use-location-services>

Apple Developer. n.d.d. “In-app purchase.” Apple Developer Documentation.

<https://developer.apple.com/design/human-interface-guidelines/in-app-purchase>

Apple Developer. n.d.e. “Handling location updates in the background.” Apple Developer

Documentation. <https://developers.apple.com/documentation/corelocation/handling-location-updates-in-the-background>

Apple Developer. n.d.f. “Getting the current location of a device.” Apple Developer Documentation.

<https://developer.apple.com/documentation/corelocation/getting-the-current-location-of-a-device>

Apple Inc. 2019b. Apple Inc. v. Robert Pepper, et al., No. 17-204, U.S. Supreme Court, May 13.

[https://scholar.google.com/scholar\\_case?case=15157888566008153651&q=apple+v+pepper&hl=en&as\\_sdt=6,48&as\\_vis=1](https://scholar.google.com/scholar_case?case=15157888566008153651&q=apple+v+pepper&hl=en&as_sdt=6,48&as_vis=1)

Apple Inc. 2020–2025. Annual Reports and SEC Filings. SEC.

Apple Inc. 2024a. Annual Report 2024 (Form 10-K). Securities and Exchange Commission.

<https://www.sec.gov/Archives/edgar/data/320193/000032019324000123/aapl-20240928.htm>

Apple Inc. 2023b. “Apple unveils iPhone 15 Pro and iPhone 15 Pro Max.” Press Release. Apple

Newsroom. <https://www.apple.com/newsroom/2023/09/apple-unveils-iphone-15-pro-and-iphone-15-pro-max/>

Apple Inc. 2024b. "iPhone 16 – Technical Specifications." <https://www.apple.com/iphone-16/specs/>

Apple Inc. 2024c. "Introducing Apple Intelligence for iPhone, iPad, and Mac." Apple Newsroom. <https://www.apple.com/newsroom/2024/06/introducing-apple-intelligence-for-iphone-ipad-and-mac/>

Apple Support. n.d. "Use Continuity to work across Apple devices." Apple Support. <https://support.apple.com/guide/iphone/intro-to-continuity-iphf5fa30b66/ios>

Apple. 2021a. "A Day in the Life of Your Data." [https://www.apple.com/privacy/docs/A\\_Day\\_in\\_the\\_Life\\_of\\_Your\\_Data.pdf](https://www.apple.com/privacy/docs/A_Day_in_the_Life_of_Your_Data.pdf)

Apple. 2021b. "App Tracking Transparency." Apple Developer Documentation. <https://developer.apple.com/documentation/apptrackingtransparency>

Apple. 2021c. "Data Privacy Day at Apple." Apple Newsroom. <https://www.apple.com/newsroom/2021/01/data-privacy-day-at-apple-improving-transparency-and-empowering-users>

Apple. 2022. "Deploying Transformers on the Apple Neural Engine." Apple Machine Learning Research. <https://machinelearning.apple.com/research/neural-engine-transformers>

Apple. 2024a. Apple Platform Security. [https://help.apple.com/pdf/security/en\\_US/apple-platform-security-guide.pdf](https://help.apple.com/pdf/security/en_US/apple-platform-security-guide.pdf)

Apple. 2024b. "Introducing Apple's On-Device and Server Foundation Models." Apple Machine Learning Research. <https://machinelearning.apple.com/research/introducing-apple-foundation-models>

Apple. 2024c. "Private Cloud Compute: A new frontier for AI privacy in the cloud." Apple Security Research. <https://security.apple.com/blog/private-cloud-compute>

Apple. 2024d. "Apple Intelligence Foundation Language Models." arXiv preprint: arXiv 2407.21075. <https://arxiv.org/pdf/2407.21075>

Apple. 2025. "Apple Intelligence Foundation Language Models: Tech Report 2025." arXiv preprint: arXiv 2507.13575. <https://arxiv.org/pdf/2507.13575>

Applegate, Lynda, and Kevin Davis. 1995. "Xerox: Outsourcing Global Information Technology Resources." Harvard Business School Case 9-195-158, Harvard Business School Publishing. Available at: <https://hbsp.harvard.edu/product/195158-PDF-ENG>

Apptronik. 2024. Apptronik partners with Google DeepMind robotics to accelerate advancement on AI-powered humanoid robots. Apptronik Press Release. <https://apptronik.com/news-collection/apptronik-partners-with-google-deepmind-robotics>

Areeda, Phillip, and Herbert Hovenkamp. 1978. *Antitrust Law: An Analysis of Antitrust Principles and Their Application*. New York: Aspen Publishers.

Areeda, Phillip, and Herbert Hovenkamp. 2003. *Antitrust Law: An Analysis of Antitrust Principles and Their Application* (2nd ed.). New York: Aspen Publishers.

Areeda, Phillip, and Herbert Hovenkamp. 2023. *Antitrust Law: An Analysis of Antitrust Principles and Their Application* (5th ed.). New York: Wolters Kluwer.

Areeda, Phillip, and Donald Turner. 1975. "Predatory Pricing and Related Practices Under Section 2 of the Sherman Act." *Harvard Law Review* 88(4): 697–733.

Areeda, Phillip, and Donald Turner. 1978. *Antitrust Law*. Boston: Little, Brown.

Areeda, Phillip, and Herbert Hovenkamp. 2020. *Antitrust Law*. 5th ed. New York: Wolters Kluwer.

Areeda, Phillip. 1989. "Essential Facilities." *Antitrust Law Journal* 58(3): 841–53.

Argenton, Cédric, and Jens Prüfer. 2012. "Search Engine Competition with Network Externalities." *Journal of Competition Law and Economics* 8(1): 73-105.

Aridor, Guy, Yeon-Koo Che, Brett Hollenbeck, Maximilian Kaiser, and Daniel McCarthy. 2025. "Evaluating the Impact of Privacy Regulation on E-Commerce Firms." Working paper. [https://www.anderson.ucla.edu/sites/default/files/document/2025-05/att\\_privacy.pdf](https://www.anderson.ucla.edu/sites/default/files/document/2025-05/att_privacy.pdf)

Arista Networks. 2024. Arista Unveils Etherlink AI Networking Platforms. Arista Networks Press Release. <https://investors.arista.com/Communications/Press-Releases-and-Events/Press-Release-Detail/2024/Arista-Unveils-Etherlink-AI-Networking-Platforms/default.aspx>

Arm Holdings plc. 2023. Registration Statement on Form F-1. U.S. Securities and Exchange Commission. <https://investors.arm.com/node/6921/html>

Armbrust, Michael, Ali Ghodsi, Reynold Xin, and Matei Zaharia. 2021. "Lakehouse: A New Generation of Open Platforms that Unify Data Warehousing and Advanced Analytics." Proceedings of CIDR 2021.

Armstrong, Mark. 2006. "Competition in Two-Sided Markets." *The RAND Journal of Economics* 37(3): 668–91.

Arnold, Thurman Wesley 1941. *The Bottlenecks of Business*. New York: Reynal and Hitchcock.

Arora, Ashish, Andrea Fosfuri, and Alfonso Gambardella. 2001. *Markets for Technology*. Cambridge, MA: MIT Press.

Arrow, Kenneth. 1962. "Economic Welfare and the Allocation of Resources for Invention." In *The Rate and Direction of Inventive Activity: Economic and Social Factors*, edited by Richard Nelson. Princeton, NJ: Princeton University Press, 609–25.

Arroyo, Carmen, and Edward Ludlow. 2025. "xAI to Raise \$20 Billion After Nvidia and Others Boost Round." Bloomberg. <https://www.bloomberg.com/news/articles/2025-10-07/musk-s-xai-nears-20-billion-capital-raise-tied-to-nvidia-chips?embedded-checkout=true>

ASME. 2025. "Special Issue: Human–Robot Collaboration in Industry 5.0." <https://asmejcise.org/2025/04/07/special-issue-human-robot-collaboration-in-industry-5-0/>

ASML. 2024. *ASML Annual Report 2024*. ASML Holding N.V. Investor Relations. <https://www.asml.com/en/investors/annual-report/2024>

ASML. 2025. "EUV lithography." ASML Products – EUV lithography systems. <https://www.asml.com/en/products/euv-lithography-systems>

ASPDAC. 2025. "Chiplet Package Design: Complexity, Reliability, and Advanced Packaging." <https://www.aspdac.com/aspdac2025/archive/pdf/7F-3.pdf>

ASTC. 2025. "New ASTC Survey Shows Gaps Between Public Support for Science and Understanding of Federal Actions." ASTC News. <https://www.astc.org/astc-news-announcements/new-astc-survey->

shows-gaps-between-public-support-for-science-and-understanding-of-how-science-is-impacted-by-federal-actions/

AT&T 2018. *United States v. AT&T Inc.*, 310 F. Supp. 3d 161 (D.D.C.).

<https://law.justia.com/cases/federal/district-courts/district-of-columbia/dcdce/1:2017cv02511/182300/172>

AT&T 2019. *United States v. AT&T Inc.*, No. 18-5214, 916 F.3d 1029 (D.C. Cir.).

<https://law.justia.com/cases/federal/appellate-courts/cadc/18-5214/18-5214-2019-02-26.html>

AT&T 1999. *AT&T Corp. v. Excel Communications, Inc., Excel Communications Marketing, Inc., and Excel Telecommunications, Inc.* No. 98-1338. US Court of Appeals, Federal Circuit, April 14.

<https://law.justia.com/cases/federal/appellate-courts/F3/172/1352/599511/>

Atkinson, Robert, Daniel Castro, and Stephen Ezell. 2008. "The Digital Road to Recovery: A Stimulus Plan to Create Jobs, Boost Productivity and Revitalize America." Information Technology and Innovation Foundation. [https://d1bcsfjk95uj19.cloudfront.net/files/Digital\\_Road\\_to\\_Recovery.pdf](https://d1bcsfjk95uj19.cloudfront.net/files/Digital_Road_to_Recovery.pdf)

Atkinson, Robert. 2025. "No, AI Robots Won't Take All Our Jobs: Instead, they will boost productivity, lower prices and spur the evolution of the labor market." *The Wall Street Journal*.

<https://www.wsj.com/opinion/no-ai-robots-wont-take-all-our-jobs-employment-productivity-innovation-41068792>

ATOM Project Substack. 2025. One Small Step for ATOM, one giant leap for OpenAI.

<https://atomproject.substack.com/p/one-small-step-for-atom-one-giant>

Austin, Steven. 2025. How many people work at Anthropic (Claude)? Comprehensive workforce analysis for 2025. *Marketingscoop*. <https://www.marketingscoop.com/website/seo/how-many-people-work-at-anthropic-claude-comprehensive-workforce-analysis-for-2025/>

Authers, John. 2026. Greenland Isn't as Big as it Looks in Risk. *Points of Return*. Bloomberg.

<https://www.bloomberg.com/opinion/newsletters/2026-01-20/greenland-isn-t-as-big-as-it-looks-in-risk?srnd=phx-opinion>

Autor, David, David Dorn, Lawrence Katz, Christina Patterson, and John Van Reenen. 2020. "The Fall of the Labor Share and the Rise of Superstar Firms." *Quarterly Journal of Economics* 135(2): 645–709.

Autor, David, and Michael Handel. 2013. "Putting Tasks to the Test: Human Capital, Job Tasks, and Wages." *Journal of Labor Economics* 31(1): S59–S96.

Autor, David. 2015. "Why Are There Still So Many Jobs? The History and Future of Workplace Automation." *Journal of Economic Perspectives* 29(3): 3–30.

Autor, David. 2024. "Applying AI to Rebuild Middle Class Jobs." NBER Working Paper No. 32140. Cambridge, MA: National Bureau of Economic Research.

AWS. 2024a. High Performance Computing on AWS. AWS. <https://aws.amazon.com/hpc>

AWS. 2024b. Generative AI inference architecture and best practices on AWS (AWS inference stack). Available at: <https://docs.aws.amazon.com/prescriptive-guidance/latest/gen-ai-inference-architecture-and-best-practices-on-aws/aws-inference-stack.html>

AWS. 2024c. "Anthropic's Claude 3 Opus model now available on Amazon Bedrock." AWS What's New. <https://aws.amazon.com/about-aws/whats-new/2024/04/anthropics-claude-3-opus-amazon-bedrock>

AWS. 2025. Amazon Bedrock Documentation. AWS. <https://docs.aws.amazon.com/bedrock>

Ayodele, Abiola. 2025. OSAT Semiconductor Services: The Backbone of Outsourced Chip Assembly and Testing. Wevolver. <https://www.wevolver.com/article/osat-semiconductor-services-the-backbone-of-outsourced-chip-assembly-testing>

Azure AI Studio. 2025. Endpoints for inference in production. Azure AI Studio. <https://learn.microsoft.com/en-us/azure/machine-learning/concept-endpoints>

Baack, Stefan. 2024. "A Critical Analysis of the Largest Source for Generative AI Training Data: Common Crawl." Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency (FAccT '24). Association for Computing Machinery. <https://dl.acm.org/doi/10.1145/3630106.3659033>

Backus, Matthew, Christopher Conlon, and Michael Sinkinson. 2021. "Common Ownership in America: 1980–2017." *American Economic Journal: Microeconomics* 13(3): 273–308.

Baek, Jinheon, Nirupama Chandrasekaran, Silviu Cucerzan, and Allen Herring. 2024. Knowledge-Augmented Large Language Models for Personalized Contextual Query Suggestion. Proceedings of the ACM Web Conference (WWW '24). <https://arxiv.org/pdf/2311.06318>

Bahar, Dany. 2018. "The Middle Productivity Trap: Dynamics of Productivity Dispersion." *Economics Letters* 167: 60–66.

Bailey, Martha, John DiNardo, and Bryan Stuart. 2021. "The Economic Impact of a High National Minimum Wage: Evidence from the 1966 Fair Labor Standards Act." *Journal of Labor Economics*, Vol. 39(S2): S329-S367.

Bain, Joe. 1956. *Barriers to New Competition*. Cambridge, MA: Harvard University Press.

Baker, Jonathan. 1997. The Problem with Baker Hughes and Syufy: On the Role of Entry in Merger Analysis. *Antitrust Law Journal* 65(2): 353-374.

Baker, Jonathan and Carl Shapiro. 2010. Reinvigorating Horizontal Merger Enforcement. In Robert Pitofsky (Ed.), *How the Chicago School Overshot the Mark: The Effect of Conservative Economic Analysis on U.S. Antitrust* New York: Oxford University Press, 235-291.

Baker, Jonathan. 2007. "Market Definition." *Antitrust Law Journal* 74(1): 129–173.

Bakos, Yannis. 1997. "Reducing Buyer Search Costs: Implications for Electronic Marketplaces." *Management Science* 43(12): 1676–1692.

Baldwin, Richard. 2016. *The Great Convergence*. Cambridge, MA: Harvard University Press.

Ball Memorial Hospital. 1986. *Ball Memorial Hospital, Inc. v. Mutual Hospital Insurance, Inc.*, No. 85-1481. US Court of Appeals, Seventh Circuit, March 4. <https://law.justia.com/cases/federal/appellate-courts/F2/784/1322>

Balp, Laura. 2025. Storage Next: Why AI Workloads Need a Kitchen Built for Speed. ScaleFlux Blog. <https://scaleflux.com/storage/storage-next-serving-data-like-a-world-class-restaurant>

Balser, Jimmy, Minon Schwartz and Jon Shimabukuro. 2025. "A New Civil Service 'Policy/Career' Schedule." CRS Report LSB11262. <https://www.congress.gov/crs-product/LSB11262>

Bandyopadhyay, Subhayu, Maximiliano Dvorkin, Fernando Leibovici, and Ana Maria Santacreu. 2022. The Shifting Tides of Global Trade. Federal Reserve Bank of St. Louis. <https://www.stlouisfed.org/annual-report/2022/shifting-tides-global-trade>.

Banga, Rashmi. 2013. Measuring value in global value chains. UNCTAD. [https://unctad.org/system/files/official-document/ecidc2013misc1\\_bp8.pdf](https://unctad.org/system/files/official-document/ecidc2013misc1_bp8.pdf)

Bantock, Asher 2025. The SignalFire State of Tech Talent Report – 2025. SignalFire. <https://www.signalfire.com/blog/signalfire-state-of-talent-report-2025>

Barath Harithas, Kyle Meng, Evan Brown, and Catharine Mouradian. 2025. The Liberation Day Tariffs. Washington, DC: Center for Strategic and International Studies. <https://www.csis.org/analysis/liberation-day-tariffs-explained>

Barenholtz, Lior. 2025. ChatGPT monthly visits: SEO analysis and competitor comparisons. Similarweb. <https://www.similarweb.com/blog/marketing/geo/chatgpt-monthly-visits/>

Barkai, Simcha. 2020. “Declining Labor and Capital Shares.” *Journal of Finance* 75 (5): 2421–2463.

Barnes and Noble. 2014. *Nguyen v. Barnes and Noble Inc.*, No. 12-56628, U.S. Court of Appeals, Ninth Circuit, August 16. [https://scholar.google.com/scholar\\_case?case=11003811139217543321](https://scholar.google.com/scholar_case?case=11003811139217543321)

Barnes. 2009. *Barnes v. Yahoo!, Inc.* No. 06-56662. US Court of Appeals, Ninth Circuit, June 22. <https://law.justia.com/cases/federal/appellate-courts/F3/570/1096>

Barnett, Jonathan. 2020. *Innovators, Firms, and Markets: The Organizational Logic of Intellectual Property*. New York: Oxford University Press.

Barnett, Jonathan. 2017. "Has the Academy Led Patent Law Astray?" *Berkeley Technology Law Journal* 32(4): 1313–1386.

Barron, David. 2010. Best Practices for OLC Legal Advice and Written Opinions. Memorandum for Attorneys of the Office. DoJ, Office of Legal Counsel. <https://www.justice.gov/olc/page/file/1511836/dl>

Baroudy, Kim, Jonatan Janmark, Tobias Strålin, Abhi Satyavarapu, and Zeno Ziemke. 2020. “Europe’s start-up ecosystem: Heating up, but still facing challenges.” McKinsey & Company.

<https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/europes-start-up-ecosystem-heating-up-but-still-facing-challenges>

Barroso, Luiz André, Jeffrey Dean, and Urs Holzle. 2003. Web search for a planet: The Google cluster architecture. *IEEE micro*, 23(2): 22-28.

Barry Wright. 1983. *Barry Wright Corp. v. ITT Grinnell Corp.*, No. 83-1292. US Court of Appeals, First Circuit, December 29. <https://law.justia.com/cases/federal/appellate-courts/F2/724/227/265312/>

Bartz et al. 2025. *Bartz et al. v. Anthropic PBC*, No. 3:24-cv-05417. US District Court for the Northern District of California, August 19, 2024. <https://docs.justia.com/cases/federal/district-courts/california/candce/3:2024cv05417/434709/231>

Battelle, John. 2005. *The Search*. New York: Portfolio.

Baumol, William. 1982. "Contestable Markets: An Uprising in the Theory of Industry Structure." *The American Economic Review* 72(1): 1–15.

Harrison, Joseph, Matthew Josefy, Matias Kalm, and Ryan Krause. 2023. "Using supervised machine learning to scale human-coded data: A method and dataset in the board leadership context." *Strategic Management Journal* 44(7): 1780-1802.

Bayard, Thomas, and Kimberly Ann Elliott. 1994. *Reciprocity and Retaliation in U.S. Trade Policy*. Washington, DC: Institute for International Economics.

BEA. 2024. Trade in Services in 2023. BEA. <https://apps.bea.gov/scb/issues/2024/10-october/1024-international-services.htm>

BEA. 2025. U.S. International Trade in Goods and Services, December and Annual 2024. <https://www.bea.gov/news/2025/us-international-trade-goods-and-services-december-and-annual-2024>

Beauchene, Vinciane, Sylvain Duranton, Nipun Kalra, and David Martin. 2025. "AI at Work 2025: Momentum Builds, but Gaps Remain." Boston Consulting Group, Slideshow. <https://web-assets.bcg.com/fd/0d/bcc5dfae4cbaa08c718b95b16cf5/ai-at-work-2025-slideshow-june-2025-edit-02.pdf>

Bell, Emily, and Taylor Owen. 2017. "The Platform Press." Tow Center for Digital Journalism, Columbia Journalism School. [https://www.cjr.org/tow\\_center\\_reports/platform-press-how-silicon-valley-reengineered-journalism.php](https://www.cjr.org/tow_center_reports/platform-press-how-silicon-valley-reengineered-journalism.php)

Bendavid, Naftali, and Abha Bhattarai. 2025. "Trump's Risky Role as the Would-Be CEO of America." The Washington Post. <https://www.washingtonpost.com/politics/2025/08/12/trump-corporations-ceo-economy-tariffs/>

Benesch, Friedlander, Coplan and Aronoff LLP. 2008. "Qualcomm Highlights the Need for In-House and Outside Counsel to Work Closely to Conduct E-Discovery." E-Discovery Bulletin. [https://www.beneschlaw.com/a/web/687/IP\\_eDoc\\_Qualcomm.pdf](https://www.beneschlaw.com/a/web/687/IP_eDoc_Qualcomm.pdf)

Bengio, Yoshua, Sasha Luccioni, Edward Raff, Abhishek Gupta, Daniel McDuff, David Rolnick, Emma Strubell, Irene Solaiman, Margaret Mitchell, and Timnit Gebru. 2023. "Collective Governance of Open Foundation Models." arXiv preprint: arXiv.2307.08774. <https://doi.org/10.48550/arXiv.2307.08774>

Benjamin, Dan, and Matthews, Kent. 1992. U.S. and U.K. Unemployment Between the Wars: A Doleful Story. Institute for Economic Affairs, London.

Benston, George. 1990. The Separation of Commercial and Investment Banking. New York: Oxford University Press.

Beraja, Martin, Andrew Kao, David Yang, and Noam Yuchtman. 2023. "AI-tocracy." Quarterly Journal of Economics 138(3): 1349–1402.

Bergen, Mark, and Jennifer Surane. 2017. "Google Will Stop Reading Your Emails for Gmail Ads." Bloomberg. <https://www.bloomberg.com/news/articles/2017-06-23/google-will-stop-reading-your-emails-for-gmail-ads>

Berkey Photo. 1979. Berkey Photo, Inc. v. Eastman Kodak Co., Nos. 1019, 1070. US Court of Appeals, Second Circuit, June 25. <https://law.justia.com/cases/federal/appellate-courts/F2/603/263/105215>

Bessen, James, and Michael Meurer. 2008. Patent Failure: How Judges, Bureaucrats, and Lawyers Put Innovators at Risk. Princeton: Princeton University Press.

Bessen, James, and Robert Hunt. 2007. "An Empirical Look at Software Patents." *Journal of Economics and Management Strategy* 16(1): 157–189.

Bessen, James, and Stefano Calligaris. 2018. "Proprietary IT Intensity and the Persistence of Superstar Returns." NBER Working Paper No. 24259. Cambridge, MA: National Bureau of Economic Research.

Bessen, James. 2015. *Learning by Doing: The Real Connection between Innovation, Wages, and Wealth*. New Haven, CT: Yale University Press.

Bessen, James. 2020. Industry Concentration and Information Technology. *Journal of Law and Economics* 63(3): 531–570.

Bick, Alexander, Adam Blandin, and David Deming. 2025. The Rapid Adoption of Generative AI, Federal Reserve Bank of St. Louis Working Paper 2024-027. <https://doi.org/10.20955/wp.2024.027>

BigCommerce. (n.d.). What is a Facebook Custom Audience and how can they grow online stores?

BigCommerce. <https://www.bigcommerce.com/glossary/facebook-customer-audience/>

Biswas, Prassana. 2025. CUDA vs SYCL: A Comparison of GPU Programming Models. LinkedIn.

<https://www.linkedin.com/pulse/cuda-vs-sycl-comparison-gpu-programming-models-prasanna-biswas-pdwsc>

BitLaw. 2025. "History of Software Patents, from Benson, Flook, and Diehr to Bilski and Alice." Beck and Tysver. <https://www.bitlaw.com/software-patent/history.html>.

Bjorhovde, Runar. 2024. "How much money do smartphone vendors actually earn per device?" LinkedIn.

Blackstone. (2024, May 17). CoreWeave Secures \$7.5 Billion Debt Financing Facility led by Blackstone and Magnetar. <https://www.blackstone.com/news/press/coreweave-secures-7-5-billion-debt-financing-facility-led-by-blackstone-and-magnetar/>

Blevins, Emily, Alice Grossman, and Karen Sutter. 2023. CHIPS Act of 2022. CRS Report R47523.

<https://www.congress.gov/crs-product/R47523>

Blind, Knut. 2004. *The Economics of Standards: Theory, Evidence, Policy*. Cheltenham, England: Edward Elgar Publishing.

Block, Fred 2008. *Swimming Against the Current: The Rise of a Hidden Developmental State in the United States*. *Politics and Society*, 36(2): 169–206.

Block, Fred, and Matthew Keller. 2011. *State of Innovation: The U.S. Government’s Role in Technology Development*. New York: Routledge.

Block. 2025. "Block, Anthropic, and OpenAI Launch the Agentic AI Foundation." *Block Inside*.  
<https://block.xyz/inside/block-anthropic-and-openai-launch-the-agentic-ai-foundation>

Bloom, Benjamin 1984. “The 2 Sigma Problem: The Search for Methods of Group Instruction as Effective as One-to-One Tutoring.” *Educational Researcher* 13(6): 4–16.

Bloomberg News. 2025. “Klarna Turns From AI to Real-Person Customer Service.” Bloomberg.  
<https://www.bloomberg.com/news/articles/2025-05-08/klarna-turns-from-ai-to-real-person-customer-service>

BLS. 2011. *Handbook of Methods*. DoL. <https://www.bls.gov/opub/hom/pdf/inp-20110509.pdf>

BLS. 2025a. “Year-over-Year Inflation Rate: CPI-U (All Items), Annual Average, 2002–2024.” DoL.  
[https://www.bls.gov/pir/spm/spm\\_chart5\\_2024data.htm](https://www.bls.gov/pir/spm/spm_chart5_2024data.htm)

BLS 2025b. *Productivity and Costs: Fourth Quarter and Annual Averages 2024, Preliminary*. DoL.  
[https://www.bls.gov/news.release/archives/prod2\\_02062025.htm](https://www.bls.gov/news.release/archives/prod2_02062025.htm)

BLS. 2025c. *Productivity and Costs: Second Quarter 2025, Revised*. DoL.  
<https://www.bls.gov/news.release/prod2.nr0.htm>

BLS. 2025d. “Office of Productivity and Technology: Concepts.” *BLS Handbook of Methods*. DoL.  
<https://www.bls.gov/opub/hom/opt/concepts.htm>

BLS. 2025e. “Office of Productivity and Technology: Data Sources.” *BLS Handbook of Methods*. DoL..  
<https://www.bls.gov/opub/hom/opt/data.htm>

BLS. n.d. *Consumer Price Index (CPI)*. DoL. <https://www.bls.gov/cpi/>

Bohn, Robert, Christopher Greer, and Jason Kahn. 2023. *Report of the Advanced Communications Technologies Working Group (NIST Interagency Report 8483)*. NIST.  
[https://tsapps.nist.gov/publication/get\\_pdf.cfm](https://tsapps.nist.gov/publication/get_pdf.cfm)

Bonacich, Edna, and Khaleelah Hardie. 2006. "Wal-Mart and the Logistics Revolution." In *Wal-Mart: The Face of Twenty-First-Century Capitalism*, edited by Nelson Lichtenstein, New York: The New Press, 163–187.

Bonatesta, Lara. 2025. "Forty-seven states have enacted deepfake legislation since 2019." *Ballotpedia News*. <https://news.ballotpedia.org/2025/07/22/forty-seven-states-have-enacted-deepfake-legislation-since-2019/>

Bonnet, Alexandre. 2024. *ONNX Standardized Format: The Universal Translator for AI Models*. Encord. <https://encord.com/blog/onnx-open-neural-network-exchange-format/>

Borin, Alessandro, and Michele Mancini. 2021. "Supply Chain Fragmentation and the Global Trade Elasticity: A New Accounting Framework." *IMF Economic Review* 69 (4): 656–680.

Bork, Robert. 1978. *The Antitrust Paradox*. New York: Basic Books.

Borkar, Shekhar, Saurabh Dighe, and Nirav Kahn. 2024. "Accelerating Industry-Wide Innovations in Datacenter Infrastructure and Security." *Microsoft Azure Blog*. <https://azure.microsoft.com/en-us/blog/accelerating-industry-wide-innovations-in-datacenter-infrastructure-and-security/>

Bort, Julie. 2025a. "Google Launched Its Deepest AI Research Agent Yet—on the Same Day OpenAI Dropped GPT-5.2." *TechCrunch*. <https://techcrunch.com/2025/12/11/google-launched-its-deepest-ai-research-agent-yet-on-the-same-day-openai-dropped-gpt-5-2/>

Bort, Julie. 2025b. "In another chess move with Microsoft, OpenAI is pouring \$12B into CoreWeave." <https://techcrunch.com/2025/03/10/in-another-chess-move-with-microsoft-openai-is-pouring-12b-into-coreweave/>

Boudreau, Kevin J., and Hagiu, Andrei 2009. Platform rules: Multi-sided platforms as regulators." *Platforms, markets and innovation* 1(0): 163-191.

Bousquette, Isabelle, Belle Lin, Tom Loftus, and Steven Rosenbush. 2025. "What Are Companies Actually Doing With AI? Our Reporters Talk It Out." *Wall Street Journal*. <https://www.wsj.com/articles/what-are-companies-actually-doing-with-ai-our-reporters-talk-it-out-a12dd305>

Bowman, Ward. 1957. Tying Arrangements and the Leverage Problem. *Yale Law Journal*, 67(1): 19-36.

Boyle, Alan, Taylor Soper, and Todd Bishop. 2020. "Apple acquires Xnor.ai, edge AI spin-out from Paul Allen's AI2, for price in \$200M range." *GeekWire*. <https://www.geekwire.com/2019/ai2-china-us-research/>

Bradford, Anu. 2020. *The Brussels Effect*. New York: Oxford University Press.

Bradford, Anu. 2023. *Digital Empires*. Oxford: Oxford University Press.

Brantle, Thomas, and M. Hosein Fallah. 2007. "Complex Innovation Networks, Patent Citations and Power Laws." In *PICMET '07: Portland International Conference on Management of Engineering and Technology*, August 5–9, Portland, OR. IEEE. <https://doi.org/10.1109/PICMET.2007.4349390>

Bratton, Laura. 2025. "Big Tech's AI Investments Set to Spike to \$364 Billion in 2025." *Yahoo Finance*. <https://finance.yahoo.com/news/big-techs-ai-investments-set-to-spike-to-364-billion-in-2025-as-bubble-fears-ease-143203885.html>

Braun, Alan. 2025. Scaling responsibly: evolving our API rate limits to power the next generation of Atlassian Cloud. *Atlassian Blog*. <https://www.atlassian.com/blog/platform/evolving-api-rate-limits>

Bresnahan, Timothy, and Manuel Trajtenberg. 1995. "General Purpose Technologies: 'Engines of Growth'?" *Journal of Econometrics* 65(1): 83–108.

Bresnahan, Timothy, and Shane Greenstein. 1996. Technical Progress and Co-Invention in Computing and in the Uses of Computers. *Brookings Papers on Economic Activity: Microeconomics*: 1–77. [https://www.brookings.edu/wp-content/uploads/1996/01/1996\\_bpeamicro\\_bresnahan.pdf](https://www.brookings.edu/wp-content/uploads/1996/01/1996_bpeamicro_bresnahan.pdf)

Bresnahan, Timothy, Erik Brynjolfsson, and Lorin Hitt. 2002. "Information Technology, Workplace Organization, and the Demand for Skilled Labor." *The Quarterly Journal of Economics* 117(1): 339-376.

Bresnahan, Timothy. 2010. General Purpose Technologies. In *Handbook of the Economics of Innovation*, Vol. 2, edited by Bronwyn Hall and Nathan Rosenberg. Amsterdam, Holland: Elsevier, 761-791.

Breyer, Stephen, Richard Stewart, Cass Sunstein, Adrian Vermeule, and Michael Herz. 2022. *Administrative Law and Regulatory Policy: Problems, Text, and Cases*. 9th ed. Burlington, MA: Aspen Publishing.

Bria, Francesca, Paul Timmers and Fausto Gernone. 2025. EuroStack – A European Alternative for Digital Sovereignty. Bertelsmann Stiftung. <https://www.bertelsmann-stiftung.de/en/our-projects/reframetech-algorithmen-fuers-gemeinwohl/project-news/eurostack-a-european-alternative-for-digital-sovereignty>

Brock, William, and Steven Durlauf. 2001. "Discrete Choice with Social Interactions." *Review of Economic Studies* 68(2): 235-260.

Brooke Group Ltd. 1993. *Brooke Group Ltd. v. Brown and Williamson Tobacco Corp.*, No. 92-466. U.S. Supreme Court, June 21. <https://supreme.justia.com/cases/federal/us/509/209>

Brougher, Jim. 2025. "Reddit's New AI Licensing Deal Shows How Content Co.s Get Paid to Train LLMs." *Media + the Machine (Substack)*, September 29, 2025. <https://mediaandthemachine.substack.com/p/reddits-new-ai-licensing-deal-shows>.

Brown Shoe. 1962. *Brown Shoe Co. v. United States*, No. 4. US Supreme Court, June 25. <https://supreme.justia.com/cases/federal/us/370/294/>

Brown, A. B. 2020. Record 2.7M industrial robots are in operation, but demand has slowed. *Supply Chain Dive*.

Brynjolfsson, Erik. 1993. The Productivity Paradox of Information Technology. *Communications of the ACM* 36(12): 66-77.

Brynjolfsson, Erik, Avinash Collis, W. Erwin Diewert, Felix Eggers, and Kevin Fox. 2019. *GDP-B: Accounting for the value of new and free goods in the digital economy*. NBER Working Paper No. 25695. Cambridge, MA: National Bureau of Economic Research.

Brynjolfsson, Erik, and Lorin Hitt. 2000. "Beyond Computation." *Journal of Economic Perspectives* 14(4): 23–48.

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen. 2025. *Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence*. Stanford Digital Economy Lab Working Paper. <https://digitaleconomy.stanford.edu/publications/canaries-in-the-coal-mine/>

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2021. The Productivity J-Curve. *American Economic Journal: Macroeconomics* 13(1): 333–372.

Buchanan, James, and Gordon Tullock. 1962. *The Calculus of Consent: Logical Foundations of Constitutional Democracy*. Ann Arbor: University of Michigan Press.

Buiten, Miriam. 2019. "Towards intelligent regulation of artificial intelligence." *European Journal of Risk Regulation* 10(1): 41-59.

BEA. 2021. "Gross Domestic Product by Industry." DoC. <https://www.bea.gov/news/2022/gross-domestic-product-fourth-quarter-and-year-2021-second-estimate>

Burg, David, and Jesse Ausubel 2021. Moore's Law revisited through Intel chip density. *PLOS ONE* 16(8): e0256245.

Burrell, Jennifer 2016. How the machine ‘thinks’: Understanding opacity in machine learning algorithms. *Big Data and Society* 3(1). <https://journals.sagepub.com/doi/10.1177/2053951715622512>

Bush, George W. 2007. Executive Order 13422: Further Amendment to Executive Order 12866 on Regulatory Planning and Review. 72 FR: 2763. <https://www.federalregister.gov/documents/2007/01/23/07-293/further-amendment-to-executive-order-12866-on-regulatory-planning-and-review>

Business Wire. 2025. “SambaNova Expands Deployment with SoftBank Corp. to Offer Fast AI Inference Across APAC.” Business Wire. <https://www.businesswire.com/news/home/20250305799834/en/SambaNova-Expands-Deployment-with-SoftBank-Corp.-to-Offer-Fast-AI-Inference-Across-APAC/>

Byrne, David, and Carol Corrado. 2017. "ICT Services and their Prices: What do they tell us about Productivity and Technology?" Finance and Economics Discussion Series 2017-015. Board of Governors of the Federal Reserve System. <https://doi.org/10.17016/FEDS.2017.015>

C2PA. 2023. C2PA Technical Specification, Version 2.0. C2PA. <https://c2pa.org/specifications/specifications/2.0/>

Caballero, Ricardo, and Adam Jaffe. 1993. "How High Are the Giants' Shoulders: An Empirical Assessment of Knowledge Spillovers and Creative Destruction in a Model of Economic Growth." In *NBER Macroeconomics Annual 1993, Volume 8*, edited by Olivier J. Blanchard and Stanley Fischer. Cambridge, MA: MIT Press, 15–74.

CAC. 2022. "Provisions on the Administration of Deep Synthesis of Internet-based Information Services." Order No.12. <https://www.chinalawtranslate.com/en/deep-synthesis/>

CAC. 2023. "Interim Measures for the Management of Generative Artificial Intelligence Services." Order No.15. <https://www.chinalawtranslate.com/en/generative-ai-interim/>

Cadwalladr, Carole, and Emma Graham-Harrison. 2018. "Revealed: 50 Million Facebook Profiles Harvested for Cambridge Analytica in Major Data Breach." *The Guardian*. <https://www.theguardian.com/news/2018/mar/17/cambridge-analytica-facebook-influence-us-election>

Cai, Kenrick. 2025. "Meta introduces Llama application programming interface to attract AI developers." *Reuters*. (Reprint: *Investing.com*.) <https://www.investing.com/news/stock-market-news/meta-introduces-llama-application-programming-interface-to-attract-ai-developers-4011605>

Caines, Colin, Sharon Jeon, and Cheyenne Quijano. 2025. "Developments in Chinese Chipmaking." *FEDS Notes*. U.S. Federal Reserve. <https://www.federalreserve.gov/econres/notes/feds-notes/developments-in-chinese-chipmaking-20250117.html>

California Dental Association. 1999. *California Dental Association v. Federal Trade Commission*, 526 U.S. 756. <https://supreme.justia.com/cases/federal/us/526/756>

Campbell-Kelly, Martin, Daniel Garcia-Swartz, Richard Lam, and Yilei Yang. 2015. "Economic and Business Perspectives on Smartphones as Multi-sided Platforms." *Telecommunications Policy* 39(8): 717-34.

Cantwell, Maria 2025. "Pipeline Safety Enforcement Plummets Under Trump." U.S. Senate Committee on Commerce, Science, and Transportation. <https://www.commerce.senate.gov/2025/5/cantwell-pipeline-safety-enforcement-plummets-under-trump>

Carleton, Dennis. 2007. Does antitrust need to be modernized?." *Journal of Economic Perspectives* 21(3): 155-176.

Carlton, Dennis W, and Jeffrey Perloff. 2005. *Modern Industrial Organization*, Vol. 4. Boston: Pearson/Addison Wesley.

Carr, David. 2024. ChatGPT topped 3 billion visits in September. Similarweb. <https://www.similarweb.com/blog/insights/ai-news/chatgpt-topped-3-billion-visits-in-september/>

Carr, Nicholas. 2003. "IT Doesn't Matter." *Harvard Business Review* 81(5): 41-49.

Carr, Nicholas. 2008. *The Big Switch*. New York: W.W. Norton.

Carter, James. 1977. Executive Order 12022—Establishing the National Commission for the Review of Antitrust Laws and Procedures. 42 FR 61441. <https://www.federalregister.gov/executive-order/12022>

Carter, James. 1979. Trade Agreements Program for 1979: Message to the Congress Transmitting a Report. The American Presidency Project. <https://www.presidency.ucsb.edu/documents/trade-agreements-program-for-1979-message-the-congress-transmitting-report>

Carter, James. 1980a. "Economic Renewal Program Remarks Announcing the Program." The American Presidency Project at UCSB. <https://www.presidency.ucsb.edu/documents/economic-renewal-program-remarks-announcing-the-program>

Carter, Jimmy. 1980b. "Patent and Trademark System Reform: Statement on Signing H.R. 6933 Into Law." <https://www.presidency.ucsb.edu/documents/patent-and-trademark-system-reform-statement-signing-hr-6933-into-law>

Carter, James. 1980c. Executive Order 12188—International Trade Functions. 45 FR: 989. <https://www.archives.gov/federal-register/codification/executive-order/12188.html>

Casado, Martin, and Anne Neuberger. 2025. "China Is Winning the Race for AI Robots: Americans who want to compete have to spend more time with regulatory lawyers than engineers." *The Wall Street Journal*. <https://www.wsj.com/opinion/china-is-winning-the-race-for-intelligent-robots-7b2416fe>

Casanova, Lourdes, and Anne Miroux. 2018. "The Competitive Edge in China's Overseas Investments: The Golden Hand of the State." Notes Internacionals CIDOB, no. 205 (November). Barcelona: Barcelona Centre for International Affairs.

Case, Jeffrey., Mark Fedor, Martin Lee Schoffstall, and James Davin. 1990. "A Simple Network Management Protocol (SNMP)." RFC 1157. Internet Engineering Task Force. Unpublished Paper.

Casey, Christopher. 2024. Presidential Authority to Address Tariff Barriers in Trade Agreements. CRS Report IF11400. [https://www.everycrsreport.com/files/2024-10-15\\_IF11400\\_bbbd27645444fe543bee2c23cfb5d6163aaeade1.pdf](https://www.everycrsreport.com/files/2024-10-15_IF11400_bbbd27645444fe543bee2c23cfb5d6163aaeade1.pdf)

Castro, Daniel, and Michael McLaughlin. 2024. "How Innovative Is China in AI?" Information Technology and Innovation Foundation (ITIF). <https://itif.org/publications/2024/08/26/how-innovative-is-china-in-ai/>

Caudel, Edward, and James Magar 1986. Microcomputer system for digital signal processing (U.S. Patent No. 4,577,282). USPTO. <https://image-ppubs.uspto.gov/dirsearch-public/print/downloadPdf/4577282>

Cazzaniga, Mauro, Florence Jaumotte, Longji Li, Giovanni Melina, Augustus Panton, Carlo Pizzinelli, Emma Rockall, and Marina Mendes Tavares. 2024. Gen-AI: Artificial Intelligence and the Future of Work. IMF Staff Discussion Note No. SDN/2024/001. International Monetary Fund. <https://doi.org/10.5089/9798400262548.006>

CB Insights. 2018. Artificial Intelligence Trends To Watch In 2018. CB Insights. [https://www.cbinsights.com/reports/CB-Insights\\_State-of-Artificial-Intelligence-2018.pdf](https://www.cbinsights.com/reports/CB-Insights_State-of-Artificial-Intelligence-2018.pdf)

CB Insights. 2025a. "Who's Winning the AI Coding Race?" CB Insights. <https://www.cbinsights.com/research/report/ai-coding-race/>

CB Insights. 2025b. The AI agent market map. CB Insights. <https://www.cbinsights.com/research/ai-agent-market-map-2025/>

CB Insights. 2025c. "AI Agent Bible: The Ultimate Guide to Agent Disruption." CB Insights. <https://www.cbinsights.com/research/report/ai-agent-bible/>

CB Insights. 2025d. “The Complete List of Unicorn Companies.” CB Insights.  
<https://www.cbinsights.com/research-unicorn-companies>

CBO 1979. The Effects of the Tokyo Round of Multilateral Trade Negotiations on the U.S. Economy: An Updated View. U.S. Government Printing Office.

CEA. 1979. The 1979 Joint Economic Report. U.S. Government Printing Office.  
[https://www.jec.senate.gov/reports/96th%20Congress/The%201979%20Joint%20Economic%20Report%20\(930\).pdf](https://www.jec.senate.gov/reports/96th%20Congress/The%201979%20Joint%20Economic%20Report%20(930).pdf)

CEA. 1996. Economic Report of the President. Washington, D.C.: U.S. Government Printing Office.

CEA. 2024. What Drives the U.S. Services Trade Surplus? Growth in Digitally Enabled Services Exports. White House. <https://bidenwhitehouse.archives.gov/cea/written-materials/2024/06/10/what-drives-the-u-s-services-trade-surplus-growth-in-digitally-enabled-services-exports/>

Center for AI Safety. 2023. “AI Extinction Statement Press Release.” <https://safe.ai/work/press-release-ai-risk>

CFPB. 2023. Consumer Financial Protection Circular 2023-03: Adverse action notification requirements and the proper use of the CFPB’s sample forms provided in Regulation B.  
[https://files.consumerfinance.gov/f/documents/cfpb\\_adverse\\_action\\_notice\\_circular\\_2023-09.pdf](https://files.consumerfinance.gov/f/documents/cfpb_adverse_action_notice_circular_2023-09.pdf)

Challapally, Aditya, Chris Pease, Ramesh Raskar, and Pradyumna Chari. 2025. “The GenAI Divide: State of AI in Business 2025.” MIT Media Lab, NANDA Initiative.  
[https://nanda.media.mit.edu/ai\\_report\\_2025.pdf](https://nanda.media.mit.edu/ai_report_2025.pdf)

Chalmers, Stephanie. 2025. “Commonwealth Bank Backtracks on AI Job Cuts, Apologises for ‘Error’ as Call Volumes Rise.” ABC News (Australia). <https://www.abc.net.au/news/2025-08-21/cba-backtracks-on-ai-job-cuts-as-chatbot-lifts-call-volumes/105679492>

Chander, Anupam, Margot E. Kaminski, and William McGeeveran. 2021. “Catalyzing Privacy Law.” *Minnesota Law Review* 105: 1733-1802.

Chamberlin, Edward. 1927. The Theory of Monopolistic Competition. Unpublished Manuscript.

Chandler, Alfred, Jr. (1977). *The Visible Hand: The Managerial Revolution in American Business*. Cambridge, MA: Harvard University Press.

Chang, Wei, Priscilla Owusu-Mensah, Jordan Everson, and Chelsea Richwine. 2025. "Hospital Trends in the Use, Evaluation, and Governance of Predictive AI, 2023–2024." ASTP Data Brief no. 80. Office of the Assistant Secretary for Technology Policy/Office of the National Coordinator for Health Information Technology. HealthIT.gov. <https://www.healthit.gov/data/data-briefs/hospital-trends-use-evaluation-and-governance-predictive-ai-2023-2024>

Chari, Varadarajan, and Hugo Hopenhayn. 1991. "Vintage human capital, growth, and the diffusion of new technology." *Journal of political Economy* 99(6): 1142-1165.

Chavez, Anthony. 2024. "A New Path for Privacy Sandbox on the Web." Google The Keyword Blog. [https://privacysandbox.com/intl/en\\_us/news/privacy-sandbox-update/](https://privacysandbox.com/intl/en_us/news/privacy-sandbox-update/)

Chen, Mark, et al. 2021. "Evaluating Large Language Models Trained on Code." arXiv preprint arXiv:2107.03374. <https://arxiv.org/abs/2107.03374>

Chevron. 1984. *Chevron U.S.A., Inc. v. Natural Resources Defense Council, Inc.*, 467 U.S. 837. U.S. Supreme Court, June 25. <https://supreme.justia.com/cases/federal/us/467/837>

Chia, Austin. 2024. Open Neural Network Exchange (ONNX) Explained. Splunk. [https://www.splunk.com/en\\_us/blog/learn/open-neural-network-exchange-onnx.html](https://www.splunk.com/en_us/blog/learn/open-neural-network-exchange-onnx.html)

Chicago Board of Trade. 1918. *Board of Trade of Chicago v. United States*, 246 U.S. 231. <https://supreme.justia.com/cases/federal/us/246/231/>

Chien, Colleen, Jorge Contreras, Thomas Cotter, Brian Love, Christopher Seaman, and Norman Siebrasse. 2019. "Enhanced Damages, Litigation Cost Recovery, and Interest." In *Patent Remedies and Complex Products: Toward a Global Consensus*. Edited by C. Bradford Biddle, Jorge Contreras, Brian Love, and Norman Siebrasse. Cambridge: Cambridge University Press, 90-114.

Chin, Josh, and Liza Lin. 2022. *Surveillance State: Inside China's Quest to Launch a New Era of Social Control*. New York: St. Martin's Press.

**Chin, Josh, and Raffaele Huang (2025). “The AI Cold War That Will Redefine Everything.” The Wall Street Journal. <https://www.wsj.com/tech/ai/the-ai-cold-war-that-will-redefine-everything-4e1810b2?msockid=099ba79ac5be67760541b13ac47066c7>**

Chollet, François 2019. On the Measure of Intelligence. arXiv preprint *arXiv:1911.01547*.

<https://ms456000.wpcomstaging.com/wp-content/uploads/2020/02/2020-02-25-d0a1d182d0b0d182d18cd18f-d0bfd180d0be-d098d098.pdf>

Christiano, Paul, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. 2017. “Deep Reinforcement Learning from Human Preferences.” arXiv preprint arXiv:1706.03741.

<https://arxiv.org/abs/1706.03741>

City of Arlington. 2013. City of Arlington v. Federal Communications Commission. No. 11-1545. US Supreme Court, May 20. <https://supreme.justia.com/cases/federal/us/569/290>

CJEU. 2024a. Google LLC and Alphabet Inc. v European Commission (Google Shopping). Case C-48/22 P. Judgment, 10 September 2024.

<https://curia.europa.eu/juris/document/document.jsf?docid=289925&doclang=EN>.

CJEU. 2024b. European Commission v Ireland and Apple Sales International (Apple State Aid). Case C-465/20 P. Judgment, 10 September 2024. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:62020CJ0465>

Clark, Kim. 1985. The interaction of design hierarchies and market concepts in technological evolution. *Research Policy* 14(5): 235-251.

Clark, Kate. 2025. “Nvidia Licenses Groq’s AI Technology as Demand for Cutting-Edge Chips Grows.” The Wall Street Journal. <https://www.wsj.com/articles/nvidia-licenses-groqs-ai-technology-as-demand-for-cutting-edge-chips-grows-87990cbe856818d5eddac44c7b1cdeb8>

Clayton, Lewis, and Eric Stone. 2018. "Claim Drafting, Damages Apportionment in Multi-Component Product Cases." *New York Law Journal* 259(8).

<https://www.paulweiss.com/media/cvbjgfj/11jan2018nylj.pdf>

Clifford, Stephanie, and Quentin Hardy. 2013. "Attention, Shoppers: Store Is Tracking Your Cell." *New York Times*. <https://www.nytimes.com/2013/07/15/business/attention-shopper-stores-are-tracking-your-cell.html>

Clinton White House. 1995. National Economic Council Function. The White House. <https://clintonwhitehouse5.archives.gov/WH/EOP/nec/html/function.html>

Clinton, William J. 1993. Executive Order 12866: Regulatory Planning and Review. *Federal Register* 58(190): 51735-51744.

Clinton, William, and Albert Gore. 1997. A Framework for Global Electronic Commerce. The White House. <https://clintonwhitehouse3.archives.gov/WH/New/Commerce/read.html>

Clinton, William. 1993. Executive Order 12866—Regulatory Planning and Review. 58 FR: 51735. <https://www.archives.gov/federal-register/executive-orders/pdf/12866.pdf>

CMA. 2020. Online Platforms and Digital Advertising Market Study: Final Report. London: Competition and Markets Authority. [https://assets.publishing.service.gov.uk/media/5efc57ed3a6f4023d242ed56/Final\\_report\\_1\\_July\\_2020\\_.pdf](https://assets.publishing.service.gov.uk/media/5efc57ed3a6f4023d242ed56/Final_report_1_July_2020_.pdf)

Coase, Ronald. 1937. "The Nature of the Firm." *Economica*, New Series 4(16): 386–405.

Coase, Ronald. 1960. "The Problem of Social Cost." *Journal of Law and Economics* 3 (October): 1-44.

Cobb, Charles, and Paul Douglas. 1928. "A Theory of Production." *American Economic Review* 18(1): 139-165.

Cochrane, Willard. 1979. *The Development of American Agriculture*. Minneapolis: University of Minnesota Press.

Cohen, Ben. 2025. "It's Waymo's World. We're All Just Riding in It." *The Wall Street Journal*. <https://www.wsj.com/tech/waymo-cars-self-driving-robotaxi-tesla-uber-0777f570>

Cohen, Benjamin and Daniel Sisle. 1970. Export of developing countries in the 1960's. Discussion Paper No. 94, Yale University, Economic Growth Center.

Cohen, Julie, and Mark Lemley. 2001. "Patent Scope and Innovation in the Software Industry." *California Law Review* 89(1): 1–57.

Cohen, Wesley, Richard Nelson, and John Walsh. 2002. "Links and Impacts: The Influence of Public Research on Industrial R&D." *Management Science* 48(1): 1-23.

Columbus, Louis 2017. McKinsey's State Of Machine Learning And AI, 2017.  
<https://www.forbes.com/sites/louiscolumbus/2017/07/09/mckinseys-state-of-machine-learning-and-ai-2017/>

Comai, Sam. 2025. "AI Pilot Programs in K-12 Settings." Education Commission of the States. ECS.  
<https://www.ecs.org/ai-artificial-intelligence-pilots-k12-schools/>

Comin, Diego, and Bart Hobijn. 2010. "An Exploration of Technology Diffusion." *American Economic Review* 100(5): 2031-2059.

Compulife. 2020. *Compulife Software Inc. v. Newman*, 959 F.3d 1288 (11th Cir.).  
<https://law.justia.com/cases/federal/appellate-courts/ca11/18-12004/18-12004-2020-05-20.html>

Constine, Josh. 2018. "Facebook Cuts Off Data Brokers to Avoid Scandals Like Cambridge Analytica's." *TechCrunch*. <https://techcrunch.com/2018/03/28/facebook-will-cut-off-access-to-third-party-data-for-ad-targeting>

Continental Can. 1964. *United States v. Continental Can Co.*, No. 367. US Supreme Court, June 22.  
<https://supreme.justia.com/cases/federal/us/378/441>

Continental T.V. 1977. *Continental T.V., Inc. v. GTE Sylvania Inc.*, No. 76-15. US Supreme Court, June 23. <https://supreme.justia.com/cases/federal/us/433/36>

Contino. 2021. What is compliance as code? Benefits, use cases and tools.  
<https://www.contino.io/insights/compliance-as-code>

Contreras, Jorge. 2015. "A Brief History of FRAND: Analyzing Current Debates in Standard-Setting and Antitrust Law." *Antitrust Law Journal* 80(1): 39–120.

Contreras, Jorge. 2017. "From Private Ordering to Public Law." *Harvard Journal of Law and Technology* 30(Special Symposium): 211-231.

Contreras, Jorge. 2022. "Technical Standards: Fair, Reasonable and Non-Discriminatory (FRAND) Licensing." In *Intellectual Property Licensing and Transactions*, edited by Jorge L. Contreras. Cambridge, UK: Cambridge University Press: 196–247.

Cook, Lisa. 2024. "Artificial Intelligence, Big Data, and the Path Ahead for Productivity." Speech, Board of Governors of the Federal Reserve System.  
<https://www.federalreserve.gov/newsevents/speech/cook20241001a.htm>

Copeland, Curtis. 2013. Federal Rulemaking. CRS Report RL32397. <https://www.congress.gov/crs-product/RL32397>

Corrado, Carol, Charles Hulten, and Daniel Sichel. 2009. "Intangible Capital and U.S. Economic Growth." *Review of Income and Wealth* 55(3): 661–685.

CoSN. 2023. *An Introduction to Interoperability Standards for Education Leaders: Strategically Connecting the K–12 Enterprise*. Washington, DC: Consortium for School Networking.  
[https://www.cosn.org/wp-content/uploads/2023/06/COSN\\_Interop\\_Intro\\_F3.pdf](https://www.cosn.org/wp-content/uploads/2023/06/COSN_Interop_Intro_F3.pdf)

Cowen, Tyler. 2011. *The Great Stagnation*. New York: Dutton.

Cowen, Tyler. 2025. "Why I Think AI Take-off Is Relatively Slow." *Marginal Revolution*.  
<https://marginalrevolution.com/marginalrevolution/2025/02/why-i-think-ai-take-off-is-relatively-slow.html>

Cramton, Peter. 1997. The FCC Spectrum Auctions: An Early Assessment. *Journal of Economics and Management Strategy* 6(3): 431–495.

Cramton, Peter. 2010. "700 MHz Device Flexibility Promotes Competition." Working Paper.  
[www.cramton.umd.edu/papers2010-2014/cramton-700-mhz-device-flexibility-promotes-competition.pdf](http://www.cramton.umd.edu/papers2010-2014/cramton-700-mhz-device-flexibility-promotes-competition.pdf)

Crandall, Robert, and Clifford Winston. 2003. Does Antitrust Policy Improve Consumer Welfare? Assessing the Evidence. *Journal of Economic Perspectives* 17(4): 3-26.

Crandall, Robert. 2001. "Managed Competition in U.S. Telecommunications." *The Brookings Review* 19(2): 19–21.

Crane, Daniel. 2007. Antitrust Modesty. *Michigan Law Review* 105(8): 1193-1212.

Crawford, Susan. 2013. *Captive Audience*. New Haven, CT: Yale University Press.

CRFCAS. 1975. "Structure and Internal Procedures: Recommendations for Change." CRFCAS. <https://www.fordlibrarymuseum.gov/library/document/0019/4520540.pdf>

Crissinger, Cathy. 2000. "AT&T Corp. v. Excel Communications, Inc." *Berkeley Tech. Law Journal* 15(165): 165–184.

Crouzet, Nicolas, and Janice C. Eberly. 2019. "Understanding Weak Capital Investment: The Role of Market Concentration and Intangibles." NBER Working Paper No. 25869. Cambridge, MA: National Bureau of Economic Research.

Crémer, Jacques, Yves-Alexandre De Montjoye, and Heike Schweitzer. 2019. *Competition policy for the digital era*. Brussels: Publications Office of the European Union.

CSIRO. 2015. *Commonwealth Scientific and Industrial Research Organisation v. Cisco Systems, Inc.*, No. 15-1066. US Court of Appeals for the Federal Circuit, December 1. <https://cafc.uscourts.gov/sites/default/files/opinions-orders/15-1066.OPINION.12-1-2015.1.PDF>

Cunningham, Waylon 2024. Big Mac goes Big Tech, with a few hiccups. Reuters. <https://www.reuters.com/business/retail-consumer/big-mac-goes-big-tech-with-few-hiccups-2024-03-16/>

Cuofano, Genaro. 2025. AI Companies Hit \$2.8M Revenue Per Employee: The Death of Traditional Business Models. *FourWeekMBA*. <https://fourweekmba.com/ai-companies-hit-2-8m-revenue-per-employee-the-death-of-traditional-business-models/>

Custers, Bart, Henning Lahmann and Benjamyn Scott. 2025. From liability gaps to liability overlaps: Shared responsibilities and fiduciary duties in AI and other complex technologies. *AI and Society* 40(5): 4035–4050.

Cutter, Chip, and Haley Zimmerman. 2025. CEOs Start Saying the Quiet Part Out Loud: AI Will Wipe Out Jobs. *The Wall Street Journal*. <https://www.wsj.com/tech/ai/ai-white-collar-job-loss-b9856259>

Cyber Creative Institute. 2013. "Evaluation of LTE Essential Patents Declared to ETSI." Version 3.0. Cyber Creative Institute. [https://www.cybersoken.com/wp\\_kanri/wp-content/uploads/2011/12/lte03EN.pdf](https://www.cybersoken.com/wp_kanri/wp-content/uploads/2011/12/lte03EN.pdf)

Cyberspace Administration of China. 2025. "关于印发《人工智能生成合成内容标识办法》的通知."

[https://www.cac.gov.cn/2025-03/14/c\\_1743654684782215.htm](https://www.cac.gov.cn/2025-03/14/c_1743654684782215.htm)

Dabir, Rupesh. 2025. The future of code: How AI is transforming software development. Forbes.

<https://www.forbes.com/councils/forbestechcouncil/2025/04/04/the-future-of-code-how-ai-is-transforming-software-development/>

Dadaboyev, Sherzodbek Murodilla Ugli, Jasmina Abdullayeva, Naval Abbosova, Afina Suleymenova and

Komila Mamadjanova. 2025. "Role of artificial intelligence in employee recruitment: systematic review and future research directions." Discover Global Society 3(99).

Daffodil Software. 2025. Top 10 Image Recognition APIs in 2025.

<https://insights.daffodilsw.com/blog/top-10-image-recognition-apis-for-app-development>

Dahlman, Erik, Stefan Parkvall, and Johan Sköld. 2011. 4G: LTE/LTE-Advanced for Mobile Broadband.

Oxford: Academic Press.

Dally, Bill. 2023. Heeding Huang's Law: Video Shows How Engineers Keep the Speedups Coming.

NVIDIA Blog. <https://blogs.nvidia.com/blog/huangs-law-dally-hot-chips/>

Damle, Sy, Alli Stillman, Britt Lovejoy, Ivana Dukanovic. 2024. "The ELVIS Act: Tennessee Shakes Up

Its Right of Publicity Law and Takes On Generative AI." Latham and Watkins LLP.

<https://www.lw.com/admin/upload/SiteAttachments/The-ELVIS-Act-Tennessee-Shakes-Up-Its-Right-of-Publicity-Law-and-Takes-On-Generative-AI.pdf>

Dance, Gabriel, Michael LaForgia, and Nicholas Confessore. 2018. "As Facebook Raised a Privacy Wall, It Carved an Opening for Tech Giants." The New York Times.

<https://www.nytimes.com/2018/12/18/technology/facebook-privacy-data-sharing-ebay-microsoft-amazon-netflix-spotify.html>

Darveau-Garneau, Matt, Manuel Frey, Beth Goldberg, and Nick South. 2022. "As the Cookie Crumbles:

Three Strategies for Advertisers to Thrive." McKinsey & Company.

<https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/as-the-cookie-crumbles-three-strategies-for-advertisers-to-thrive>

Dastin, Jeffrey. 2018. Amazon scraps secret AI recruiting tool that showed bias against women. Reuters.

<https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>

Databricks. 2024. Databricks Unified Data and AI Platform. <https://databricks.com/>

DataCenterDynamics. 2025. Iron Mountain breaks ground on first data center in Miami.

DataCenterDynamics. <https://www.datacenterdynamics.com/en/news/iron-mountain-breaks-ground-on-first-data-center-in-miami/>

David, Paul. 1990. "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox." *American Economic Review* 80(2): 355-361.

Davies, Paul. 2023. "Potentially Disastrous for Innovation: Tech Sector Reacts to the EU AI Act."

Euronews. <https://www.euronews.com/next/2023/12/15/potentially-disastrous-for-innovation-tech-sector-says-eu-ai-act-goes-too-far>

Daws, Ryan. 2025. "AI and its impact on software development jobs." *Developer Tech*.

<https://www.developer-tech.com/news/ai-impact-on-software-development-jobs/>

de Boer. 2025. Nvidia Crushes Competition With 94% GPU Market Share. *Yahoo Finance*.

<https://finance.yahoo.com/news/nvidia-crushes-competition-94-gpu-144404532.html>

De Loecker, Jan, and Jan Eeckhout. 2017. "The Rise of Market Power and the Macroeconomic Implications." NBER Working Paper No. 23687. Cambridge, MA: National Bureau of Economic Research.

De Loecker, Jan, József Eeckhout, and Gabriel Unger. 2020. "The Rise of Market Power and the Macroeconomic Implications." *Quarterly Journal of Economics* 135(2): 561–644.

Dean, Jeffrey, et al. 2012. "Large Scale Distributed Deep Networks." *Advances in Neural Information Processing Systems* 25 (NIPS 2012): 1223–1231.

Dedrick, Jason, Kenneth Kraemer and Greg Linden. 2010. "Who profits from innovation in global value chains?: a study of the iPod and notebook PC." *Industrial and Corporate Change* 19(1): 81-116.

DeLamarter, Richard. 1986. *Big Blue*. New York: Dodd, Mead and Company.

Delrahim, Makan 2018. DoJ Business Review Letter for Avanci LLC. DoJ.  
<https://www.justice.gov/atr/page/file/1298626/download>

Demsetz, Harold 1973. Industry Structure, Market Rivalry, and Public Policy. *The Journal of Law and Economics* 16(1): 1-9.

Dentsply International. 2005. *United States v. Dentsply International, Inc.*, No. 03-4097, 3d Cir., Feb. 24.  
<https://law.justia.com/cases/federal/appellate-courts/F3/399/181/527982>

Desmarais, Anna. 2025. "Two city governments in Denmark are moving away from Microsoft amid Trump and US Big Tech concerns." *Euronews Next*. <https://www.euronews.com/next/2025/06/12/two-city-governments-in-denmark-are-moving-away-from-microsoft-amid-trump-and-us-big-tech>

Destler, I. M. (Mac) 1996. *The National Economic Council*. Washington, DC: Peterson Institute for International Economics.

Deuel. 1995. *In re Deuel*. No. 94-1202. US Federal Circuit Court, March 28.  
<https://law.justia.com/cases/federal/appellate-courts/F3/51/1552/617820/>

Deutsche Telekom 2020. *New York v. Deutsche Telekom AG*, No. 1:19-cv-05434, 2020 WL 1879475 (S.D.N.Y.). <https://law.justia.com/cases/federal/district-courts/new-york/nysdce/1:2019cv05434/517350/409>

Devine, Warren. 1983: "From Shafts to Wires: Historical Perspectives on Electrification." *The Journal of Economic History* 43(2): 347-372.

Dezan Shira and Associates. 2024. "Pre-Tax Super Deduction of R&D Expenses in China – An Explainer." *China Briefing*. <https://www.china-briefing.com/news/china-rd-expenses-pre-tax-super-deduction-explainer/>

Diab, Wael William, and Mike Mullane. 2024. “How the ISO and IEC Are Developing International Standards for the Responsible Adoption of AI.” UNESCO. <https://www.unesco.org/en/articles/how-iso-and-iec-are-developing-international-standards-responsible-adoption-ai>

Diamond. 1981. *Diamond, Commissioner of Patents and Trademarks v. Diehr et al.* No. 79-1112. US Supreme Court, March 3. <https://supreme.justia.com/cases/federal/us/450/175/>

Dickinson, Grace. 2023. Are AI-powered robots the future of dishwashing in the restaurant industry? Back of House. <https://backofhouse.io/resources/ai-powered-robots-future-dishwashing-restaurant-industry-automation-technology>

DICOM Standards Committee. 2025. “Artificial Intelligence and DICOM (WG-23): Interoperability Guidance and Resources.” DICOM Standard. <https://www.dicomstandard.org/ai/>

Disu, Opemipo. 2024. Amazon Bedrock. Microtica (Dev.to). <https://dev.to/microtica/amazon-bedrock-a-practical-guide-for-developers-and-devops-engineers-4k8g>

Dodge, Jesse, Maarten Sap, Ana Marasović, William Agnew, Gabriel Ilharco, Dirk Groeneveld, Margaret Mitchell, and Matt Gardner. 2021. “Documenting Large Webtext Corpora.” Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, 1286–1305. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.emnlp-main.98>.

DoE. 2015. Frequently Asked Questions—Cloud Computing. Privacy Technical Assistance Center (PTAC). DoE. [https://studentprivacy.ed.gov/sites/default/files/resource\\_document/file/FAQ\\_Cloud\\_Computing\\_0.pdf](https://studentprivacy.ed.gov/sites/default/files/resource_document/file/FAQ_Cloud_Computing_0.pdf)

Doe. 2016. *Doe No. 1 v. Backpage.com, LLC.* No. 15-1724. US Court of Appeals, First Circuit, March 14. <https://law.justia.com/cases/federal/appellate-courts/ca1/15-1724/15-1724-2016-03-14.html>

DoJ and FTC. 1992. Horizontal Merger Guidelines. <https://www.justice.gov/archives/atr/1992-merger-guidelines>

DoJ and FTC. 1995. Antitrust Guidelines for the Licensing of Intellectual Property. DoJ and FTC. <https://www.justice.gov/atr/archived-1995-antitrust-guidelines-licensing-intellectual-property>

DoJ and FTC. 2010. "Horizontal Merger Guidelines." DoJ and FTC.  
<https://www.justice.gov/atr/horizontal-merger-guidelines-08192010>

DOJ and USPTO. 2025. "Statement of Interest of the United States." Radian Memory Systems LLC v. Samsung Electronics Co., No. 2:24-cv-1073 (E.D. Tex.). <https://justice.gov/atr/media/1404506/dl>

DoJ. 1982. Merger Guidelines. <https://www.justice.gov/archives/atr/1982-merger-guidelines>

DoJ. 1984. "Non-Horizontal Merger Guidelines." <https://www.justice.gov/archives/atr/1984-merger-guidelines>

DoJ. 1988. Antitrust Enforcement Guidelines for International Operations. Washington, D.C.: DoJ.

DoJ. 1997. "Business Review Letter to MPEG LA." DoJ Antitrust Division. Washington, D.C.: DoJ.

DoJ. 2004. "Justice Department Implements the Standards Development Organization Advancement Act of 2004." Press Release. [https://www.justice.gov/archive/atr/public/press\\_releases/2004/204345.htm](https://www.justice.gov/archive/atr/public/press_releases/2004/204345.htm)

DoJ. 2008. Competition and Monopoly: Single-Firm Conduct Under Section 2 of the Sherman Act. Washington, DC: DoJ.

DoJ. 2009. Justice Department Withdraws Report on Antitrust Monopoly Law. Press Release. <https://www.justice.gov/archives/opa/pr/justice-department-withdraws-report-antitrust-monopoly-law>

DoJ. 2015. "Business Review Letter to Institute of Electrical and Electronics Engineers, Incorporated." Press Release. [https://www.justice.gov/sites/default/files/opa/press-releases/attachments/2015/02/02/ieee\\_business\\_review\\_letter.pdf](https://www.justice.gov/sites/default/files/opa/press-releases/attachments/2015/02/02/ieee_business_review_letter.pdf)

DoJ. 2019. "Complaint." United States v. Sabre Corp., et al., No. 1:19-cv-01548 (D. Del.), filed Aug. 20, 2019. Washington, D.C.: DoJ.

DoJ. 2020. "Plaintiff's Pretrial Brief [Redacted Public Version]." United States v. Sabre Corp., et al., No. 1:19-cv-01548 (D. Del.), filed Jan. 15, 2020. Washington, D.C.: DoJ.

DoJ. n.d.a. Criminal Enforcement. <https://www.justice.gov/atr/criminal-enforcement>

DoJ. n.d.b. Civil Enforcement. <https://www.justice.gov/atr/civil-enforcement>

DoJ/USPTO 2025. "Joint Comment on the Public Interest in Certain DRAM Devices, Products Containing the Same and Components Thereof." ITC Investigation No. 337-TA-3854.

<https://justice.gov/atr/media/1419496/dl>

DoJ/USPTO/NIST. 2019. Policy Statement on Remedies for Standards-Essential Patents Subject to Voluntary F/RAND Commitments. DoJ/USPTO/NIST

<https://www.justice.gov/atr/page/file/1228016/dl?inline>

DoJ/USPTO/NIST. 2022. Withdrawal of 2019 Policy Statement on Remedies for Standards-Essential Patents Subject to Voluntary F/RAND Commitments. DoJ/USPTO/NIST.

<https://www.uspto.gov/sites/default/files/documents/SEP2019-Withdrawal.pdf>

DoL. 2024. History of Federal Minimum Wage Rates Under the Fair Labor Standards Act, 1938–2009.

<https://www.dol.gov/agencies/whd/minimum-wage/history/chart>

Domen. 2021. *Domen v. Vimeo, Inc.* No. 20-616. US Court of Appeals, Second Circuit, March 11.

<https://law.justia.com/cases/federal/appellate-courts/ca2/20-616/20-616-2021-03-11.html>

Domingos, Pedro. 2015. *The Master Algorithm*. New York: Basic Books.

Donnellan, Douglas, Andy Lawrence, Daniel Bizo, Peter Judge, John O'Brien, Jacqueline Davis, Max Smolaks, Jabari Williams-George, and Rose Weinschenk. 2024. Uptime Institute Global Data Center Survey 2024. Uptime Institute. [https://datacenter.uptimeinstitute.com/rs/711-RIA-](https://datacenter.uptimeinstitute.com/rs/711-RIA-145/images/2024.GlobalDataCenterSurvey.Report.pdf)

[145/images/2024.GlobalDataCenterSurvey.Report.pdf](https://datacenter.uptimeinstitute.com/rs/711-RIA-145/images/2024.GlobalDataCenterSurvey.Report.pdf)

DoS. 1980. 1980 Editorial Note: Executive Order 12188 and the Trade Reorganization Plan. *Foreign Relations of the United States, 1977–1980, Volume XXVIII. Organization and Management of Foreign Policy*. Document 121. <https://history.state.gov/historicaldocuments/frus1977-80v28/d121>

Doss, Christopher, Robert Bozick, Heather Schwartz, Lisa Chu, Lydia Rainey, Ashley Woo, Justin Reich, Jesse Dukes. 2025. *AI Use in Schools Is Quickly Increasing but Guidance Lags Behind: Findings from the RAND Survey Panels*. Santa Monica, CA: RAND Corporation.

[https://www.rand.org/pubs/research\\_reports/RRA4180-1.html](https://www.rand.org/pubs/research_reports/RRA4180-1.html)

Dr. Miles Medical. 1911. *Dr. Miles Medical Co. v. John D. Park and Sons Co.*, 220 U.S. 373.  
<https://supreme.justia.com/cases/federal/us/220/373>

Draghi, Mario. 2024. *The Future of European Competitiveness (Executive Summary Highlights)*.  
European Commission. [https://commission.europa.eu/document/download/97e481fd-2dc3-412d-be4c-f152a8232961\\_en](https://commission.europa.eu/document/download/97e481fd-2dc3-412d-be4c-f152a8232961_en)

Drew, Rocket. 2025. *AI Agenda. The Information*. <https://www.theinformation.com/articles/chinese-tech-behind-amazons-humanoid-robots?rc=tuorsb>

Dreyfuss, Rochelle Cooper. 1989. "The Federal Circuit: A Case Study in Specialized Courts." *New York University Law Review* 64(1): 1–77.

Dudley, Susan. 2020. "OMB's Circular A-4 and the History of Regulatory Impact Analysis." *Administrative Law Review* 72(1): 11–32.

Duffer, Sara, Singh Ashish, Denis Batalov, and Peter Hallinan. 2025. "Building trust in AI: The AWS approach to the EU AI Act." *AWS Machine Learning Blog*. <https://aws.amazon.com/blogs/machine-learning/building-trust-in-ai-the-aws-approach-to-the-eu-ai-act>

Dumont, Andrew. 2024. "Changes in the U.S. Economy and Rural-Urban Employment Disparities." *FEDS Notes, Board of Governors of the Federal Reserve System (January 19, 2024)*.  
<https://www.federalreserve.gov/econres/notes/feds-notes/changes-in-the-us-economy-and-rural-urban-employment-disparities-20240119.html> [harvardmagazine.com]

Dupont. 1956. *United States v. E. I. du Pont de Nemours and Co.*, 351 U.S. 377.  
<https://supreme.justia.com/cases/federal/us/351/377>

Durie, Daralyn, and Mark Lemley. 2008. "A Realistic Approach to the Obviousness of Inventions." *William and Mary Law Review* 50(3): 989–1019.

Dwork, Cynthia, and Aaron Roth. 2014. "The Algorithmic Foundations of Differential Privacy." *Foundations and Trends in Theoretical Computer Science* 9(3–4): 211–407.

Dyroff. 2019. *Dyroff v. Ultimate Software Group, Inc.* No. 18-15175. US Court of Appeals, Ninth Circuit, August 20. <https://law.justia.com/cases/federal/appellate-courts/ca9/18-15175/18-15175-2019-08-20.html>

Dzieza, Josh, and Hayden Field. 2025. "Feeding the Machine." The Verge.  
<https://www.theverge.com/2025/12/15/feeding-the-machine-ai-training-data-mercors-scale-surge>

D'Silva, Charmaine. 2020. "11 Weeks of Android: Privacy and Security." Android Developers Blog.  
<https://android-developers.googleblog.com/2020/06/11-weeks-of-android-privacy-and-security.html>

Easterbrook, Frank. 1984. The Limits of Antitrust. *Texas Law Review* 63(1): 1-40.

Eastman Kodak. 1992. *Eastman Kodak Co. v. Image Technical Services, Inc.*, 504 U.S. 451.  
<https://supreme.justia.com/cases/federal/us/504/451>

eBay. 2000. *eBay, Inc. v. Bidder's Edge, Inc.* No. C-99-21200 RMW. US District Court, N.D. California, May 24. <https://law.justia.com/cases/federal/district-courts/FSupp2/100/1058/2478126>

eBay. 2006. *eBay, Inc. v. MercExchange, L.L.C.*, No. 05-130. US Supreme Court, May 15.  
<https://supreme.justia.com/cases/federal/us/547/388>

Edelman, Benjamin. 2015. "Does Google Leverage Market Power Through Tying and Bundling?" *Journal of Competition Law and Economics* 11(2): 365–400.

Edlin, Aaron. 2002. "Stopping Above-Cost Predatory Pricing." *Yale Law Journal* 111(4): 941–991.

EDM Council. n.d. *Cloud Data Management Capabilities Framework*. EDM Council.  
<https://edmcouncil.org/>

Edwards, Chris. 2005. *Agricultural Policy*. Cato Handbook on Policy (6th ed).  
<https://www.cato.org/sites/cato.org/files/serials/files/cato-handbook-policymakers/2005/9/hb109-31.pdf>

Efrati, Amir, and Stephanie Palazzolo. 2025. "OpenAI Targets AGI with System That Thinks Like a Pro Engineer." *The Information*. <https://www.theinformation.com/articles/openai-targets-agi-with-system-that-thinks-like-a-pro-engineer>

EIA. 2018. "PURPA-qualifying capacity increases, but it's still a small portion of added renewables." *Today in Energy*. <https://www.eia.gov/todayinenergy/detail.php?id=36912>

Eichengreen, Barry. 1992. *Golden Fetters: The Gold Standard and the Great Depression, 1919–1939*. New York, NY: Oxford University Press.

Eisenach, Jeffrey and Robert Kulick. 2020. Economic impacts of mobile broadband innovation: Evidence from the transition to 4G. AEI Economics Working Paper 2020-05. American Enterprise Institute. <https://aei.org/wp-content/uploads/2020/06/Eisenach-Kulick-Mobile-Broadband-Innovation-WP.pdf>

Elhauge, Einer 2003. Defining Better Monopolization Standards. *Stanford Law Review* 56: 253-234.

Elmore, John. 2024. "The Foxconn Connection: Unraveling Apple's Outsourcing Strategy." *TheTechyLife*. <https://thetechylife.com/why-does-apple-outsource-to-foxconn/>

Elsner, Mark, Grace Atkinson, and Saadia Zahidi. 2025. *The Global Risks Report 2025: 20th Edition*. World Economic Forum. <https://www.weforum.org/publications/global-risks-report-2025/>

Elzinga, Kenneth, and Thomas Hogarty. 1973. "The Problem of Geographic Market Delineation in Antimerger Suits." *The Antitrust Bulletin* 18(1): 45–81.

Emberson, Luke. 2025. Frontier open models may surpass  $10^{26}$  FLOP of training compute before 2026. Epoch AI. <https://epochai.org/data-insights/frontier-open-models-may-surpass-10-26-flop-of-training-compute-before-2026>

Emory Solutions. 2024. *The Problem with AI Hallucinations: Why They Happen, How to Prevent, and What Risks Exist*. Emory Solutions. <https://www.emorysolutions.com/the-problem-with-ai-hallucinations-why-they-happen-how-to-prevent-and-what-risks-exist/>

Engler, Alex 2023. "The Evolution of AI Regulation in the United States." Washington, DC: Brookings Institution Report.

EPA. 2014. *Regulatory Impact Analysis: Control of Air Pollution from Motor Vehicles: Tier 3 Motor Vehicle Emission and Fuel Standards*. EPA-420-R-14-005. <https://www.govinfo.gov/content/pkg/FR-2014-04-28/pdf/2014-06954.pdf>

Epoch AI. 2024. "Data on Machine Learning Hardware." Epoch AI. <https://epoch.ai/data/machine-learning-hardware>

Epstein, Richard. 2006. *How Progressives Rewrote the Constitution*. Washington, D.C.: Cato Institute.

ETSI. 2024. *ETSI IPR Policy*. Sophia Antipolis: ETSI.

ETSI. 2019. "Intellectual Property Rights Policy." ETSI Rules of Procedure, Annex 6. Sophia Antipolis: European Telecommunications Standards Institute.  
[https://portal.etsi.org/directives/40\\_directives\\_apr\\_2019.pdf](https://portal.etsi.org/directives/40_directives_apr_2019.pdf)

EU General Court. 2022. Google and Alphabet v. European Commission. Case T-604/18, ECLI:EU:T:2022:541. September 14 2022.  
<https://curia.europa.eu/juris/document/document.jsf?docid=265435>

EU General Court. 2024. Google and Alphabet v European Commission (Google AdSense for Search). Case T-334/19. Judgment, 18 September 2024. <https://eur-lex.europa.eu/eli/C/2024/6420/oj/eng>

EU. 2022. Regulation (EU) 2022/1925 (Digital Markets Act). OJ L 265, 12.10.2022, 1–66. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32022R1925>

EU. 2024. Regulation (EU) 2024/1689 of the European Parliament and of the Council (Artificial Intelligence Act). [https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=OJ:L\\_202401689](https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=OJ:L_202401689)

EuroHPC JU. 2024. Consolidated Annual Activity Report 2024. [https://www.eurohpc-ju.europa.eu/document/download/a8889156-6b46-4e91-bf70-e4baf5cdb68f\\_en](https://www.eurohpc-ju.europa.eu/document/download/a8889156-6b46-4e91-bf70-e4baf5cdb68f_en)

European Commission. 2024. "Commission sends Statement of Objections to Microsoft over possibly abusive tying practices regarding Teams." Press release.  
[https://ec.europa.eu/commission/presscorner/api/files/document/print/en/ip\\_24\\_3446/IP\\_24\\_3446\\_EN.pdf](https://ec.europa.eu/commission/presscorner/api/files/document/print/en/ip_24_3446/IP_24_3446_EN.pdf)

European Commission. 2025a. "Commission launches two strategies to speed up AI uptake in European industry and science." Press release.  
[https://ec.europa.eu/commission/presscorner/api/files/document/print/en/ip\\_25\\_2299/IP\\_25\\_2299\\_EN.pdf](https://ec.europa.eu/commission/presscorner/api/files/document/print/en/ip_25_2299/IP_25_2299_EN.pdf)

European Commission. 2025b. "Commission finds Apple and Meta in breach of the Digital Markets Act." Press release.  
[https://ec.europa.eu/commission/presscorner/api/files/document/print/en/ip\\_25\\_1085/IP\\_25\\_1085\\_EN.pdf](https://ec.europa.eu/commission/presscorner/api/files/document/print/en/ip_25_1085/IP_25_1085_EN.pdf)

European Court of Auditors. 2025. Special Report 12/2025 — The EU’s strategy for microchips: Ambitious targets, but funding and implementation risks remain. Luxembourg: Publications Office of the European Union. [https://www.eca.europa.eu/ECAPublications/SR-2025-12/SR-2025-12\\_EN.pdf](https://www.eca.europa.eu/ECAPublications/SR-2025-12/SR-2025-12_EN.pdf)

European Parliament. “Artificial Intelligence Act.” European Parliament legislative resolution of 13 March 2024, P9\_TA(2024)0138. [https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=EP%3AP9\\_TA%282024%290138](https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=EP%3AP9_TA%282024%290138)

European Union. 2024. Regulation (EU) 2024/1689 laying down harmonised rules on Artificial Intelligence (AI Act). Official Journal of the European Union. [https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=OJ%3AL\\_202401689](https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=OJ%3AL_202401689)

Evans, Benedict. 2018. “Presentation: The End of the Beginning.” Annual presentation. a16z Summit. <https://www.ben-evans.com/benedictevans/2018/11/16/the-end-of-the-beginning>

Evans, David. 2008. Antitrust Issues Raised by the Emerging Global Internet Economy. *Northwestern University Law Review* 102(4): 1987-2008.

Evans, David, and Richard Schmalensee. 2001. "Some Economic Aspects of Antitrust Analysis in Dynamically Competitive Industries." NBER Working Paper No. 8268. Cambridge, MA: National Bureau of Economic Research.

Evans, David, and Richard Schmalensee. 2016. *Matchmakers*. Boston, MA: Harvard Business Review Press.

Evans, David, Andrei Hagiu, and Richard Schmalensee. 2006. *Invisible Engines*. Cambridge, MA: MIT Press.

Exmark Manufacturing. 2018. *Exmark Manufacturing Co. v. Briggs and Stratton Power Products Group, LLC*. No. 2016-2197. US Court of Appeals, Federal Circuit, January 12. <https://law.justia.com/cases/federal/appellate-courts/cafc/16-2197/16-2197-2018-01-12.html>

Facebook 2022. "About Lookalike Audiences." Meta for Business. <https://www.facebook.com/business/help/164749007013531>

Facebook, Inc. 2021. "Annual Report 2020 (Form 10-K)."

<https://www.sec.gov/Archives/edgar/data/1326801/000132680121000014/fb-20201231.htm>

Facebook. 2007a. "Facebook Unveils Platform for Developers of Social Applications." Facebook Newsroom. <https://about.fb.com/news/2007/05/facebook-unveils-platform-for-developers-of-social-applications/>

Facebook. 2007b. "Facebook Develops Network Portals, New Inbox and Updates Site Design." Facebook Newsroom. <https://about.fb.com/news/2007/04/facebook-develops-network-portals-new-inbox-and-updates-site-design/>

Facebook. 2007c. "Facebook Unveils Facebook Ads." Facebook Newsroom.

<https://about.fb.com/news/2007/11/facebook-unveils-facebook-ads/>

Facebook. 2013. "Announcing News Feed FYI: A Series of Blogs on News Feed Ranking." Facebook Newsroom. <https://about.fb.com/news/2013/08/announcing-news-feed-fyi-a-series-of-blogs-on-news-feed-ranking/>

Facebook. 2015. "Meta Blueprint (Facebook Blueprint): Learn New Skills to Build Your Brand or Business." Meta Blueprint (training portal). <https://www.facebookblueprint.com/student/catalog>

Facebook. 2017. "Celebrating 5 Million Active Advertisers." Facebook Newsroom.

<https://about.fb.com/news/2017/04/celebrating-5-million-active-advertisers>

Facebook. 2019. In re Facebook, Inc., Consumer Privacy User Profile Litigation. No. 18-md-02843-VC.

US District Court, N.D. California, September 9. <https://law.justia.com/cases/federal/district-courts/california/candce/5:2018md02843/376023/233>

Facebook. 2025. "About Custom Audiences." Meta Business Help Center.

<https://www.facebook.com/business/help/744354708981227>

Fair Housing Council. 2008. Fair Housing Council of San Fernando Valley v. Roommates.com, LLC. No. 04-57115. US Court of Appeals, Ninth Circuit, April 3. <https://law.justia.com/cases/federal/appellate-courts/F3/521/1157>

Fairfield, Joshua. 2017. Owned. New York: Cambridge University Press.

Fairless, Tom, and David Luhnow. 2025. "The Tech Industry Is Huge—and Europe's Share of It Is Very Small." Wall Street Journal. <https://www.wsj.com/tech/europe-big-tech-ai-1f3f862c>

Farrell, Joseph, and Carl Shapiro. 2010. "Antitrust Evaluation of Horizontal Mergers." *B.E. Journal of Theoretical Economics* 10(1): Article 9.

Farrell, Joseph, and Garth Saloner. 1985. "Standardization, Compatibility, and Innovation." *The RAND Journal of Economics* 16(1): 70-83.

Farrell, Joseph, and Paul Klemperer. 2007. "Coordination and Lock-In: Competition with Switching Costs and Network Effects." In *Handbook of Industrial Organization*, Vol. 3, edited by Mark Armstrong and Robert Porter. Amsterdam: Elsevier, 1967-2072.

Farrell, Joseph, and Timothy Simcoe. 2012. "Choosing the Rules for Formal Standardization." *Review of Industrial Organization* 40(3): 159–88.

FCC. 2003. Review of the Section 251 Unbundling Obligations of Incumbent Local Exchange Carriers. Report and Order and Order on Remand and Further Notice of Proposed Rulemaking (FCC 03-36). <https://docs.fcc.gov/public/attachments/FCC-03-36A1.pdf>

FCC. 2005. Appropriate Framework for Broadband Access to the Internet over Wireline Facilities. Report and Order and Notice of Proposed Rulemaking (FCC 05-150). <https://docs.fcc.gov/public/attachments/FCC-05-150A2.pdf>

FCC. 2008. Auction 73: 700 MHz Band. Web page. <https://www.fcc.gov/auction/73>

FCC. 2009. Petition for Declaratory Ruling to Clarify Provisions of Section 332(c)(7)(B) to Ensure Timely Siting Review. Declaratory Ruling (FCC 09-99). [https://apps.fcc.gov/edocs\\_public/attachmatch/FCC-09-99A1.pdf](https://apps.fcc.gov/edocs_public/attachmatch/FCC-09-99A1.pdf)

FCC. 2013. Annual Assessment of the Status of Competition in the Market for the Delivery of Video Programming. Report (FCC 13-99). <https://docs.fcc.gov/public/attachments/FCC-13-99A1.pdf>

FCC. 2014. Protecting and Promoting the Open Internet. Notice of Proposed Rulemaking (FCC 14-61). [https://apps.fcc.gov/edocs\\_public/attachmatch/FCC-14-61A1\\_Rcd.pdf](https://apps.fcc.gov/edocs_public/attachmatch/FCC-14-61A1_Rcd.pdf)

Federal Reserve Bank of St. Louis. 2024. Private Business Sector: Total Factor Productivity. Retrieved January 26, 2025, from <https://fred.stlouisfed.org/series/MPU4900013>

Federico, Giovanni, and Antonio Tena-Junguito. 2018. "World Trade, 1800-1938: A New Synthesis." *Revista de Historia Económica / Journal of Iberian and Latin American Economic History* 37(1): 9-41.

Federighi, Craig. 2019. "Apple previews iOS 13." Apple Newsroom.  
<https://www.apple.com/newsroom/2019/06/apple-previews-ios-13/> [apple.com]

Feigenbaum, James, and Daniel Gross. 2022. "Answering the Call of Automation: How the Labor Market Adjusted to the Mechanization of Telephone Operation." NBER Working Paper No. 28061. Cambridge, MA: National Bureau of Economic Research.

Feist. 1991. *Feist Publications, Inc. v. Rural Telephone Service Co.* No. 89-1909. US Supreme Court, March 18. <https://supreme.justia.com/cases/federal/us/499/340/>

Feldman. 2007. *Feldman v. Google, Inc.* No. 06-cv-02570-JG. US District Court, E.D. Pennsylvania, March 30. <https://law.justia.com/cases/federal/district-courts/pennsylvania/paedce/2:2006cv02570/206160/46>

FHWA (Federal Highway Administration). 2002. Appendix A Economic Effects of Transportation: The Freight Story. [https://ops.fhwa.dot.gov/freight/freight\\_analysis/improve\\_econ/appa.htm](https://ops.fhwa.dot.gov/freight/freight_analysis/improve_econ/appa.htm)

Field, Hayden. 2025a. "Google agrees to new \$1 billion investment in Anthropic." CNBC.  
<https://www.cnbc.com/2025/01/22/google-agrees-to-new-1-billion-investment-in-anthropic.html>

Field, Hayden. 2025b. "Perplexity AI wrapping \$500 million raise at \$14 billion valuation." CNBC.  
<https://www.cnbc.com/2025/05/12/perplexity-funding-round-comet.html>

Fielding, Sarah. 2023. Amazon says its new AI-powered robots reduce fulfilment time by 25 percent. Engadget. <https://www.engadget.com/amazon-says-its-new-ai-powered-robots-reduce-fulfilment-time-by-25-percent-122517342.html>

Filippucci, Francesco, Peter Gal, and Cecilia Susanna Jona Lasinio. 2024. "The impact of Artificial Intelligence on productivity, distribution and growth." Working Paper.  
<https://iris.luiss.it/handle/11385/240418>

Filistrucchi, Lapo, Damien Geradin, Eric Van Damme, and Pauline Affeldt. "Market definition in two-sided markets: Theory and practice." *Journal of Competition Law and Economics* 10(2): 293-339.

FillApp. 2025. The State of AI Browser Agents in 2025. FillApp Blog. <https://fillapp.ai/blog/the-state-of-ai-browser-agents-2025/>

Findlay, Ronald, and Kevin O'Rourke. 2007. *Power and Plenty: Trade, War, and the World Economy in the Second Millennium*. Princeton: Princeton University Press.

Finn, Chelsea, Sergey Levine, and Pieter Abbeel. "Guided cost learning: Deep inverse optimal control via policy optimization." In *International conference on machine learning (PMLR, 2016)*: 49-58.

FinRegLab. 2025. Advancing the credit ecosystem: Machine learning and cash flow data in consumer underwriting. Empirical White Paper. FinRegLab. <https://finreglab.org/research/advancing-the-credit-ecosystem-machine-learning-cash-flow-data-in-consumer-underwriting/>

Fishback, Price. 2007. "The New Deal." *The Concise Encyclopedia of Economics*. Indianapolis: Liberty Fund.

Fishman, Charles. 2006. *The Wal-Mart Effect*. New York: Penguin.

Fitch, Asa. 2025. "AI's Big Leaps Are Slowing. That Could Be a Good Thing." *Wall Street Journal*. <https://www.wsj.com/tech/ai/ais-big-leaps-are-slowing-that-could-be-a-good-thing-34c87619>

Force. 2019. *Force v. Facebook, Inc.* No. 18-397. US Court of Appeals, Second Circuit, July 31. <https://law.justia.com/cases/federal/appellate-courts/ca2/18-397/18-397-2019-07-31.html>

Foremost Pro Color. 1983. *Foremost Pro Color, Inc. v. Eastman Kodak Co.*, 703 F.2d 534. <https://law.justia.com/cases/federal/appellate-courts/F2/703/534/11835>

Forgash, Emily, and Agnee Ghosh. 2025. "OpenAI, Nvidia Fuel \$1 Trillion AI Market With Web of Circular Deals." *Bloomberg News*. <https://www.bloomberg.com/news/features/2025-10-07/openai-s-nvidia-amd-deals-boost-1-trillion-ai-boom-with-circular-deals>

Forlini, Emily. 2023. Samsung Software Engineers Busted for Pasting Proprietary Code Into ChatGPT. *PC Magazine*. <https://www.pcmag.com/news/samsung-software-engineers-busted-for-pasting-proprietary-code-into-chatgpt>

Free Enterprise Fund 2010. Free Enterprise Fund and Beckstead and Watts, LLP v. Public Company Accounting Oversight Board et al., No. 08-861, U.S. Supreme Court, June 28.  
[https://scholar.google.com/scholar\\_case?case=6019616206791654633](https://scholar.google.com/scholar_case?case=6019616206791654633)

Freeman, Mike. 2018. "FTC Antitrust Case Against Qualcomm Puts Business Model Under Microscope." San Diego Union-Tribune. <https://www.sandiegouniontribune.com/2018/12/30/ftc-antitrust-case-against-qualcomm-puts-business-model-under-microscope/>

Freeman. 1978. In re Freeman. No. 77-598. US Court of Customs and Patent Appeals, March 30.  
<https://digital-law-online.info/cases/197PQ464.htm>

Frey, Carl Benedikt, and Michael Osborne. 2013. The Future of Employment: How Susceptible Are Jobs to Computerisation? Oxford Martin School. <https://www.oxfordmartin.ox.ac.uk/publications/the-future-of-employment>

Frieden, Jeffrey. 2006. Global Capitalism: Its Fall and Rise in the Twentieth Century. New York: W. W. Norton.

Friedman, Milton, and Anna Jacobson Schwartz. 1963. A Monetary History of the United States, 1867–1960. Princeton, NJ: Princeton University Press.

Froeb, Luke. 2003. The Role of Expert Economic Testimony in Antitrust Litigation. FTC, Bureau of Economics. [https://www.ftc.gov/sites/default/files/documents/public\\_statements/role-expert-economic-testimony-antitrust-litigation/031105froeb.pdf](https://www.ftc.gov/sites/default/files/documents/public_statements/role-expert-economic-testimony-antitrust-litigation/031105froeb.pdf)

FTC. 1966. In the Matter of National Dairy Products Corp. Order and Opinions. 70 F.T.C. 79. Docket No. 7018. [https://www.ftc.gov/sites/default/files/documents/commission\\_decision\\_volumes/volume-70/ftcd-vol70july-december1966pages79-223.pdf](https://www.ftc.gov/sites/default/files/documents/commission_decision_volumes/volume-70/ftcd-vol70july-december1966pages79-223.pdf)

FTC. 1975. In the Matter of Xerox Corp. Decision and Order. 86 F.T.C. 364. Washington, DC: FTC.

FTC. 1982. Statement of the Federal Trade Commission Concerning Horizontal Mergers. Washington, DC: FTC.

FTC. 1996. In the Matter of Dell Computer Corp. Consent Order. 121 F.T.C. 616. Docket No. C-3658. <https://www.ftc.gov/system/files/documents/cases/960617dellconsentorder.pdf>

FTC. 2000. "Privacy of Consumer Financial Information; Final Rule." F.R. 65 (101): 33646–33689.  
<https://www.govinfo.gov/content/pkg/FR-2000-05-24/html/00-12755.htm>

FTC. 2002. In the Matter of Rambus Inc. Administrative Complaint. Docket No. 9302. FTC Matter/File No. 011 0017. <https://www.ftc.gov/legal-library/browse/cases-proceedings/011-0017-rambus-inc-matter>

FTC. 2003. To Promote Innovation: The Proper Balance of Competition and Patent Law and Policy. Washington, DC: FTC.

FTC. 2007. "Statement of Federal Trade Commission Concerning Google/DoubleClick." FTC File No. 071-0170. Washington, D.C.: FTC.

FTC. 2009. "Self-Regulatory Principles for Online Behavioral Advertising." FTC Staff Report.  
<https://www.ftc.gov/sites/default/files/documents/reports/federal-trade-commission-staff-report-self-regulatory-principles-online-behavioral-advertising/p085400behavadreport.pdf>

FTC. 2012. In the Matter of Robert Bosch GmbH. Analysis to Aid Public Comment. Docket No. C-4377. FTC File No. 121-0081.  
<https://www.ftc.gov/sites/default/files/documents/cases/2013/04/121126boschanalysis.pdf>

FTC. 2013a. In the Matter of Motorola Mobility LLC and Google Inc. Decision and Order. Docket No. C-4410. FTC File No. 121-0120.  
<https://www.ftc.gov/sites/default/files/documents/cases/2013/01/130103googlemotorolado.pdf>

FTC. 2013b. In the Matter of Google Inc. Statement of the Federal Trade Commission Regarding Google's Search Practices. Matter No. 111-0163. FTC File No. 111-0163. <https://www.ftc.gov/legal-library/browse/cases-proceedings/public-statements/statement-federal-trade-commission-regarding-googles-search-practices-matter-google-inc>

FTC. 2017a. "Statement of Federal Trade Commission's Acting Director of the Bureau of Competition on the Agency's Review of Amazon.com, Inc.'s Acquisition of Whole Foods Market Inc." Washington, D.C.: FTC.

FTC. 2017b. Cross-Device Tracking: A Federal Trade Commission Staff Report.  
<https://www.ftc.gov/reports/cross-device-tracking-federal-trade-commission-staff-report-january-2017>

FTC. 2025. A Brief Overview of the Federal Trade Commission's Investigative, Law Enforcement, and Rulemaking Authority. <https://www.ftc.gov/about-ftc/mission/enforcement-authority>

FTC. n.d.a. The Antitrust Laws (Guide to Antitrust Laws). <https://www.ftc.gov/advice-guidance/competition-guidance/guide-antitrust-laws/antitrust-laws>

FTC. n.d.b. Clayton Act (Legal Library: Statutes). <https://www.ftc.gov/legal-library/browse/statutes/clayton-act>

FTSE Russell. 2025. Russell 3000 Index constituent list [Data set]. <https://www.ftserussell.com/index-series/index-tools/russell-3000-index>

Fu, Hui. 2025. Open-source LLMs surpass closed models in LMArena ranking. LinkedIn Post. [https://www.linkedin.com/posts/huifu\\_the-latest-lmarena-ranking-of-open-source-activity-7383889353653035008-JsBJ](https://www.linkedin.com/posts/huifu_the-latest-lmarena-ranking-of-open-source-activity-7383889353653035008-JsBJ)

Fues, Eric. 2007. Implications of eBay v. MercExchange. Patent World. <https://www.finnegan.com/en/insights/articles/implications-of-ebay-v-mercexchange.html>

Funk, Jeffrey. 2009. "The Emerging Value Network in the Mobile Phone Industry: The Case of Japan and Its Implications for the Rest of the World." *Telecommunications Policy* 33(1–2): 4–18.

Furman, Jason, and Peter Orszag. 2018. "A Firm-Level Perspective on the Role of Rents in the Rise in Inequality." In *Toward a Just Society: Joseph Stiglitz and Twenty-First Century Economics*, edited by Martin Guzman. New York: Columbia University Press, 89–110.

Furman, Jason. 2025. September 26. 92% of GDP Growth. X (formerly Twitter). <https://x.com/jasonfurman/status/1971995367202775284>

Gaffney, Jonathan. 2024. "Judicial Review Under the Administrative Procedure Act (APA)." Congressional Research Service Legal Sidebar LSB10558. <https://crsreports.congress.gov/product/pdf/LSB/LSB10558>.

Galetovic, Alexander, and Stephen Haber. 2017. "The Fallacies of Patent Holdup Theory." *Journal of Competition Law and Economics* 13(1): 1–44.

Galetovic, Alexander, Stephen Haber, and Lew Zaretski. 2024. "Cellular SEP Royalties and 5G: What Should Competition Policy Be?" In *5G and Beyond: Intellectual Property and Competition Policy in the Internet of Things*, edited by Jonathan M. Barnett and Seán M. O'Connor. Cambridge: Cambridge University Press, 53–78.

Galetovic, Alexander, Stephen Haber, and Ross Levine. 2018. "An Empirical Examination of Patent Holdup." *Journal of Competition Law and Economics* 11(3): 549-578.

Galetovic, Alexander. 2021. "Patents in the History of the Semiconductor Industry: The Ricardian Rent Hypothesis." In *The Battle Over Patents*, edited by Naomi Lamoreaux and Stephen Haber. Oxford: Oxford University Press: 27–68.

Gans, Joshua, David Hsu, and Scott Stern. 2002. "When Does Start-Up Innovation Spur the Gale of Creative Destruction?" *RAND Journal of Economics* 33(4): 571-586.

GAO. 1994. *General Agreement on Tariffs and Trade: Uruguay Round Final Act Should Produce Overall U.S. Economic Gains*. GAO. GGD-94-83B. <https://www.govinfo.gov/content/pkg/GAOREPORTS-GGD-94-83B/html/GAOREPORTS-GGD-94-83B.htm>

GAO. 1998. *Technology Transfer: Administration of the Bayh-Dole Act by Research Universities*. GAO/RCED-98-126. GAO. <https://www.gao.gov/assets/rced-98-126.pdf>

GAO. 2003. *Electronic Rulemaking: Efforts to Facilitate Public Participation Can Be Improved*. GAO-03-901. <https://www.govinfo.gov/app/details/GAOREPORTS-GAO-03-901>

Gardizy, Anissa, and Qianer Liu. 2025. "Google Convinces OpenAI to Use TPU Chips in Win Against Nvidia." *The Information*. <https://www.theinformation.com/articles/google-convinces-openai-use-tpu-chips-win-nvidia?rc=tuorsb>

Gardner, Bruce. 2002. *American Agriculture in the Twentieth Century*. Cambridge: Harvard University Press.

Gates, Bill. 1976. "An Open Letter to Hobbyists." *Homebrew Computer Club Newsletter*, no. 2.1 (January): 2. [https://commons.wikimedia.org/wiki/File:Bill\\_Gates\\_Letter\\_to\\_Hobbyists\\_ocr.pdf](https://commons.wikimedia.org/wiki/File:Bill_Gates_Letter_to_Hobbyists_ocr.pdf)

Gatford, Molly. 2023. "Conversational AI Market Report 2025-29." Juniper Research.  
<https://www.juniperresearch.com/research/telecoms-connectivity/communication-services/conversational-ai-research-report/>

Gavil, Andrew, William Kovacic, and Jonathan Baker. (2014). *Antitrust Law in Perspective*. 3rd ed. St. Paul: West Academic Publishing.

Gawer, Annabelle, and Michael Cusumano. 2002. *Platform Leadership*. Boston, MA: Harvard Business School Press.

Gehan, Ann. 2025. "Amazon Locks Down Against Google's AI Shopping Agents." *The Information*.  
<https://www.theinformation.com/articles/amazon-locks-google-ai-shopping-agents>

Gemini Team Google. 2023. Gemini. arXiv preprint arXiv:2312.11805. <https://arxiv.org/abs/2312.11805>

Gentry, Craig. 2009. "Fully Homomorphic Encryption Using Ideal Lattices." *Proceedings of the 41st Annual ACM Symposium on Theory of Computing*. New York: ACM, 169–78.

Geroski, Paul. 2000. "Models of Technology Diffusion." *Research Policy* 29(4-5): 603-625.

Gerstenzang, James. 1993. Senate Approves NAFTA on 61-38 Vote: Trade.  
<https://www.latimes.com/archives/la-xpm-1993-11-21-mn-59485-story.html>

Gerstle, Gary. 2022. *The Rise and Fall of the Neoliberal Order*. Oxford: Oxford University Press.

Gervase, Richard, Michael Renaud and Tinny Song. 2020. "Ninth Circuit Reverses FTC Win in FTC v. Qualcomm, Finding No Antitrust Violations from Qualcomm's Licensing of its Standard-Essential Patents." *Mintz Insights*. <https://www.mintz.com/insights-center/viewpoints/2231/2020-08-13-ninth-circuit-reverses-ftc-win-ftc-v-qualcomm-finding-no>

Gestalt Diagnostics. 2025. "Gestalt Demonstrates Advanced Interoperability and Standards Compliance at 2025 DICOM WG-26 Connectathon." Company news release.  
<https://www.gestaltdiagnostics.com/gestalt-demonstrates-advanced-interoperability-and-standards-compliance-at-2025-dicom-wg-26-connectathon>

Geva, Benjamin. 2020. *Electronic Payments: Guide on Legal and Regulatory Reforms and Best Practices for Developing Countries*. Osgoode Legal Studies Research Paper. <https://ssrn.com/abstract=3631155>

Gheller, Jonathan. 2015. "Introducing On This Day: A New Way to Look Back at Photos and Memories on Facebook." Facebook Newsroom. <https://about.fb.com/news/2015/03/introducing-on-this-day-a-new-way-to-look-back-at-photos-and-memories-on-facebook/>

Gibson, Bryan. 2015. Electronic Signatures: Review and Analysis. Kentucky Transportation Center Technical Assistance Report, University of Kentucky.

[https://uknowledge.uky.edu/cgi/viewcontent.cgi?article=1003&context=ktc\\_technicalassistancereports](https://uknowledge.uky.edu/cgi/viewcontent.cgi?article=1003&context=ktc_technicalassistancereports)

Gilbert, Ben and David Rosenthal. 2024. How ARM Became the World's Default Chip Architecture (with ARM CEO Rene Haas). ACQ2 by Acquired, podcast audio. <https://www.acquired.fm/episodes/how-arm-became-the-worlds-default-chip-architecture-with-arm-ceo-rene-haas>

Gilbert, Richard, Carl Shapiro, Louis Kaplow, and Robert Gertner. 1997. "Antitrust issues in the licensing of intellectual property: The nine no-no's meet the nineties." *Brookings Papers on Economic Activity. Microeconomics* 1997: 283-349.

Gilbert, Richard, and Steven Sunshine. 1995. Incorporating Dynamic Efficiency Concerns in Merger Analysis: The Use of Innovation Markets. *Antitrust Law Journal* 63(2): 569-601.

Gilhousen, Klein, Irwin Jacobs, Roberto Padovani, Andrew Viterbi, Lindsay Weaver Jr., and Charles Wheatley III. 1991. "On the Capacity of a Cellular CDMA System." *IEEE Transactions on Vehicular Technology* 40(2): 303-312.

Gilhousen, Klein, Irwin Jacobs, Andrew Viterbi, Charles Wheatley III, Roberto Padovani, and Lindsay Weaver Jr. 1992. System and Method for Generating Signal Waveforms in a CDMA Cellular Telephone System (U.S. Patent No. 5,103,459). USPTO. [https://ptacts.uspto.gov/ptacts/public-informations/petitions/1471034/download-documents?artifactId=u1dPKxpYy6ShiP9yeA\\_0OfKLmaBu-U6jVd22CDZ9LrOyCkVzpKhtdYw](https://ptacts.uspto.gov/ptacts/public-informations/petitions/1471034/download-documents?artifactId=u1dPKxpYy6ShiP9yeA_0OfKLmaBu-U6jVd22CDZ9LrOyCkVzpKhtdYw)

Gilhousen, Klein, Irwin Mark Jacobs, Paul Padovani, Andrew Viterbi, Louis Weaver, and Charles Wheatley. 1987. Spread spectrum wireless PBX (U.S. Patent No. 4,672,658). USPTO. <https://patents.justia.com/patent/4672658>

Gillespie, Tarleton. 2018. *Custodians of the Internet*. New Haven, CT: Yale University Press.

Gochberg, Will and Victor Menaldo. 2022. To rent or not to rent? Mechanics, causes and consequences of Ricardian and Quasi-rents in the oil industry. *Resources Policy* 78(September): 102826.

Golden, Jeffrey. 2012. "The Single Agreement Concept in Derivatives Law." *Capital Markets Law Journal* 7(1): 1–23.

Goldfarb, Charles. 2007. *Telecommunications Act: Competition, Innovation, and Reform*. CRS Report RL33034. <https://www.everycrsreport.com/reports/RL33034.html>

Goldin, Claudia, and Lawrence Katz. 2008. *The Race between Education and Technology*. Cambridge, MA: Harvard University Press.

Goldstone, Jack. 2002. "Efflorescences and Economic Growth in World History: Rethinking the 'Rise of the West' and the Industrial Revolution." *Journal of World History* 13(2): 323–389.

Gomes, Ben. 2017. Our latest quality improvements in Search. *The Keyword* (Google). <https://blog.google/products/search/our-latest-quality-improvements-search/>

Gonzalez. 2023. *Gonzalez v. Google LLC*. No. 21-1333. US Supreme Court, May 18. <https://supreme.justia.com/cases/federal/us/598/21-1333/>

Good, I.J. (Irving John). 1965. "Speculations Concerning the First Ultraintelligent Machine." *Advances in Computers* 6: 31–88.

Goode, Bryan. 2025. From systems of record to systems of action: Dynamics 365, agentic business applications for the frontier. *Microsoft Dynamics 365 Blog*. <https://www.microsoft.com/en-us/dynamics-365/blog/business-leader/2025/10/21/from-systems-of-record-to-systems-of-action-dynamics-365-agentic-business-applications-for-the-frontier/>

Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. 2016. *Deep Learning*. Cambridge, MA: MIT Press.

Google Cloud. 2024. *MLOps: Continuous Delivery and Automation Pipelines in Machine Learning*. <https://cloud.google.com/architecture/mlops-continuous-delivery-and-automation-pipelines-in-machine-learning>

Google Cloud. 2025a. Overview of getting inferences on Vertex AI. <https://docs.cloud.google.com/vertex-ai/docs/predictions/overview>

Google Cloud. 2025b. AI and Machine Learning APIs Overview. Google Cloud Documentation. <https://cloud.google.com/products/ai>

Google. 2025b. A New Era of Intelligence with Gemini 3. The Keyword. <https://blog.google/products/gemini/gemini-3/>

Google. n.d. "How the Google Ads Auction Works." Google Ads Help Center. <https://support.google.com/google-ads/answer/6366577?hl=en>.

Gordon, Robert. 2000. "Does the 'New Economy' Measure Up to the Great Inventions of the Past?" *Journal of Economic Perspectives* 14(4): 49-74.

Gordon, Robert. 2016. *The Rise and Fall of American Growth: The U.S. Standard of Living since the Civil War*. Princeton, NJ: Princeton University Press.

Goswami, Rohan, and Kif Leswing. 2024. Intel stock jumps on plan to turn foundry business into subsidiary and allow for outside funding. CNBC. <https://www.cnbc.com/2024/09/16/intel-stock-jumps-on-plan-to-turn-foundry-business-into-subsidiary.html>

Gottschalk. 1972. *Gottschalk v. Benson and Tabb*. No. 71-485. US Supreme Court, November 20. <https://supreme.justia.com/cases/federal/us/409/63/>

Gowa, Joanne. 1983. *Closing the Gold Window*. Ithaca: Cornell University Press.

Grace, Katja, et al. 2024. Thousands of AI Authors on the Future of AI (2023 Expert Survey on Progress in AI). AI Impacts. [https://aiimpacts.org/wp-content/uploads/2023/04/Thousands\\_of\\_AI\\_authors\\_on\\_the\\_future\\_of\\_AI.pdf](https://aiimpacts.org/wp-content/uploads/2023/04/Thousands_of_AI_authors_on_the_future_of_AI.pdf)

Graetz, Georg, and Guy Michaels. 2018. "Robots at Work." *The Review of Economics and Statistics* 100(5): 753–768.

Graham, Stuart, and Saurabh Vishnubhakat. 2013. "Of Smart Phone Wars and Software Patents." *Journal of Economic Perspectives* 27(1): 67–86.

Gravelle, Jane. 2014. Bonus Depreciation: Economic and Budgetary Issues. CRS Report R43432. [https://www.congress.gov/crs\\_external\\_products/R/PDF/R43432/R43432.7.pdf](https://www.congress.gov/crs_external_products/R/PDF/R43432/R43432.7.pdf)

Grbovic, Mihajlo, Nemanja Djuric, Vladan Radosavljevic, Fabrizio Silvestri, and Ricardo Baeza-Yates. 2016. "Scalable Semantic Matching of Queries to Ads in Sponsored Search Advertising." Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '16). <https://arxiv.org/pdf/1607.01869>.

Green, Alastair, Andrea Del Miglio, Humayun Tai, Jesse Noffsinger, Mark Patel, Marc Sorel, and Pankaj Sachdeva. 2025. What Is a Data Center? McKinsey Explainer, McKinsey & Company. [https://www.mckinsey.com/featured-insights/mckinsey-explainers/what-is-a-data-center#](https://www.mckinsey.com/featured-insights/mckinsey-explainers/what-is-a-data-center#/)

Greenstein, Shane. 2015. How the Internet Became Commercial. Princeton, NJ: Princeton University Press.

Griffith, Erin. 2025. "In A.I. Boom, Venture Capital Firms Are Raising Loads More Money." The New York Times. <https://www.nytimes.com/2025/12/15/technology/ai-venture-capital-big-funds.html>

Griliches, Zvi. 1992. "The search for R&D spillovers." *Scandinavian Journal of Economics* 94 (Supplement): S29–S47.

Griliches, Zvi. 1957. "Hybrid Corn: An Exploration in the Economics of Technological Change." *Econometrica* 25(4): 501–522.

Grindley, Peter, and David Teece. 1997. "Managing Intellectual Capital: Licensing and Cross-Licensing in Semiconductors and Electronics." *California Management Review* 39(2): 8–41.

Grinnell. 1966. *United States v. Grinnell Corp.*, No. 73. US Supreme Court, June 13. <https://supreme.justia.com/cases/federal/us/384/563/>

Guenther, Gary. 2022. Research Tax Credit. CRS Report RL31181. <https://www.congress.gov/crs-product/RL31181>

Gundlach, Hans, Jayson Lynch, Matthias Mertens, and Neil Thompson. 2025. The price of progress: Algorithmic efficiency and the falling cost of AI inference arXiv preprint: arXiv2511.23455. <https://arxiv.org/html/2511.23455v1>

Guo, Daya, et al. 2025. "DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Large-Scale Reinforcement Learning." DeepSeek-AI. arXiv preprint arXiv:2501.12948.  
<https://arxiv.org/abs/2501.12948>

Gupta, Kirti. 2015. "Technology Standards and Competition in the Mobile Wireless Industry." *George Mason Law Review* 22: 865–893.

Gupta, Nikhil. 2025. MLOps in the Cloud-Native Era — Scaling AI/ML Workloads with Kubernetes and Serverless Architectures. Cloud Native Now.  
<https://cloudnativenow.com/topics/cloudnativedevelopment/kubernetes/mlops-in-the-cloud-native-era-scaling-ai-ml-workloads-with-kubernetes-and-serverless-architectures/>

Gurman, Mark. 2024. Gurman, Mark. 2024. "Apple to Power AI Servers With In-House Server Chips This Year." Bloomberg. <https://www.bloomberg.com/news/articles/2024-05-09/apple-to-power-ai-features-with-in-house-server-chips-this-year>

Gutiérrez, Germán, and Thomas Philippon. 2018. "Ownership, Concentration, and Investment." *AEA Papers and Proceedings* 108 (May): 432–37.

H2O.ai. 2025. Use AI to Assist Trading Strategies. <https://h2o.ai/solutions/use-case/trading-strategies/>

Haber, Stephen and Naomi Lamoreaux. 2021. "Introduction." In *The Battle Over Patents: Historical Perspectives on Current Debates*, eds., Stephen Haber and Naomi Lamoreaux. Oxford: Oxford University Press: 1-26.

Haber, Stephen. 2016. "Patents and the Wealth of Nations." *George Mason Law Review* 23(4): 811–835.

Haber, Stephen and Naomi Lamoreaux. 2021. "Introduction." In *Patents, Inventors, and Politics: Historical Perspectives on Current Debates*, eds., Stephen Haber and Naomi Lamoreaux. Oxford: Oxford University Press, 1-26.

Haggard, Stephan, and Robert Kaufman. 1992. *The Politics of Economic Adjustment: International Constraints, Distributive Conflicts, and the State*. Princeton, NJ: Princeton University Press.

Hall, Bronwyn H. 1996. "The private and social returns to research and development." In *Technology, R&D, and the Economy*, edited by Bruce Smith and Carl Barfield. Washington, DC: Brookings Institution, 140-183.

Hall, Bronwyn, and Rosemarie Ham Ziedonis. 2001. *The Patent Paradox Revisited*. *RAND Journal of Economics* 32(1): 101-128.

Hall, Bronwyn. 2002. "The Financing of Research and Development." *Oxford Review of Economic Policy* 18(1): 35–51.

Hall, Robert. 1988. "The Relation between Price and Marginal Cost in U.S. Industry." *Journal of Political Economy* 96(5): 921-947.

Hall, Robert. 2018. "New Evidence on the Markup of Prices Over Marginal Costs and the Role of Mega-Firms in the U.S. Economy." NBER Working Paper No. 24574. Cambridge, MA: National Bureau of Economic Research.

Hamid, Jai. 2025. Visa, Mastercard build payment rails as AI agents ready to shop and pay autonomously. *Cryptopolitan*. <https://www.cryptopolitan.com/visa-mastercard-build-payments-ai-agents/>

Hammond, George, Tabby Kinder, and Antoine Gara. 2025. "OpenAI Forges \$12bn Contract with CoreWeave." *Financial Times*. <https://www.ft.com/content/4b52fdbb-ca8e-4208-bb99-f1e7f9313863>

Handler, Milton. 1957. "Annual review of antitrust developments." *Record of the Association of the Bar of the City of New York* 12: 185–202.

Handshake. 2025. *State of the Graduate: Class of 2025*. Handshake. <https://joinhandshake.com/themes/handshake/dist/assets/downloads/network-trends/class-of-2025.pdf>

Harbridge-Yong, Laurel. 2015. *Is Bipartisanship Dead? Policy Agreement and Agenda-Setting in the House of Representatives*. New York: Cambridge University Press

Harris, Barry, and Joseph Simons. 1989. "Focusing Market Definition." *Research in Law and Economics* 12: 207–26.

Harvey, David. 2005. *A Brief History of Neoliberalism*. Oxford: Oxford University Press.

Haskel, Jonathan, and Stian Westlake. 2017. *Capitalism without Capital: The Rise of the Intangible Economy*. Princeton, NJ: Princeton University Press.

Hauck, Ronny. 2015. "FTC v. Google." In *Competition on the Internet*, edited by Gintarė Surblytė. Berlin: Springer, 83-108.

Hausman, Jerry, and J. Gregory Sidak. 1999. "Did Mandatory Unbundling Achieve Its Purpose? Empirical Evidence from Five Countries." *Federal Communications Law Journal* 52(1): 3–81.

Hausman, Jerry. 1997. "Valuing the Effect of Regulation on New Services in Telecommunications." *Brookings Papers on Economic Activity: Microeconomics 1997*: 1–38.

Hausmann, Ricardo, and Dani Rodrik. 2003. "Economic Development as Self-Discovery." *Journal of Development Economics* 72(2): 603-633.

Hawley, Ellis. 1966. *The New Deal and the problem of monopoly: A study in economic ambivalence*. Princeton, NJ: Princeton University Press.

Hayes, Connor. 2017. "More Ways to Share With the Facebook Camera." Facebook Newsroom (Meta). <https://about.fb.com/news/2017/03/more-ways-to-share-with-the-facebook-camera/>

Hazlett, Thomas W. 1998. "Economic and political consequences of the 1996 Telecommunications Act." *Journal of Regulation and Economic Change* 11 (1): 21–42.

HealthIT.gov. 2025. "Health Information Technology Standards and Certification." Office of the National Coordinator for Health Information Technology. <https://www.healthit.gov/>

Heaven, Will. 2025. "The great AI hype correction of 2025." MIT Technology Review. <https://www.technologyreview.com/2025/12/15/1129174/the-great-ai-hype-correction-of-2025>

Helpman, Elhanan, and Manuel Trajtenberg. 1996. *Diffusion of General Purpose Technologies*. NBER Working Paper No. 5773. Cambridge, MA: National Bureau of Economic Research.

Henderson, Rebecca, Adam Jaffe, and Manuel Trajtenberg. 1998. "Universities as a Source of Commercial Technology: A Detailed Analysis of University Patenting, 1965–1988." *The Review of Economics and Statistics* 80(1): 119–127.

Hendrycks, Dan, Norman Mu, Ekin D. Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. 2021. "The Many Faces of Robustness: A Critical Analysis of Out-of-Distribution Generalization." arXiv preprint arXiv:2006.16241. <https://arxiv.org/abs/2006.16241>

Hennessy, John and David Patterson. 2019. Computer Architecture: A Quantitative Approach. 6th ed. San Francisco: Morgan Kaufmann.

Henry, Alan. 2011. "Facebook Is Tracking Your Every Move on the Web; Here's How to Stop It." Lifehacker. <https://lifehacker.com/facebook-is-tracking-your-every-move-on-the-web-heres-5843969>.

Henshall, William. 2023. "E.U.'s AI Regulation Could Be Softened After Pushback from Biggest Members." Time. <https://time.com/6338602/eu-ai-regulation-foundation-models/>

Herrick. 2018. Herrick v. Grindr, LLC. No. 1:17-cv-01747. US District Court, S.D. New York, March 27. <https://casetext.com/case/herrick-v-grindr-llc-1>

Hestness, Joel, Sharan Narang, Newsha Ardalani, Gregory Diamos, Heewoo Jun, Hassan Kianinejad, Md. Mostofa Ali Patwary, Yang Yang, Yanqi Zhou. 2017. "Deep Learning Scaling Is Predictable, Empirically." arXiv preprint arXiv:1712.00409. <https://arxiv.org/abs/1712.00409>

Hiller, Jennifer. 2025. "AI Data Centers, Desperate for Electricity, Are Building Their Own Power Plants." The Wall Street Journal. <https://www.wsj.com/business/energy-oil/ai-data-centers-desperate-for-electricity-are-building-their-own-power-plants-291f5c81>

hiQ Labs. 2022. hiQ Labs, Inc. v. LinkedIn Corp. No. 17-16783. US Court of Appeals, Ninth Circuit, April 18. <https://cdn.ca9.uscourts.gov/datastore/opinions/2022/04/18/17-16783.pdf>

HJC. 1980. Patent and Trademark Law Amendments Act of 1980. House Report No. 96-1307. 96th Congress, 2nd Session. Washington, DC: U.S. Government Printing Office.

HJC. 2020. Investigation of Competition in Digital Markets: Majority Staff Report and Recommendations. Washington, DC: U.S. Government Printing Office.

HL7 International. 2025. "FHIR Overview." <https://www.hl7.org/fhir/>

Hoffman, Sarah. 2025. "Generative AI and the Future of Entry-Level Jobs." AlphaSense. <https://www.alpha-sense.com/blog/trends/generative-ai-entry-level-jobs/>

Hoffmann, Jordan, et al. 2022. "Training Compute-Optimal Large Language Models." arXiv preprint arXiv:2203.15556. <https://arxiv.org/abs/2203.15556>

Hofstadter, Richard 1955. *The Age of Reform: From Bryan to F.D.R.* New York: Alfred A. Knopf.

Holma, Harri, and Antti Toskala. 2012. *LTE for UMTS: Evolution to LTE-Advanced*. Chichester: John Wiley and Sons.

Holmes, Aaron, and Aaron Tilley. 2025. "Developers Eye Apple's Models for 'Invisible' AI Features." *The Information*. <https://www.theinformation.com/articles/developers-eye-apples-models-invisible-ai-features>

Holmes, Aaron, and Sri Muppidi. 2025. "OpenAI Takes a Page from Palantir, Doubles Down on Consulting Services." *The Information*. [www.theinformation.com/articles/openai-takes-page-palantir-doubles-consulting-services](http://www.theinformation.com/articles/openai-takes-page-palantir-doubles-consulting-services)

Holmes, Thomas. 2001. "Bar Codes Lead to Frequent Deliveries and Superstores." *The RAND Journal of Economics* 32(4): 708–725.

Hounshell, David. 1984. *From the American System to Mass Production, 1800-1932*. Baltimore: Johns Hopkins University Press.

Hovenkamp, Herbert. 2009. *The Neal Report and the Crisis in Antitrust*. University of Iowa Legal Studies Research Paper No. 09-09.

Hovenkamp, Herbert. 2018. "The Rule of Reason." *Florida Law Review* 70(1): 47–118.

Hovenkamp, Herbert. 2005. *The Antitrust Enterprise: Principle and Execution*. Cambridge, MA: Harvard University Press.

Hoyt, Bradley, Glenn Hinton, Andrew Glew, and Subramanian Natarajan. 1996. Branch target buffer for dynamically predicting branch instruction outcomes using a predicted branch history (U.S. Patent No. 5,584,001). USPTO. <https://patents.google.com/patent/US5584001A>

Hu, Krystal, Kenrick Cai, and Stephen Nellis. 2025. "Google works to erode Nvidia's software advantage with Meta's help." *Reuters*. <https://www.reuters.com/business/google-works-erode-nvidias-software-advantage-with-metas-help-2025-12-17/>

Huang, David, Avidan Shah, Alexandre Araujo, David Wagner, and Chawin Sitawarin. 2025. "Stronger Universal and Transferable Attacks by Suppressing Refusals." NAACL 2025 (Long Papers).  
<https://aclanthology.org/2025.naacl-long.302.pdf>.

Huang, Raffaele, and Liza Lin. 2025. "Chinese AI Companies Dodge U.S. Chip Curbs by Flying Suitcases of Hard Drives Abroad." Wall Street Journal. <https://www.wsj.com/tech/china-ai-chip-curb-suitcases-7c47dab1>

Hubbard, Patrick. 2018. "The Artificial Intelligence 'Black Box' and the Failure of Intent and Causation." Harvard Journal of Law and Technology 31(2): 890-938.

Huescas, Jose Carlos. 2024. Unleashing the Power of AI: MLPerf Benchmarking Outcomes Show. Lenovo Press. <https://lenovopress.lenovo.com/lp1969-unleashing-the-power-of-ai-at-the-edge-mlperf-benchmarking-outcomes>

Hugging Face. 2025a. Inference Endpoints. Hugging Face Documentation.  
<https://huggingface.co/docs/inference-endpoints>

Hugging Face. 2025b. deepseek-ai/DeepSeek-R1 (model card; accessed Nov. 2025).  
<https://huggingface.co/deepseek-ai/DeepSeek-R1>

Hugging Face. 2025c. meta-llama/Llama-3.1-8B-Instruct (model card; accessed Nov. 2025).  
<https://huggingface.co/meta-llama/Llama-3.1-8B-Instruct>

Hughes, Thomas. 1983. Networks of power: Electrification in Western society, 1880–1930. Baltimore: Johns Hopkins University Press.

Humphrey’s Executor 1935. Humphrey's Executor v. United States, No. 667, U.S. Supreme Court, May 27. <https://supreme.justia.com/cases/federal/us/295/602/>

Hunter, Matt. 2016. "Instagram launches 'Stories,' a product to take on Snapchat." CNBC.  
<https://www.cnbc.com/2016/08/02/instagram-launches-stories-a-product-to-take-on-snapchat.html>

Hussein, Ahmed, Mohamed Medhat Gaber, Eyad Elyan, and Chrisina Jayne. 2017. "Imitation learning: A survey of learning methods." ACM Computing Surveys 50 (2): 21.

Hwang, Tim. 2020. Subprime Attention Crisis. New York: FSG Originals.

IBM. 1969. United States v. International Business Machines Corporation. 1969. Civil Action No. 69 Civ. 200 (S.D.N.Y.). <https://www.justice.gov/atr/case-document/united-states-memorandum-1969-case>

IBM. 1982. United States v. International Business Machines Corp., No. 69 Civ. 200 (S.D.N.Y.), Stipulation of Dismissal, Jan. 8. <https://law.justia.com/cases/federal/district-courts/FSupp/539/473/2151775>

IBM Research. 2025. Agent Communication Protocol (ACP). <https://research.ibm.com/projects/agent-communication-protocol>

IDC. 2023. Worldwide Global DataSphere Forecast, 2023–2027: It’s an AI Data World. IDC. <https://www.marketresearch.com/IDC-v2477/Worldwide-IDC-Global-DataSphere-Forecast-33986214/>

Iden, Virginius Gilmore. 1913. The Federal Reserve Act of 1913: History and Digest. FRASER. <https://fraser.stlouisfed.org/title/federal-reserve-act-1913-962/fulltext>

IEA. 2025. Energy and AI. Paris: IEA. <https://www.iea.org/reports/energy-and-ai>

IETF. 2022. "Robots Exclusion Protocol." RFC 9309. <https://datatracker.ietf.org/doc/html/rfc9309>

IFR. 2020. "Record 2.7 Million Robots Work in Factories Around the Globe." IFR Press Release. <https://ifr.org/ifr-press-releases/news/record-2.7-million-robots-work-in-factories-around-the-globe>

Illinois Tool Works Inc. 2006. Illinois Tool Works Inc. et al. v. Independent Ink, Inc., No. 04-1329, U.S. Supreme Court, March 1. [https://scholar.google.com/scholar\\_case?case=12062423076831474839](https://scholar.google.com/scholar_case?case=12062423076831474839)

InfoQ. 2025. "OpenAI and Anthropic Donate AGENTS.md and Model Context Protocol to New Agentic AI Foundation." InfoQ. <https://www.infoq.com/news/2025/12/agentic-ai-foundation/>

Innovation Alliance. 2022. Public Submission on SEPs. <https://innovationalliance.net/wp-content/uploads/2022/02/2022-0204-Innovation-Alliance-SEP-Public-Submission.pdf>.

InsideHPC. 2024. Nvidia Announces Installations of 9 Grace Hopper-Powered Supercomputers. InsideHPC. <https://insidehpc.com/2024/05/nvidia-announces-installations-of-9-grace-hopper-powered-supercomputers/>

Insilico Medicine. 2023. "ReLEHF: Reinforcement Learning with Expert Human Feedback to Advance Pharmaceutical Generative AI." Insilico Medicine Blog. <https://insilico.com/blog/relehf>

Intel Corp. 2024. Form 10-K: Annual report for the fiscal year ended December 30, 2023. U.S. SEC.  
<https://www.intc.com/filings-reports/annual-reports>

Ip, Greg. 2025. "How Global Trade Could Survive Trump's Tariffs." The Wall Street Journal.  
[https://www.wsj.com/economy/trade/how-global-trade-could-survive-trumps-tariffs-24d74d08?mod=Searchresults\\_pos1&page=1](https://www.wsj.com/economy/trade/how-global-trade-could-survive-trumps-tariffs-24d74d08?mod=Searchresults_pos1&page=1).

Irwin, Douglas. 1996. "The Smoot-Hawley Tariff: A Quantitative Assessment." NBER Working Paper No. 5509. Cambridge, MA: National Bureau of Economic Research.

Irwin, Douglas. 2013. "The Nixon Shock after Forty Years." *World Trade Review* 12(1): 29–56.

Irwin, Douglas, and Peter Klenow. 1996. "High-Tech R&D Subsidies: Estimating the Effects of Sematech." *Journal of International Economics* 40(3–4): 323–344.

ISA. 2025. "ISA95, Enterprise-Control System Integration." <https://www.isa.org/standards-and-publications/isa-standards/isa-standards-committees/isa95>

Isaac. 2025. How much does it cost to manufacture a mobile phone worth more than 1000 euros? Android Guias. <https://en.androidguias.com/How-much-does-it-cost-to-manufacture-a-mobile-phone-worth-more-than-1000-euros%3F/>

Isaacson, Walter 2014. *The Innovators: How a Group of Hackers, Geniuses, and Geeks Created the Digital Revolution*. New York: Simon and Schuster.

ISDA. 2002. *ISDA Master Agreement*. New York: ISDA.

ISO/IEC JTC 1/SC 42. 2024. Artificial Intelligence. ISO/IEC JTC 1/SC 42. <https://jtc1info.org/sd-2-history/jtc1-subcommittees/sc-42>

ITC. 1993 Potential Impact on the U.S. Economy and Selected Industries of the North American Free-Trade Agreement. ITC Publication 2596. <https://www.usitc.gov/publications/332/pub2596.pdf>

ITC. 1978. Annual Report 1978. ITC. <https://www.usitc.gov/publications/annualreport/pub982.pdf>

ITC. 1994. The Uruguay Round of Multilateral Trade Negotiations: Summary of the Final Act. ITC Publication 2790. <https://www.usitc.gov/publications/332/pub2790.pdf>

Ivanov, Dmitry. 2021. "Supply Chain Resilience: Modelling, Management, and Control." *International Journal of Production Research* 59(8): 2535–2551.

Jaffe, Adam, Manuel Trajtenberg, and Rebecca Henderson. 1993. "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations." *Quarterly Journal of Economics* 108(3): 577–598.

Janakiram, MSV. 2025. "Nvidia Wants To Build A Planet-Scale AI Factory With DGX Cloud Lepton." *Forbes*. <https://www.forbes.com/sites/janakirammsv/2025/06/13/nvidia-wants-to-build-a-planet-scale-ai-factory-with-dgx-cloud-lepton/>

Janc, Victor, and Steven Jasper. 1988. Digital Global Positioning System Receiver (US Patent No. 4,785,463). USPTO. <https://patents.google.com/patent/US4785463A/en>

Jefferson Parish Hospital District No. 2. 1984. *Jefferson Parish Hospital District No. 2 et al. v. Hyde*. No. 82-1031. US Supreme Court, March 27. <https://supreme.justia.com/cases/federal/us/466/2>

Jiang, Albert, et al. 2023. "Mistral 7B." arXiv preprint arXiv:2310.06825. <https://arxiv.org/abs/2310.06825>

Jie, Yang. 2025. Intel, SoftBank Discussed a Sale. Now, a Cash Infusion Will Accelerate American AI. *The Wall Street Journal*. <https://www.wsj.com/tech/intel-gets-2-billion-infusion-from-softbank-accelerating-ai-push-in-u-s-2129b3c4>

Jin, Berber, Peter Rudegeair, and Matt Wirz. 2025. "Musk Allies to Raise Up to \$12 Billion for xAI Chips as Startup Burns Through Cash." *Wall Street Journal*. <https://www.wsj.com/tech/ai/elon-musk-x-ai-funding-feecedel>

John Mac Ghlionn. 2025. "A blue-collar bloodbath is inevitable." *The Hill*. <https://thehill.com/opinion/technology/5445562-ai-threat-to-skilled-trades/>

Johnson, Jennifer and George Hay. 2025. "Europe will struggle to get Big Tech off its cloud." *Reuters Breakingviews*. <https://www.reuters.com/commentary/breakingviews/europe-will-struggle-get-big-tech-off-its-cloud-2025-06-26/>

Jon Peddie Research. 2024. Summary report on the worldwide total GPU market. JPR.  
<https://www.jonpeddie.com/store/summary-report-on-the-worldwide-total-gpu-market/>

Jonas, Adam, et al. 2024. Humanoids: Investment Implications of Embodied AI. Morgan Stanley Blue Paper. <https://www.futuremanagementgroup.com/wp-content/uploads/240626-Humanoid-Robots-Morgan-Stanley.pdf>

Jonas, Adam, William Tackett, Sheng Zhong, Daniela Haigian, and Elizabeth Tso. 2025. The Humanoid 100: Mapping the Humanoid Robot Value Chain. Morgan Stanley.

Joose, Tess. 2022. "December 1945: The ENIAC Computer Runs Its First, Top-Secret Program." American Physical Society: APS News/This Month in Physics History.  
<https://www.aps.org/apsnews/2022/11/eniac-first-top-secret-program>

Jorgenson, Dale. 2001. "Information Technology and the U.S. Economy." *American Economic Review* 91(1): 1-32.

Jorgenson, Dale, and Kevin Stiroh. 2000. "Raising the Speed Limit: U.S. Economic Growth in the Information Age." *Brookings Papers on Economic Activity* 2000(1): 125-211.

Jouppi, N. P., et al. 2017. In-datacenter performance analysis of a tensor processing unit. In *Proceedings of the 44th annual international symposium on computer architecture*: 1-12.

Jovanovic, Boyan, and Peter Rousseau. 2005. "General Purpose Technologies." In *Handbook of Economic Growth, Volume 1, Part B*. Edited by Philippe Aghion and Steven Durlauf. Amsterdam: Elsevier, 1181–1224.

Jumper, John, et al. 2021. Highly accurate protein structure prediction with AlphaFold. *Nature* (5967873): 583–589.

Junarkar, Sandeep. 2003. "HP seals P&G outsourcing deal." *ZDNet*. <https://www.zdnet.com/article/hp-seals-p-g-outsourcing-deal/>

Kadrey et al. 2023. *Kadrey et al. v. Meta Platforms Inc.*, No. 3:23-cv-03417. US District Court for the Northern District of California, July 7. <https://www.courtlistener.com/docket/67569326/kadrey-v-meta-platforms-inc>

Kagan, Julia. 2024. Lump of Labor Fallacy: Definition and How It Works. Investopedia.  
<https://www.investopedia.com/terms/l/lump-of-labour-fallacy.asp>

Kahn, Jeremy. 2025. "OpenAI Builds Apps into ChatGPT, in a Bold Bid to Make AI the Universal Interface." Fortune. <https://fortune.com/2025/10/07/openais-chatgpt-apps-universal-interface-platform-shift-amd-deal-eye-on-ai>

Kang, Cecilia, Cade Metz, and Stuart Thompson. 2024. "Four Takeaways on the Race to Amass Data for A.I." New York Times. <https://www.nytimes.com/2024/04/06/technology/ai-data-tech-takeaways.html>.

Kaplan, Jared, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling Laws for Neural Language Models. arXiv preprint: arXiv:2001.08361 <https://doi.org/10.48550/arXiv.2001.08361>

Kaplow, Louis. 1987. "Antitrust, Law and Economics, and the Courts." *Law and Contemporary Problems* 50(4): 181–216.

Kaplow, Louis. 2010. "Why (Ever) Define Markets?" *Harvard Law Review* 124(2): 437–517.

Karlsson, Kasper. 2023. "Arm's Smartphone Monopoly: IPO and Beyond." *Quartr Insights*.  
<https://quartr.com/insights/company-research/arm-s-smartphone-monopoly-setting-the-stage-for-its-ipo>

Katsurashima, Wataru, Sid Nag, and Evan Zeng. 2024. Competitive Behaviors of Hyperscale Providers in the AI Infrastructure Market. Gartner Research. <https://www.gartner.com/document/competitive-behaviors-hyperscale-providers-ai-infrastructure-market>

Katz, Michael, and Jonathan Sallet. 2018. Multisided Platforms and Antitrust Enforcement. *Yale Law Journal* 127(7): 2142-2175.

Katz, Michael, and Carl Shapiro. 1994. "Systems Competition and Network Effects." *Journal of Economic Perspectives* 8(2): 93–115.

Katz, Michael, and Carl Shapiro. 2003. "Critical Loss: Let's Tell the Whole Story." *Antitrust* (Spring): 49–56.

Katz, Michael, and Howard Shelanski. 2005. "'Schumpeterian' Competition and Antitrust Policy in High-Tech Markets." *Competition* 14(2): 47-91.

Katz, Michael, and Howard Shelanski. 2007. "Schumpeterianism and Antitrust." *Villanova Law Review* 52(3): 513–562.

Katz, Richard. 1998. *Japan: The System That Soured: The Rise and Fall of the Japanese Economic Miracle*. Armonk, NY: M.E. Sharpe.

Katz, Richard. 2025. "A 21<sup>st</sup> Century Industrial Policy: Diffusing Technology from Leading Companies to Laggards." *Japan Economy Watch*, Dec. 4th 2025 Entry. <https://richardkatz.substack.com/p/a-21st-century-industrial-policy>

Katzen, Sally. 2011. "OIRA at Thirty." *Administrative Law Review* 63 (Special Edition): 103–114.

Kauper, Thomas. 1981. "The Burger Court and Antitrust Law: An Unhappy Marriage." *Antitrust Law Journal* 50(4): 783–801

Kay, Grace. 2025. Internal xAI org chart: Names, positions, team size. *Business Insider*.  
<https://www.businessinsider.com/xai-org-chart-employees-elon-musk-direct-reports-2025-3>

Kaysen, Carl. 1956. *United States v. United Shoe Machinery Corporation*. Cambridge, MA: Harvard University Press.

Kearns, Michael, and Yuri Nevmyvaka. 2013. "Machine Learning for Market Microstructure and High Frequency Trading." In *High Frequency Trading: New Realities for Traders, Markets and Regulators*, edited by Irene Aldridge. London: Risk Books, 67–93.

Kehayias, Alex. 2023. *SaaS Revenue per Employee Benchmarks for B2B Companies*. SaaS Capital.  
<https://notes.alexkehayias.com/revenue-per-employee-benchmarks/>

Kennedy, Paul. 2017. *Vampire Capitalism* London: Palgrave Macmillan.

Kenney, Martin, and John Zysman. 2016. "The Rise of the Platform Economy." *Issues in Science and Technology* 32(3).

Kessler, Andy. 2022. Putin's vertical empire: Power and money in Russia. *The Wall Street Journal*.  
<https://www.wsj.com/articles/putins-vertical-empire-fall-ukraine-russia-business-kleptocrat-oligarchs-value-companies-11646581296>

Kessler, Andy. 2025. "Netscape's Lessons for AI Mania." *The Wall Street Journal*.  
<https://www.wsj.com/opinion/netscapes-lessons-for-ai-mania-innovation-technology-a132d8ae>

Khan, Lina M. 2019. "The Separation of Platforms and Commerce." *Columbia Law Review* 119(7): 973–1095. *Columbia Law Review*.

Khan, Lina. 2017. "Amazon's Antitrust Paradox." *Yale Law Journal* 126(3): 710–805.

Khronos Group. 2025. What is SYCL? C++ programming for heterogeneous parallel computing. Khronos Group. <https://www.khronos.org/sycl/>

Kiela, Douwe, et al. 2021. "Dynabench: Rethinking Benchmarking in NLP." In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 4110–4124.

Kilpatrick, Jim, Berckman, Lindsey, Faver, Alan D., Hardin, Kate, and Sloane, Matt. 2025. Supply chain resilience. *Deloitte Insights*. <https://www2.deloitte.com/us/en/insights/industry/manufacturing/global-supply-chain-resilience-amid-disruptions.html>

Kim, Sang Hui Michael. 1995. "In re Alappat: A Strict Statutory Interpretation Determining Patentable Subject Matter Relating to Computer Software?" *John Marshall Journal of Information Technology and Privacy Law* 13(4): 635–665.

Kincaid, Jason. 2009. "Facebook Activates 'Like' Button; FriendFeed Tires Of Sincere Flattery." *TechCrunch*. <https://techcrunch.com/2009/02/09/facebook-activates-like-button-friendfeed-tires-of-sincere-flattery/>

Kintner, Earl. 1978. *The Legislative History of the Federal Antitrust Laws and Related Statutes*. Chelsea House Publishers.

Kirchenbauer, Jonas, John Thickstun, Ananya Kumar, Dan Jurafsky, Percy Liang, and Tatsunori Hashimoto. 2023. A Watermark for Large Language Models. *arXiv preprint: arXiv.2301.10226*.  
<https://doi.org/10.48550/arXiv.2301.10226>

Kirkpatrick, David. 2010. *The Facebook Effect*. New York: Simon and Schuster.

Klarna. 2024. "Klarna AI Assistant Handles Two-Thirds of Customer Service Chats in Its First Month." Press Release. <https://www.prnewswire.com/news-releases/klarna-ai-assistant-handles-two-thirds-of-customer-service-chats-in-its-first-month-302072740.html>

Klein, Aaron. 2019. "Not all robots take your job, some become your co-worker." Brookings. <https://www.brookings.edu/articles/not-all-robots-take-your-job-some-become-your-co-worker/>

Klein, Benjamin, Robert Crawford, and Armen Alchian. 1978. "Vertical Integration, Appropriable Rents, and the Competitive Contracting Process." *Journal of Law and Economics* 21(2): 297–326.

Klein, Joel. 1998. "The Importance of Antitrust Enforcement in the New Economy." Speech delivered at the New York Law Journal Invitational Program on Antitrust, New York, NY, January 29. DoJ. <https://www.justice.gov/atr/speech/importance-antitrust-enforcement-new-economy>

Kline, Patrick, and Enrico Moretti. 2014. "Local Economic Development, Agglomeration Economies, and the Big Push." *The Quarterly Journal of Economics* 129(1): 275–331.

Klobuchar, Amy. 2024. "Klobuchar Statement Following the Supreme Court Decision to Overturn Chevron Deference." U.S. Senator Amy Klobuchar (press release). <https://www.klobuchar.senate.gov/public/index.cfm/2024/6/klobuchar-statement-following-the-supreme-court-decision-to-overturn-chevron-deference>.

Klonick, Kate. 2018. "The New Governors." *Harvard Law Review* 131(6): 1598–1670.

Klor's. 1959. *Klor's Inc. v. Broadway-Hale Stores, Inc.*, No. 76. US Supreme Court, April 6. <https://supreme.justia.com/cases/federal/us/359/207>

Kober, Jens, J. Andrew Bagnell, and Jan Peters. 2013. "Reinforcement Learning in Robotics: A Survey." *International Journal of Robotics Research* 32(11): 1238-1274.

Kolodny, Lora and Jennifer Elias. 2024. Waymo dominated U.S. robotaxi market in 2024, but Tesla, Zoox loom. CNBC. <https://www.cnbc.com/2024/12/26/waymo-dominated-us-robotaxi-market-in-2024-but-tesla-zoox-loom.html>

Kosseff, Jeff. 2019. *The Twenty-Six Words That Created the Internet*. Ithaca, NY: Cornell University Press.

Koster, Martijn. 1994. "A Standard for Robot Exclusion." Robots Pages.  
<http://www.robotstxt.org/orig.html>

Kovach, Steve. 2023. "Twitter Users Get Error Messages as Elon Musk Outlines New Limits, Blames A.I. Data Scraping." Fortune. <https://fortune.com/2023/07/01/twitter-error-messages-technical-problems-elon-musk-rate-limit-exceeded/>

Kovacic, William. 2007. The Intellectual DNA of Modern U.S. Competition Law for Dominant Firm Conduct: The Chicago/Harvard Double Helix. *Columbia Business Law Review* 2007(1): 1-80.

Kovacic, William E., and Carl Shapiro. 2000. "Antitrust Policy: A Century of Economic and Legal Thinking." *Journal of Economic Perspectives* 14(1): 43–60.

Kovacic, William. 1990. "The Antitrust Paradox Revisited." *Wayne Law Review* 36(3): 1413–1471.

Kovacic, William. 2003. "The Modern Evolution of US Competition Policy Enforcement Norms." *Antitrust Law Journal* 71(2): 377–478.

Kratz, Agatha, Piper, Lauren, and Bouchaud, Juliana. 2024. China and the future of global supply chains. Rhodium Group. <https://rhg.com/research/china-and-the-future-of-global-supply-chains/>

Kremer, Michael. 1993. "The O-Ring Theory of Economic Development." *The Quarterly Journal of Economics* 108(3): 551–575.

Krizhevsky, Alex, Ilya Sutskever, and Geoffrey Hinton. 2012. "ImageNet Classification with Deep Convolutional Neural Networks." In *Advances in Neural Information Processing Systems 25* (NeurIPS 2012), 1097–1105.

Krug, Sammi. 2016. "What the Reactions Launch Means for News Feed." Facebook Newsroom. <https://about.fb.com/news/2016/02/news-feed-fyi-what-the-reactions-launch-means-for-news-feed/>

KSR 2007. *KSR International Co. v. Teleflex, Inc.* No. 04-1350. US Supreme Court, April 30. <https://supreme.justia.com/cases/federal/us/550/398/>

Kubeflow. 2024. Kubeflow: Machine Learning Toolkit for Kubernetes. Kubeflow Documentation. <https://www.kubeflow.org/docs/>

Kubernetes. 2024. Kubernetes documentation. <https://kubernetes.io/docs/>

Kwa, Amanda, Elizabeth Barnes, Lawrence Chan, and Ben West. 2025. Measuring AI Ability to Complete Long Tasks. arXiv preprint arXiv:2503.14499. <https://arxiv.org/abs/2503.14499>

Kwoka, John, and Lawrence White, eds. 2014. *The Antitrust Revolution: Economics, Competition, and Policy*. 6th ed. Oxford: Oxford University Press.

Kyosovska, Nicoleta and Andrea Renda. 2025. "EU Plans for AI (Giga)Factories: Sanctuaries of Innovation, or Cathedrals in the Desert?" Centre for European Policy Studies (CEPS) Working Paper. <https://cdn.ceps.eu/2025/11/251027-Sanctuaries-or-Cathedrals.pdf>

Laffont, Jean-Jacques, and Jean Tirole. 2000. *Competition in Telecommunications*. Cambridge, MA: MIT Press.

Lam Research. 2024. Form 10-K and Annual Report filings (gross margin; business description). Lam Research Corporation, Investor Relations/SEC filings portal. [https://filecache.investorroom.com/mr5ir\\_lamresearch2/1411/LRCX%20%28Lam%20Research%20Corporation%29%20%20%2810-K%29%202024-08-29.pdf\\_.pdf](https://filecache.investorroom.com/mr5ir_lamresearch2/1411/LRCX%20%28Lam%20Research%20Corporation%29%20%20%2810-K%29%202024-08-29.pdf_.pdf)

Landes, David. 2008. *The unbound Prometheus*. Cambridge: Cambridge University Press.

Landes, William M., and Richard Posner. 1981. "Market Power in Antitrust Cases." *Harvard Law Review* 94(5): 937–996.

Langlois, Richard and Paul Robertson. 1995. *Firms, Markets, and Economic Change: A Dynamic Theory of Business Institutions*. London: Routledge.

Laricchia, Federica. 2024. "Apple: Expenditure on Research and Development 2007–2024." Statista. <https://www.statista.com/statistics/273006/apple-expenses-for-research-and-development/>.

Lascelles, Eric, and Josh Nye. 2025. MacroMemo – October 14 – November 3, 2025. RBC Global Asset Management. <https://institutional.rbcgam.com/en/us/research-insights/article/macromemo-october-14-november-3-2025-us/detail>

Lau, Johann. 2020. "Google Maps 101." *The Keyword* (Google blog). <https://blog.google/products/maps/google-maps-101-how-ai-helps-predict-traffic-and-determine-routes/>

Lawson, Emma. 2025. AI in drug discovery: AlphaFold to clinical breakthroughs.  
<https://www.elbuenonews.com/ai-steps-up-in-drug-discovery-real-world-breakthroughs-from-lab-to-clinic/>

Layne-Farrar, Anne, Jorge Padilla, and Richard Schmalensee. 2007. "Pricing Patents for Licensing in Standard-Setting Organizations: Making Sense of FRAND Commitments." *Antitrust Law Journal* 74(3): 671-706.

Lazonick, William. 2009. *Sustainable Prosperity in the New Economy? Business Organization and High-Tech Employment in the United States*. Kalamazoo, MI: W.E. Upjohn Institute for Employment Research.

LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. 2015. "Deep learning." *Nature* 521: 436–444.

Lee, Jinwoo, Fangming Wu, Wei Zhao, Mohsen Ghaffari, Linxia Liao, and David Siegel. 2014. "Prognostics and Health Management Design for Rotary Machinery Systems." *Mechanical Systems and Signal Processing* 42(1-2): 314–334.

Lee, Kevin and Shubho Sengupta. 2022. "Introducing the AI Research SuperCluster." *Meta AI Engineering*. <https://ai.meta.com/blog/ai-rsc/>

Lemley, Mark, and A. Douglas Melamed. 2013. "Missing the Forest for the Trolls." *Columbia Law Review* 113(8): 2117-2189.

Lemley, Mark, and Carl Shapiro. 2007. Patent Holdup and Royalty Stacking. *Texas Law Review*, 85(7): 1991–2049.

Lemmon. 2021. *Lemmon v. Snap, Inc.* No. 19-55362. US Court of Appeals, Ninth Circuit, May 4.  
<https://law.justia.com/cases/federal/appellate-courts/ca9/19-55362/19-55362-2021-05-04.html>

Leopold, Till. 2025. "Is AI Closing the Door on Entry-Level Job Opportunities?" *World Economic Forum*.  
<https://www.weforum.org/stories/2025/04/ai-jobs-international-workers-day/>

Lerner, Abba. 1934. "The Concept of Monopoly and the Measurement of Monopoly Power." *Review of Economic Studies* 1(3): 157-175.

Lerner, Josh. 1999. "The Government as Venture Capitalist." *The Journal of Business* 72(3): 285–318.

Lessard, Kate. 2025. The Dreamforce 2025 product announcements you can't miss. Salesforce Admins Blog. <https://admin.salesforce.com/blog/2025/dreamforce-product-announcements-for-admins>

Lessig, Lawrence. 2001. *The Future of Ideas*. New York: Vintage Books.

Leswing, Kif. 2021. "Apple's ad privacy change impact shows the power it wields over other industries." CNBC. <https://www.cnbc.com/2021/11/13/apples-privacy-changes-show-the-power-it-holds-over-other-industries.html>

Leswing, Kif. 2024. Nvidia's latest AI chip will cost more than \$30,000, CEO says. CNBC. <https://www.cnbc.com/2024/03/19/nvidias-blackwell-ai-chip-will-cost-more-than-30000-ceo-says.html>

Leswing, Kif. 2025. "Apple cuts App Store fee in half for 'mini apps' that integrate more of its software." CNBC. <https://www.cnbc.com/2025/11/13/apple-announces-new-program-that-cut-mini-app-fees-in-half.html>

Lev, Baruch. 2001. *Intangibles: Management, Measurement, and Reporting*. Washington, D.C.: Brookings Institution Press.

Levi Stuart, William Ridgway David Simon, Meredith Slawe and Anita Oh. 2024. "Utah Becomes First State To Enact AI-Centric Consumer Protection Law." Skadden, Arps, Slate, Meagher and Flom LLP. <https://www.skadden.com/insights/publications/2024/04/utah-becomes-first-state>

Levin, Cole and Ruiz-Reyes, Federico. 2025. From seats to success: The new economics of AI pricing in 2025. Pilot Blog. <https://pilot.com/blog/ai-pricing-economics-2025>

Levin, Jonathan. 2011. "The Economics of Internet Markets." NBER Working Paper No. 16852. Cambridge, MA: National Bureau of Economic Research.

Levy, Ari. 2025. Elon Musk says 80% of Tesla's value will eventually come from Optimus. CNBC. <https://www.cnbc.com/2025/09/02/musk-tesla-value-optimus-robot.html>

Levy, David and Steve Weizer. 1985. "System Error." *Regulation* 9(5–6): 27-30.

Lin, Andy. 2023. "The TSMC Cost, Sell Price, and R&D Cost of Chip Foundry." Andy Lin's Long-term Stock Investment Blog. <https://www.granitefirm.com/blog/us/2023/04/29/cost-of-chip-foundry/>

Lin, Hesheng, Van der Plas, Geert, Sun, Xiao, Velenis, Dimitrios, Catthoor, Francky, Lauwereins, Rudy, and Beyne, Eric. 2022. Efficient Backside Power Delivery for High-Performance Computing Systems. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 30(11): 1748–1756.

Lin, Liza, Josh Chin, and Raffaele Huang. 2025. “China Is Quickly Eroding America’s Lead in the Global AI Race.” *The Wall Street Journal* <https://www.wsj.com/tech/ai/artificial-intelligence-us-vs-china-03372176>

Linux Foundation. 2023. *2023 State of OSPOs and OSS Initiatives*. Linux Foundation Research. <https://www.linuxfoundation.org/research/sponsorship/state-of-ospos-2023>

Lipsy, Phillip 2020. COVID-19 and the politics of crisis. *International Organization* 74(S1): E98–E127.

Lipsey, Richard, Kenneth Carlaw, and Clifford Bekar. 2005. *Economic Transformations: General Purpose Technologies and Long Term Economic Growth*. Oxford University Press.

Lipsky, Abbott 1981. "Current Antitrust Division Views on Patent Licensing Practices." Remarks before the American Bar Association Antitrust Section, Washington, D.C., November 5. <https://www.justice.gov/atr/speech/current-antitrust-division-views-patent-licensing-practices>

Litjens, Geert, et al. 2017. "A Survey on Deep Learning in Medical Image Analysis." *Medical Image Analysis* 42: 60–88.

Liu, Glory. 2022. *Adam Smith's America*. Princeton: Princeton University Press.

Liu, Qianer. 2025. “Huawei Unveils AI Chip Roadmap to Challenge Nvidia Dominance.” *The Information*. <https://www.theinformation.com/briefings/huawei-unveils-ai-chip-roadmap-challenge-nvidia-dominance>

Liu, Xing, and Price Fishback. 2019. Effects of New Deal Spending and the downturns of the 1930s on private labor markets in 1939/1940. *Explorations in Economic History* 71(1): 25–54.

LiveRamp. 2024. “Walled Gardens in Advertising Explained.” *LiveRamp Blog*. <https://liveramp.com/blog/walled-gardens-advertising-explained>

Lomas, Natasha. 2024. “German LLM Maker Aleph Alpha Pivots to AI Support.” *TechCrunch*. <https://techcrunch.com/2024/09/05/german-llm-maker-aleph-alpha-pivots-to-ai-support/>

Long, Clarisa. 1997. TRIPs and Intellectual Property Protection in Emerging Markets. Intellectual Property Practice Group Newsletter, The Federalist Society 3(1).  
<https://fedsoc.org/commentary/publications/trips-and-intellectual-property-protection-in-emerging-markets>

Lopatka, John. 2000. United States v. IBM: A Monument to Arrogance. *Antitrust Law Journal* 68(1): 145-186.

Loper Bright 2024. Loper Bright Enterprises et al. v. Raimondo, Secretary of Commerce, et al., No. 22-451, U.S. Supreme Court, June 28. [https://www.supremecourt.gov/opinions/23pdf/22-451\\_7m58.pdf](https://www.supremecourt.gov/opinions/23pdf/22-451_7m58.pdf)

Louviere, Jordan, David Hensher, and Joffre Swait. 2000. *Stated Choice Methods: Analysis and Applications*. Cambridge: Cambridge University Press.

Luhnnow, David, and Tom Fairless. 2025. "Europe Is Losing." *Wall Street Journal*.  
<https://www.wsj.com/world/europe/europe-is-losing-fe179376>

Luminance. 2025. Luminance: Legal-Grade AI Platform. <https://www.luminance.com/>

Lévêque, François, and Yann Ménière. 2008. "Licensing 3G Mobile Technology: How Did We Arrive at a Royalty Stack?" CERNA, Mines ParisTech. [https://ideas.repec.org/p/cmfi/wpaper/wp2007\\_0701.html](https://ideas.repec.org/p/cmfi/wpaper/wp2007_0701.html)

Ma, Wayne, and Qianer Liu. 2024. "Apple Is Working on AI Server Chips With Broadcom." *The Information*. <https://www.theinformation.com/articles/apple-is-working-on-ai-server-chips-with-broadcom>

Machado, Adrian. 2025. How tiered pricing enhances your API monetization strategy. *DEV Community*.  
<https://dev.to/zuplo/how-tiered-pricing-enhances-your-api-monetization-strategy-12io>

Macher, Jeffrey, and David Mowery. 2004. "Vertical Specialization and Industry Structure in High Technology Industries." In *Business Strategy over the Industry Lifecycle (Advances in Strategic Management Vol. 21)*, edited by Joseph Baum and Anita McGahan. Bingley: Emerald Group Publishing Limited, 317–355.

MacMillan, Douglas. 2018. "Tech's 'Dirty Secret': The App Developers Sifting Through Your Gmail." The Wall Street Journal. <https://www.wsj.com/articles/techs-dirty-secret-the-app-developers-sifting-through-your-gmail-1530544442>

Macrotrends. n.d.a. Apple Market Cap 2020. Macrotrends. <https://www.macrotrends.net/stocks/charts/AAPL/apple/market-cap#AAPL-2020>

Macrotrends (n.d.b). Microsoft Market Cap 2020. Macrotrends. <https://www.macrotrends.net/stocks/charts/MSFT/microsoft/market-cap#MSFT-2020>

Macrotrends (n.d.c). Amazon Market Cap 2020. Macrotrends. <https://www.macrotrends.net/stocks/charts/AMZN/amazon/market-cap#AMZN-2020>

Macrotrends (n.d.d). Alphabet Market Cap 2020. Macrotrends. <https://www.macrotrends.net/stocks/charts/GOOGL/alphabet/market-cap#GOOGL-2020>

Macrotrends (n.d.e). Nvidia Market Cap 2025. Macrotrends. <https://www.macrotrends.net/stocks/charts/NVDA/nvidia/market-cap#NVDA-2025-07-08>

Macrotrends (n.d.f). Microsoft Market Cap 2025. Macrotrends. <https://www.macrotrends.net/stocks/charts/MSFT/microsoft/market-cap#MSFT-2025-07-08>

Macrotrends (n.d.g). Apple Market Cap 2025. Macrotrends. <https://www.macrotrends.net/stocks/charts/AAPL/apple/market-cap#AAPL-2025-07-08>

Macrotrends (n.d.h). Alphabet Market Cap 2025. Macrotrends. <https://www.macrotrends.net/stocks/charts/GOOGL/alphabet/market-cap#GOOGL-2025-07-08>

Macrotrends. n.d. "Historical Stock Price Data." Macrotrends. <https://www.macrotrends.net/>

Madden, Keith. 2025. File Types and Artificial Intelligence: Data Formats for AI Training. Hex Browser. <https://www.hexbrowser.com/file-types-artificial-intelligence-data/>.

Magistro, Beatrice and Victor Menaldo. 2025. How Populism Harms Prosperity: Unified Populist Rule Reduces Investment, Innovation, and Productivity. *Journal of Evolutionary Economics* 35(3): 553–586.

Mahaseth, Raushan Kumar. 2025. Reinforcement Learning with Human Feedback (RLHF): Shaping the Future of AI Alignment Roadmap 2025–2035. TechRxiv.

<https://www.techrxiv.org/users/963437/articles/1335644-reinforcement-learning-with-human-feedback-rlhf-shaping-the-future-of-ai-alignment-roadmap-2025-2035>

Malik, Aisha. 2022. "Meta says Reels now makes up over 20% of the time users spend on Instagram." TechCrunch. <https://techcrunch.com/2022/04/27/meta-says-reels-now-makes-up-over-20-of-the-time-users-spend-on-instagram/>

Mallinson, Keith. 2021. Royalty pricing dichotomy in 5G SEP patent pool for Open RAN Radio Units. IP Finance. <http://www.ip.finance/2021/12/royalty-pricing-dichotomy-in-5g-sep.htm>

Malwarebytes. 2020. Malwarebytes, Inc. v. Enigma Software Group USA, LLC. No. 19-1284. US Supreme Court, October 13. <https://www.law.cornell.edu/supremecourt/text/19-1284>

Mandler, Clare 2024. Three Mile Island nuclear power plant will reopen for Microsoft. NPR. <https://www.npr.org/2024/09/20/nx-s1-5120581/three-mile-island-nuclear-power-plant-microsoft-ai>

Manenti, Pierfrancesco. 2021. Supply Chain Strategic Planning Must Account for 5 Key Capabilities. Gartner Research. [https://www.scmr.com/article/strategic\\_planning\\_must\\_factor\\_in\\_5\\_key\\_capabilities](https://www.scmr.com/article/strategic_planning_must_factor_in_5_key_capabilities)

Mann, Ronald, and Thomas Sager. 2007. "Patents, Venture Capital, and Software Startups." *Research Policy* 36(2): 193–208.

Mansfield, Edwin. 1961. "Technical Change and the Rate of Imitation." *Econometrica* 29(4): 741–766.

Manyika, James, Sree Ramaswamy, Jacques Bughin, Jonathan Woetzel, Michael Birshan, and Zubin Nagpal. 2017. *The Age of Analytics: Competing in a Data-Driven World*. McKinsey Global Institute. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-age-of-analytics-competing-in-a-data-driven-world>.

MAPI Foundation. 2017. "U.S. Industrial Outlook: Glimmers of Light." MAPI Foundation. <https://mapifoundation.org/economic/2017/2/26/us-industrial-outlook-glimmers-of-light/>

Marcus, Gary. 2025a. AI Agents have, so far, mostly been a dud. Marcus on AI (Substack). <https://garymarcus.substack.com/p/ai-agents-have-so-far-mostly-been>

Marcus, Gary. 2025b. "A Knockout Blow for LLMs? LLM 'Reasoning' Is So Cooked They Turned My Name into a Verb." Marcus on AI (Substack). <https://garymarcus.substack.com/p/a-knockout-blow-for-llms>

Marcus, Gary. 2025c. "AI, Layoffs, Productivity and The Klarna Effect." Marcus on AI (Substack). <https://garymarcus.substack.com/p/ai-layoffs-productivity-and-the-klarna>

Marguerit, David. 2025. "Augmenting or Automating Labor? The Effect of AI Development on New Work, Employment, and Wages." arXiv preprint arXiv:2503.19159.

Markey, Edward, and Jack Fields. 2016. The Telecommunications Act of 1996. The Hill. <https://thehill.com/blogs/congress-blog/technology/272406-the-telecommunications-act-of-1996-then-now-and-beyond>

Markoff, John. 1997. "Microsoft Comes to the Aid of a Struggling Apple." The New York Times. <https://archive.nytimes.com/www.nytimes.com/library/cyber/week/080797apple.html>

Marshall Phelps and Cheryl Milone. 2012. LTE Standard Essential Patents Now and in the Future. Article One Partners. White paper. <https://1library.net/document/y82ddr5y-lte-standard-essential-patents-now-and-in-the-future.html>

Martin, Jeremy, Travis Mayberry, Collin Donahue, Lucas Foppe, Lamont Brown, Chadwick Riggins, Erik Rye, and Dane Brown. 2017. "A Study of MAC Address Randomization in Mobile Devices and When it Fails." *Proceedings on Privacy Enhancing Technologies* 4: 365–383.

Martin, Matthew. 2025. "For Some Recent Graduates, the A.I. Job Apocalypse May Already Be Here." New York Times, May 30. <https://www.nytimes.com/2025/05/30/technology/ai-jobs-college-graduates.html>

Maskus, Keith. 2000. Intellectual Property Rights in the Global Economy. Washington, DC: Institute for International Economics.

Maslej, Nestor, et al. 2024. "The AI Index 2024 Annual Report." AI Index Steering Committee, Institute for Human-Centered AI. Stanford University, Stanford, CA. [https://hai.stanford.edu/assets/files/hai\\_ai-index-report-2024-smaller2.pdf](https://hai.stanford.edu/assets/files/hai_ai-index-report-2024-smaller2.pdf)

Maslej, Nestor, et al. 2025. "The AI Index 2025 Annual Report," AI Index Steering Committee, Institute for Human-Centered AI. Stanford University, Stanford, CA. <https://doi.org/10.48550/arXiv.2504.07139>

Mason, Edward. 1964. *Economic Concentration and the Monopoly Problem*. New York: Atheneum.

Mastercard Incorporated. 2021. "2020 Annual Report."  
[https://s25.q4cdn.com/479285134/files/doc\\_financials/2020/ar/MA-2020-Annual-Report.pdf](https://s25.q4cdn.com/479285134/files/doc_financials/2020/ar/MA-2020-Annual-Report.pdf)

Matsushita Electric Industrial. 1986. *Matsushita Electric Industrial Co., LTD, et al. v. Zenith Radio Corp.* et al., No. 83-2004. U.S. Supreme Court, March 26. <https://supreme.justia.com/cases/federal/us/475/574>

Mattson, Peter, et al. 2020. "MLPerf: An Industry Standard Benchmark Suite for Machine Learning Performance." *IEEE Micro* 40(2): 8–16.

Maues, Julia. 2025. *Banking Act of 1933 (Glass-Steagall)*. Federal Reserve History.  
<https://www.federalreservehistory.org/essays/glass-steagall-act>

Mavens Global. 2025. *AI vs. Humans: What Jobs Will Always Require a Human Touch?*  
<https://www.linkedin.com/pulse/ai-vs-humans-what-jobs-always-require-human-touch-mavens-global-3anyc>

Mazurek, Jan. 1999. *Making Microchips: Policy, Globalization, and Economic Restructuring in the Semiconductor Industry*. Cambridge, MA: MIT Press.

Mazzucato, Maria. 2013. *The Entrepreneurial State: Debunking Public vs. Private Sector Myths*. London: Anthem Press.

McAfee, Andrew 2019. *More from Less* New York: Scribner.

McAfee, Andrew, and Erik Brynjolfsson. 2012. "Big Data." *Harvard Business Review*.  
<https://hbr.org/2012/10/big-data-the-management-revolution>

McCarthy, John, Marvin L. Minsky, Nathaniel Rochester, and Claude Elwood Shannon. 1956. "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence." *AI Magazine* 27(4): 12–14.

McCarty, Nolan, Keith Poole, and Howard Rosenthal. 2016. *Polarized America*. 2nd ed. Cambridge, MA: MIT Press.

McClain, Colleen, Kennedy, Brian, Gottfried, Jeffrey, Anderson, Monica, and Pasquini, Giancarlo. 2025.

“How the U.S. Public and AI Experts View Artificial Intelligence.” Pew Research Center.

[https://www.pewresearch.org/internet/2025/04/03/how-the-us-public-and-ai-experts-view-artificial-intelligence/?itid=lk\\_inline\\_enhanced-template](https://www.pewresearch.org/internet/2025/04/03/how-the-us-public-and-ai-experts-view-artificial-intelligence/?itid=lk_inline_enhanced-template)

McCloskey, Deirdre. 2016. *Bourgeois Equality: How Ideas, Not Capital or Institutions, Enriched the World*. Chicago: University of Chicago Press.

McCraw, Thomas. 1984. *Prophets of Regulation*. Cambridge, MA: Harvard University Press.

McCullough, Matthew. 2025. “16 Things to Know for Android Developers at Google I/O 2025.” *Android Developers Blog*. <https://android-developers.googleblog.com/2025/05/16-things-to-know-for-android-developers-google-io-2025.html>

McElheran, Kristina, Mu-Jeung Yang, Zachary Kroff, and Erik Brynjolfsson. 2025. "The Rise of Industrial AI in America." *CES Working Paper CES-25-27*. U.S. Census Bureau, Center for Economic Studies. <https://www2.census.gov/library/working-papers/2025/adrm/ces/CES-WP-25-27.pdf>

McKinsey & Company. 2025. “The State of AI 2025.”

<https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>

McKinsey Global Institute. 2000. "Why the Japanese Economy is Not Growing: Micro Barriers to Productivity Growth." *McKinsey Global Institute*.

[https://www.mckinsey.com/~/media/mckinsey/featured%20insights/asia%20pacific/why%20the%20japanese%20economy%20is%20not%20growing/mgi\\_why\\_the\\_japanese\\_economy\\_is\\_not\\_growing\\_report.pdf](https://www.mckinsey.com/~/media/mckinsey/featured%20insights/asia%20pacific/why%20the%20japanese%20economy%20is%20not%20growing/mgi_why_the_japanese_economy_is_not_growing_report.pdf)

McKinsey Global Institute. 2020. "The Future of Work After COVID-19." *McKinsey & Company*.

<https://www.mckinsey.com/featured-insights/future-of-work/the-future-of-work-after-covid-19>

McLaughlin, Kevin. 2025. *Why AI Isn't Giving Salesforce a Boost*. The Information.

<https://www.theinformation.com/articles/ai-giving-salesforce-boost>

McMahon, Liv. 2025. “ChatGPT-maker OpenAI Releases Web Browser to Rival Google.” *BBC News*.

<https://www.bbc.com/news/technology-68445981.amp>

McMillan, John. 2025. "2.5D vs. 3D IC: Which Chip Packaging Tech Is Right for You?" Siemens EDA Blog. <https://blogs.sw.siemens.com/semiconductor-packaging/2025/06/24/2-5d-vs-3d-ic-which-chip-packaging-technology-is-right-for-you/>

McShane, Clay, and Joel Tarr. 2007. *The Horse in the City: Living Machines in the Nineteenth Century*. Baltimore: Johns Hopkins University Press.

Mead Corp. 2001. *U.S. v. Mead Corp.*, No. 99-1434, U.S. Supreme Court, June 18. [https://scholar.google.com/scholar\\_case?case=6553117666921312576](https://scholar.google.com/scholar_case?case=6553117666921312576)

Mehta, Aashish, Jesús Felipe, and John McCombie. 2021. "Software Industry Global Growth and Productivity Dynamics." *Journal of Economic Perspectives* 35(2): 102-127.

Mehta, Amit. 2025. *United States v. Google LLC, Remedies Opinion*. U.S. District Court for the District of Columbia, September 2. <https://www.courthousenews.com/wp-content/uploads/2025/09/judge-amit-mehta-google-search-monopoly-remedy-ruling.pdf>

Meltzer, Allan. 2005. "Origins of the Great Inflation." *Federal Reserve Bank of St. Louis Review* 87(2): 145–175.

Menaldo, Victor. 2021. "Do Patents Foster International Technology Transfer? Evidence from Spanish Steelmaking, 1850–1930." In *The Battle over Patents: History and Politics of Innovation*, edited by Stephen Haber and Naomi Lamoreaux. New York: Oxford University Press: 69–111.

Menaldo, Victor. 2026. *Supplementary Appendix to History's Most Revolutionary Innovation* (Cambridge University Press). <https://faculty.washington.edu/vmenaldo/Books/Mbook26.pdf>

Menaldo, Victor and Nicolas Wittstock. 2021. "Does Technology Transfer from the US to China harm American firms and consumers? A historical and analytic investigation." *Economic and Political Studies* 9(4): 417-446.

Menaldo, Victor and Nicolas Wittstock. 2025. *U.S. Innovation Inequality and Trumpism*. New York: Cambridge University Press.

Menand, Louis. 2023. "The Rise and Fall of Neoliberalism." *The New Yorker*. <https://www.newyorker.com/magazine/2023/07/24/the-rise-and-fall-of-neoliberalism>

Meredith, Sam. 2025. "Meta and Google are laying a fast-growing web of mega subsea cables." CNBC. <https://www.cnbc.com/2025/07/17/meta-and-google-are-laying-a-fast-growing-web-of-mega-subsea-cables.html>

Merizzi, Nicholas, Chris Thomas, and Ed Burns. 2025. The AI infrastructure reckoning: Optimizing compute strategy in the age of inference economics. Deloitte Insights, Tech Trends 2026. <https://www.deloitte.com/us/en/insights/topics/technology-management/tech-trends/2026/ai-infrastructure-compute-strategy.html>

Merrill, Thomas. 2002. "The Mead Doctrine." *Administrative Law Review* 54(3): 807–835.

Meta. 2025. "Helping Indian Brands Grow Their Cross Border Business With Meta Ad Tools." Meta Newsroom. <https://about.fb.com/news/2025/05/helping-indian-brands-grow-their-cross-border-business-with-meta-ad-tools/>

Meta AI. 2024. "Introducing Llama 3: The Most Capable Openly Available LLM to Date." Meta AI Blog. <https://ai.meta.com/blog/meta-llama-3/>

Meta Platforms Inc. 2021. Form 10-K for the fiscal year ended December 31, 2020. SEC. <https://www.sec.gov/Archives/edgar/data/1326801/000132680121000014/mti-20201231.htm>

Meta Platforms Inc. 2022. Annual Report (Form 10-K). Menlo Park, CA: Meta Platforms Inc.

Meta Platforms Inc. 2023. Annual Report (Form 10-K). Menlo Park, CA: Meta Platforms Inc.

Meta Platforms Inc. 2024. "Meta Earnings Presentation Q3 2024." [https://s21.q4cdn.com/399680738/files/doc\\_financials/2024/q3/Earnings-Presentation-Q3-2024.pdf](https://s21.q4cdn.com/399680738/files/doc_financials/2024/q3/Earnings-Presentation-Q3-2024.pdf)

Meta Platforms Inc. 2025a. 2025. Q1 2025 Earnings Call Transcript/CFO outlook (CapEx guidance). [https://s21.q4cdn.com/399680738/files/doc\\_financials/2025/q1/Transcripts/META-Q1-2025-Earnings-Call-Transcript.pdf](https://s21.q4cdn.com/399680738/files/doc_financials/2025/q1/Transcripts/META-Q1-2025-Earnings-Call-Transcript.pdf)

Meta Platforms Inc. 2025b. "Hello, Bowling Green!" Meta Data Centers. <https://datacenters.atmeta.com/2025/04/hello-bowling-green/>

Metz, Cade, Cecilia Kang, Sheera Frenkel, Stuart Thompson, and Nico Grant. 2024. "How Tech Giants Cut Corners to Harvest Data for AI." *The New York Times*.

<https://www.nytimes.com/2024/04/06/technology/tech-giants-harvest-data-artificial-intelligence.html>

Micron. 2024. DDR5 DRAM. Micron Corporation. <https://www.micron.com/products/memory/dram-components/ddr5-sdram>

Microsoft Corp. 2025. Form 10-K for the fiscal year ended June 30, 2024 (Intangible assets schedule).

SEC. <https://www.sec.gov/Archives/edgar/data/789019/000156459024015124/msft-20240630.htm>

Microsoft Research. 2018. "Jabil Pilots Azure and Project Brainwave in Advanced Manufacturing Solutions." <https://www.microsoft.com/en-us/research/video/jabil-pilots-azure-project-brainwave-advanced-manufacturing-solutions/>

Microsoft. 1985. Microsoft Windows Operating Environment User's Guide. Microsoft Corp.

[https://americanhistory.si.edu/collections/object/nmah\\_1695414](https://americanhistory.si.edu/collections/object/nmah_1695414)

Microsoft. 2001. *United States v. Microsoft Corp.*, Nos. 00-5212, 5213. US Court of Appeals, District of Columbia Circuit, June 28. <https://law.justia.com/cases/federal/appellate-courts/F3/253/34>

Microsoft. 2024. Azure HPC and AI Solutions. <https://azure.microsoft.com/en-us/solutions/high-performance-computing/>

Microsoft. 2025a "The Next Chapter of the Microsoft–OpenAI Partnership." Official Microsoft Blog.

<https://blogs.microsoft.com/blog/2025/10/28/the-next-chapter-of-the-microsoft-openai-partnership>

Microsoft. 2025b. "Build Apps through Conversation with Copilot." Microsoft Learn (Power Apps documentation). <https://learn.microsoft.com/en-us/power-apps/maker/canvas-apps/ai-conversations-create-app>

Mih, Atah Nuh, Hung Cao, Joshua Pickard, Monica Wachowicz, and Rickey Dubay. 2023. "TransferD2: Automated Defect Detection Approach in Smart Manufacturing using Transfer Learning Techniques."

Working Paper. <https://www.cs.unb.ca/~hcao3/publications/2023-coins/TransferD2.pdf>

Milgrom, Paul, and John Roberts. 1990. "Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities." *Econometrica* 58(6): 1255–1277.

Milgrom, Paul, and John Roberts. 1982. "Limit Pricing and Entry under Incomplete Information: An Equilibrium Analysis." *Econometrica* 50(2): 443-459.

Milgrom, Paul. 2004. *Putting Auction Theory to Work*. Cambridge: Cambridge University Press.

Miller, Chris. 2022. *Chip War: The Fight for the World's Most Critical Technology*. New York: Scribner.

Miller, Gabby. 2024. *Unpacking New NIST Guidance on Artificial Intelligence*. Tech Policy Press.  
<https://techpolicy.press/unpacking-new-nist-guidance-on-artificial-intelligence>

Milli, Smitha, Micah Carroll, Yike Wang, Sashrika Pandey, Sebastian Zhao, and Anca Dragan. 2025. "Engagement, User Satisfaction, and the Amplification of Divisive Content on Social Media." *PNAS Nexus* 4(3): pgaf062.

Mills, Mark. 2021. *The Cloud Revolution*. New York: Encounter Books.

Mims, Christopher 2024. *Why Every Big Tech Company Has Failed to Dethrone Nvidia as King of AI*. Wall Street Journal. <https://www.wsj.com/tech/ai/ai-nvidia-apple-amd-jensen-huang-software-bb581f5a>

Mims, Christopher. 2025a. *The New Chips Designed to Solve AI's Energy Problem*. The Wall Street Journal. <https://www.wsj.com/tech/ai/the-new-chips-designed-to-solve-ais-energy-problem-1ba9cac1>

Mims, Christopher. 2025b. *Silicon Valley's new strategy: Move slow and build things*. The Wall Street Journal. <https://www.wsj.com/tech/ai/silicon-valley-ai-infrastructure-capex-cffe0431>

Miotto, Riccardo, Li Li, Brian Kidd, and Joel Dudley. 2016. "Deep Patient: An Unsupervised Representation to Predict the Future of Patients from the Electronic Health Records." *Scientific Reports* 6(1): 26094.

Mirowski, Philip, and Dieter Plehwe, eds. 2009. *The Road from Mont Pèlerin: The Making of the Neoliberal Thought Collective*. Cambridge, MA: Harvard University Press.

Misa, Thomas. 1985. "Military Needs, Commercial Realities, and the Development of the Transistor, 1948–1958." In *Military Enterprise and Technological Change*, edited by Merritt Roe Smith. Cambridge, MA: MIT Press, 253–287.

Mistral AI. 2024. "Mistral Large: Technical Report." Paris: Mistral AI. <https://mistral.ai/>

Mitchel. 1711. *Mitchel v. Reynolds*, 24 Eng. Rep. 347. Court of King's Bench.  
<https://www.quimbee.com/cases/mitchel-v-reynolds>

Mitchell, Laura, and Nicole Trotta 2025. "OFCCP To Close All Prior Section 503 and VEVRAA Compliance Reviews Following Secretary of Labor Order." *Affirmative Action and OFCCP Law Advisor*.  
<https://www.affirmativeactionlawadvisor.com/2025/07/ofccp-to-close-all-prior-section-503-and-vevraa-compliance-reviews-following-secretary-of-labor-order-reviving-enforcement-activities/>

Mitchell, Margaret, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. "Model Cards for Model Reporting." *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT)*: 220–229.

MLCommons. 2024. *MLPerf benchmarks: Driving transparent AI performance standards*. MLCommons.  
<https://mlcommons.org/en/mlperf/>

MLflow. 2024. *MLflow: Open source platform for the machine learning lifecycle*.  
<https://mlflow.org/docs/latest/index.html>

Mobley. 2025. *Mobley v. Workday, Inc.* No. 23-cv-00770-RFL. Northern District of California, June 11.  
[https://www.govinfo.gov/content/pkg/USCOURTS-cand-3\\_23-cv-00770/pdf/USCOURTS-cand-3\\_23-cv-00770-1.pdf](https://www.govinfo.gov/content/pkg/USCOURTS-cand-3_23-cv-00770/pdf/USCOURTS-cand-3_23-cv-00770-1.pdf)

Mokyr, Joel. 1990. *The Lever of Riches*. New York: Oxford University Press.

Mollick, Ethan. 2024. *Co-Intelligence: Living and Working with AI*. New York: Portfolio/Penguin.

Mollick, Ethan. n.d. *One Useful Thing*. Substack newsletter. <https://www.oneusefulthing.org/>

Montgomery, Blake and Johana Bhuiyan. 2024. "California won't require big tech firms to test safety of AI after Newsom kills bill." *The Guardian*. <https://www.theguardian.com/us-news/2024/sep/29/california-governor-gavin-newsom-vetoes-ai-safety-bill>

Moore, Gordon. 1965. *Cramming More Components onto Integrated Circuits*. *Electronics* 38(8): 114–117.

Mordor Intelligence. 2024. *Electronic Design Automation Tools Market Analysis (2025–2030)*.  
<https://www.mordorintelligence.com/industry-reports/electronic-design-automation-eda-tools-market>

Moresi, Serge, and Carl Shapiro. 2012. "Any-and-All Offers in Antitrust and Consumer Protection." *Antitrust* 26(3): 41–45.

Motor Vehicle Manufacturers Association. 1983. *Motor Vehicle Manufacturers Association of the United States, Inc. et al. v. State Farm Mutual Automobile Insurance Co. et al.*, No. 82-354, U.S. Supreme Court, June 24. <https://supreme.justia.com/cases/federal/us/463/29/>

Mowery, David and Arvis Ziedonis. 2001. The Growth of Patenting and Licensing by U.S. Universities. *Research Policy* 30(1): 99-119.

Mowery, David, and Timothy Simcoe. 2004. "Is the Internet a U.S. Invention? An Economic and Technological History of Computer Networking." *Research Policy* 31(8-9): 1369-1387.

Mowery, David, Richard Nelson, Bhaven Sampat, and Arvids Ziedonis. 2001. "The Growth of Patenting and Licensing by U.S. Universities: An Assessment of the Effects of the Bayh–Dole Act of 1980." *Research Policy* 30(1): 99–119.

Mowery, David, and Bhaven Sampat. 2005. "The Bayh–Dole Act of 1980 and University–Industry Technology Transfer." *Journal of Technology Transfer* 30(1–2): 115–127.

Mowery, David, Richard Nelson, Bhaven Sampat, and Arvids Ziedonis. 2004. *Ivory Tower and Industrial Innovation*. Stanford: Stanford University Press.

Mozilla. 2019. "Firefox Now Available with Enhanced Tracking Protection by Default." Mozilla Blog. <https://blog.mozilla.org/en/products/firefox/firefox-now-available-with-enhanced-tracking-protection-by-default/>

Mozingo, Tom. 2020. Revisiting the Enforceability of Online Contracts: The Need for Unambiguous Assent to Inconspicuous Terms. *Seattle University Law Review* 43(3): 1065.

Mukesh, Sagarika, and Jingyun Zhang. 2022. "A Review of the Gate-All-Around Nanosheet FET Process Opportunities." *Electronics* 11(21): 3589.

Muris, Timothy. 2023. "Neo-Brandeisian Antitrust: Repeating History's Mistakes." *Regulation* 46(1): 6–9.

Musk, Elon. 2023. "To address extreme levels of data scraping and system manipulation, we've applied the following temporary limits." X (formerly Twitter).

<https://x.com/elonmusk/status/1675187969420828672%5C%3E%5B2>

Myung, Hyung, Junsung Lim, and David Goodman. 2006. "Single Carrier FDMA for Uplink Wireless Transmission." *IEEE Vehicular Technology Magazine* 1(3): 30-38.

Na, Hang Ryeol 2025. "Environmental Governance Under Trump: From EPA Decline to Judicial Ascendancy." NiCHE. <https://niche-canada.org/2025/06/25/environmental-governance-under-trump-from-epa-decline-to-judicial-ascendancy/>

Nada Golmie, Alhussein Abouzeid, Thyagarajan Nandagopal, Murat Torlak, Marc Leh, and Miller Higgins. 2023. NextG Communications Research and Development Gaps Report (NIST Special Publication 1293). NIST. <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1293.pdf>

Nagabandi, Aravind, Kurt Konolige, Sergey Levine, and Vijay Kumar. 2020. "Deep Dynamics Models for Learning Dexterous Manipulation." *Proceedings of the Conference on Robot Learning*, PMLR 100: 1101–1112.

Nagle, Frank, Robert Seamans, and Steven Tadelis. 2020. Transaction Cost Economics in the Digital Economy: A Research Agenda. Harvard Business School Strategy Unit Working Paper No. 21-009.

Nagler, Markus, Monika Schnitzer, and Martin Watzinger. 2022. "Fostering the Diffusion of General Purpose Technologies: Evidence from the Licensing of the Transistor Patents." *Journal of Industrial Economics* 70(4): 828-875.

NAHB 2025. "DOE Suspends Energy Efficiency Mandates on Key Home Appliances." National Association of Home Builders. <https://www.nahb.org/blog/2025/02/doe-suspends-energy-efficiency-mandates-on-key-home-appliances>

Nahra, Kirk, Jennings, Molly, Straky, Andrew, Jessani, Ali. 2025. Year in Review: 2024 BIPA Litigation Takeaways. WilmerHale. <https://www.wilmerhale.com/en/insights/blogs/wilmerhale-privacy-and-cybersecurity-law/20250219-year-in-review-2024-bipa-litigation-takeaways>

Naidu, Richa, Martin Coulter and Jason Lange. 2023. "Focus: ChatGPT Fever Spreads to U.S. Workplace, Sounding Alarm for Some." Reuters. <https://www.reuters.com/technology/chatgpt-fever-spreads-us-workplace-sounding-alarm-some-2023-08-11/>

Nanda, Arun. 2025. Azure OpenAI. Datacamp. <https://www.datacamp.com/tutorial/azure-openai-getting-started>

NARA. 2024. Code of Federal Regulations, Title 37: Patents, Trademarks, and Copyrights. Office of the Federal Register. <https://www.govinfo.gov/content/pkg/CFR-2024-title37-vol1/pdf/CFR-2024-title37-vol1.pdf>

Narayanan, Arvind, and Sayash Kapoor. 2025. "AI as Normal Technology." Knight First Amendment Institute at Columbia University. <https://knightcolumbia.org/content/ai-as-normal-technology>

NCCUSL. 1999. Uniform Computer Information Transactions Act. Chicago: NCCUSL.

National Science Board. 2020. *The State of U.S. Science and Engineering 2020*. Alexandria, VA: National Science Foundation. <https://nces.nsf.gov/pubs/nsb20201>

National Science Board. 2022. *The State of U.S. Science and Engineering 2022*. Alexandria, VA: National Science Foundation. <https://nces.nsf.gov/pubs/nsb20221>

NaturalGas.org. 2013. The History of Regulation. <http://naturalgas.org/regulation/history/>

Nayak, Pandu. 2019. "Understanding searches better than ever before." The Keyword (Google). <https://blog.google/products/search/search-language-understanding-bert/>

NCAA. 1984. NCAA v. Board of Regents of the University of Oklahoma, 468 U.S. 85. <https://supreme.justia.com/cases/federal/us/468/85>

NCES. 2000. Internet Access in U.S. Public Schools and Classrooms: 1994–99. Stats in Brief. <https://nces.ed.gov/pubs2000/2000086.pdf>

NCRALP. 1979. Report to the President and the Attorney General of the National Commission for the Review of Antitrust Laws and Procedures. Washington, DC: U.S. Government Printing Office.

Nebuly. 2024. OpenAI GPT-4 API pricing: The drops from 2023 to 2024. <https://www.nebuly.com/blog/openai-gpt-4-api-pricing>

Nellis, Stephen. 2023. "Qualcomm to supply Apple with 5G chips until 2026 under new deal." Reuters. <https://www.reuters.com/technology/qualcomm-supply-apple-with-5g-chips-until-2026-under-new-deal-2023-09-11/>

Nellis, Stephen. 2025. "Stargate UAE AI datacenter to begin operation in 2026." Reuters. <https://www.reuters.com/business/media-telecom/stargate-uae-ai-datacenter-begin-operation-2026-2025-05-22/>

Nelson, Richard, and Edmund Phelps. 1966. "Investment in Humans, Technological Diffusion, and Economic Growth." *American Economic Review* 56(1/2): 69-75.

Nelson, Richard. 1959. "The Simple Economics of Basic Scientific Research." *Journal of Political Economy* 67(3): 297–306.

Neumark, David, and William Wascher. 2008. *Minimum Wages*. Cambridge, MA: MIT Press.

Nicholas, Gabriel. 2022. "Shedding Light on Shadowbanning." Center for Democracy and Technology. <https://cdt.org/insights/shedding-light-on-shadowbanning/>.

Nickelsburg, Monica, Joshua McNichols, Lucy Soucek, Alec Cowan, and Carol Smith. 2025. Is our AI obsession good for small town America? Booming Podcast. <https://www.kuow.org/stories/is-our-ai-obsession-good-for-small-town-america>

Nielsen, Rasmus Kleis, and Sarah Anne Ganter. 2022. *The Power of Platforms*. New York: Oxford University Press.

Nikolić, Igor. 2017. "Who needs injunctions? Alternative remedies in standard essential patents disputes." *Journal of Intellectual Property Law and Practice* 12(2): 126–135.

Niskanen, William. 1988. *Reaganomics*. New York: Oxford University Press.

NIST. 2023. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. NIST AI 100-1. <https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf>

NIST. 2024. *Interagency Committee on Standards Policy (ICSP)*. <https://www.nist.gov/standardsgov/interagency-committee-standards-policy-icsp>

Nixon, Richard. 1971. Executive Order 11615: Stabilization of Prices, Rents, Wages, and Salaries. Federal Register 36(159): 15727.

Nordhaus, William D. 2004. "Schumpeterian Profits in the American Economy: Theory and Measurement." NBER Working Paper No. 10433. Cambridge, MA: National Bureau of Economic Research.

Nordhaus, William. 2007. "Two Centuries of Productivity Growth in Computing." *The Journal of Economic History* 67(1): 128–159.

Northern Pacific Railway. 1958. *Northern Pacific Railway Co. v. United States*, 356 U.S. 1. <https://supreme.justia.com/cases/federal/us/356/1>

Noy, Shakked, and Whitney Zhang. 2023. "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence." *Science* 3816654: 187–192.

NSB. 2020. *Science and Engineering Indicators 2020. "Academic Research and Development."* Alexandria, VA: National Science Foundation.

NSB. 2024. *Science and Technology: Public Perceptions and Awareness in the United States.* NSF. <https://nces.nsf.gov/pubs/nsb20244/assets/nsb20244.pdf>

Nuechterlein, Jonathan, and Philip Weiser. 2005. *Digital Crossroads.* Cambridge, MA: MIT Press.

Nuvolari, Alessandro. 2009. "Understanding Successive Industrial Revolutions: A 'Collective Invention' Perspective." *Rivista di Politica Economica* 99(3-4): 9–40.

Nuñez, Michael. 2025. "Anthropic Launches Enterprise 'Agent Skills' and Opens the Standard, Challenging OpenAI in Workplace AI." *VentureBeat.* <https://venturebeat.com/ai/anthropic-launches-enterprise-agent-skills-and-opens-the-standard>

NVIDIA Developer. 2025. *Deep Learning Frameworks.* NVIDIA Developer. <https://developer.nvidia.com/deep-learning-frameworks>

NVIDIA. 2024a. "NVIDIA H100 Tensor Core GPU Architecture White Paper." <https://www.advancedclustering.com/wp-content/uploads/2022/03/gtc22-whitepaper-hopper.pdf>

NVIDIA. 2024b. "NVIDIA Blackwell Platform Arrives to Power a New Era of Computing." NVIDIA Investor Relations Press Release. <https://investor.nvidia.com/news/press-release-details/2024/NVIDIA-Blackwell-Platform-Arrives-to-Power-a-New-Era-of-Computing/default.aspx>

NVIDIA. 2025a. "CFO Commentary on Fourth Quarter and Fiscal 2025 Results." Nvidia Investor Relations. [https://investor.nvidia.com/files/doc\\_financials/2025/Q425/Q4FY25-CFO-Commentary.pdf](https://investor.nvidia.com/files/doc_financials/2025/Q425/Q4FY25-CFO-Commentary.pdf)

NVIDIA. 2024c. "NVIDIA Ethernet Networking Accelerates World's Largest AI Supercomputer, Built by xAI." NVIDIA Newsroom. <https://nvidianews.nvidia.com/news/spectrum-x-ethernet-networking-xai-colossus>

NVIDIA. 2025b. NVIDIA TensorRT. Nvidia Documentation Hub. <https://docs.nvidia.com/tensorrt/index.html>

NVIDIA. 2025c. "NVIDIA Announces DGX Cloud Lepton to Connect Developers to NVIDIA's Global Compute Ecosystem." Nvidia Press Release. <https://nvidianews.nvidia.com/news/nvidia-announces-dgx-cloud-lepton-to-connect-developers-to-nvidias-global-compute-ecosystem>

NVM Express. 2024. "About NVMe Technology." NVM Express Inc. <https://nvmexpress.org/specifications/>

Nye, Joseph. 2004. *Power in the Global Information Age: From Realism to Globalization*. New York: Routledge.

Nylen, Leah. 2023. "Google Pays Apple 36% of Search Revenue from Safari, Witness Says." Bloomberg. <https://www.bloomberg.com/news/articles/2023-11-13/apple-gets-36-of-google-revenue-from-search-deal-witness-says>

Nynex. 1998. *Nynex Corp. v. Discon, Inc.*, No. 96-1570. US Supreme Court, December 14. <https://supreme.justia.com/cases/federal/us/525/128>

O'Hara, Kieron, and Wendy Hall. 2018. *Four Internets: The Geopolitics of Digital Governance*. CIGI Paper No. 206. Centre for International Governance Innovation.

O'Reilly, Tim. 2021. "Data Is the New Sand." *The Information*. <https://www.theinformation.com/articles/data-is-the-new-sand>.

O'Rourke, Kevin, and Jeffrey Williamson. 1999. *Globalization and History: The Evolution of a Nineteenth-Century Atlantic Economy*. Cambridge, MA: MIT Press.

Obama, Barack. 2011. "Executive Order 13563—Improving Regulation and Regulatory Review." 76 FR: 3821. <https://www.federalregister.gov/documents/2011/01/21/2011-1385/improving-regulation-and-regulatory-review>

OECD Statistics. 2025. "R&D Spending Growth Slows in OECD, Surges in China: Government Support for Energy and Defence R&D Rises Sharply." OECD, March 30.

OECD. 2023a. *Trade in Value Added (TiVA) Indicators: United States*. OECD. <https://www.oecd.org/content/dam/oecd/en/topics/policy-sub-issues/trade-in-value-added/tiva-2023-USA.pdf>

OECD. 2023b. *Trade in Value Added (TiVA) Database, Version 2023*. OECD. <https://www.oecd.org/sti/ind/inter-country-input-output-tables.htm>.

OECD. 2024. *OECD Principles on Artificial Intelligence*. <https://oecd.ai/en/ai-principles>

Ohlhausen, Maureen. 2013. "Antitrust Enforcement in High Technology Markets," FTC Speech, October 16. CCIA Annual Washington Caucus. <https://www.ftc.gov/news-events/news/speeches/antitrust-enforcement-high-technology-markets>

OIRA. 2024. "Historical Reports: OIRA Concluded Reviews (1981–2023)." Washington, D.C.: General Services Administration.

Oladayo, Oluwatosin. 2024. "Future-proofing AI storage infrastructure: Managing scale, performance and data diversity." *Open Access Research Journal of Science and Technology* 12(1): 170-185.

Ollinger, Michael, and Valerie Mueller. 2003. *Managing for Safer Food*. Agricultural Economic Report No. 817. Department of Agriculture, Economic Research Service. [https://ers.usda.gov/sites/default/files/\\_laserfiche/publications/41496/18901\\_aer817.pdf?v=26386](https://ers.usda.gov/sites/default/files/_laserfiche/publications/41496/18901_aer817.pdf?v=26386)

OMB. 1996. *Economic Analysis of Federal Regulations under Executive Order 12866*. OMB Memorandum for Members of the Regulatory Working Group. <https://trumpwhitehouse.archives.gov/wp->

content/uploads/2017/11/1996-Memorandum-for-Members-of-the-Regulatory-Working-Group-Economic-Analysis-of-Federal-Regulations-Under-Executive-Order-12866.pdf

OMB. 1998. "Circular A-119: Federal Participation in the Development and Use of Voluntary Consensus Standards and in Conformity Assessment Activities." OMB. FR 63: 8546–55.  
<https://www.federalregister.gov/documents/1998/02/19/98-4177/omb-circular-a-119-federal-participation-in-the-development-and-use-of-voluntary-consensus-standards>

OMB. 2003. "Circular A-4: Regulatory Analysis." OMB. FR 68: 58366.  
<https://www.federalregister.gov/documents/2003/10/09/03-25606/circular-a-4-regulatory-analysis>

OPC Foundation. 2023. OPC UA Interoperability for Industrie 4.0 and IoT. OPC Foundation.  
<https://opcfoundation.org/wp-content/uploads/2023/05/OPC-UA-Interoperability-For-Industrie4-and-IoT-EN.pdf>

OpenAI. 2023. GPT-4 Technical Report. arXiv preprint arXiv:2303.08774.  
<https://arxiv.org/abs/2303.08774>

OpenAI. 2025a. "Introducing Operator." OpenAI. <https://openai.com/index/introducing-operator/>

OpenAI. 2025b. "Computer-Using Agent." OpenAI. <https://openai.com/index/computer-using-agent/>

OpenAI. 2025c. Introducing GPT-5. OpenAI. <https://openai.com/en-GB/index/introducing-gpt-5>

OpenAI. 2025d. "Update to GPT-5 System Card: GPT-5.2." OpenAI Research.  
<https://openai.com/index/gpt-5-system-card-update-gpt-5-2/>

OpenAI. 2025e. "OpenAI Co-Founds the Agentic AI Foundation under the Linux Foundation." OpenAI Blog, December. <https://openai.com/index/agentic-ai-foundation/>

OpenAI. 2025f. "Pricing." OpenAI API Documentation. <https://platform.openai.com/docs/pricing>

OpenAI. 2025g. "Stargate advances with 4.5 GW partnership with Oracle."  
<https://openai.com/index/stargate-advances-with-partnership-with-oracle/>

OpenAI. n.d. Pricing and Business Model. <https://chatgpt.com/pricing>

Ord, Toby. 2025. "The Scaling Paradox." <https://www.tobyord.com/writing/the-scaling-paradox>

Oregon Steam Navigation. 1873. Oregon Steam Navigation Company v. Winsor, 87 U.S. 64, U.S. Supreme Court. <https://supreme.justia.com/cases/federal/us/87/64/>

Orgvue. 2025. "55% of Businesses Admit Wrong Decisions in Making Employees Redundant When Bringing AI into the Workforce." Orgvue.com. <https://www.orgvue.com/news/55-of-businesses-admit-wrong-decisions-in-making-employees-redundant-when-bringing-ai-into-the-workforce/>

Osawa, Juro and Qianer Liu. 2025. "DeepSeek: a National Treasure in China, Now Closely Guarded." The Information. <https://www.theinformation.com/articles/deepseek-national-treasure-china-now-closely-guarded?rc=tuorsb>.

OSHA. 2012. Hazard Communication Standard: Safety Data Sheets. DoL. <https://www.osha.gov/hazcom>

Ouyang, Long, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems* 35:27730–27744.

Owens, Kellie, Zachary Griffen, and Lasya Damaraju, 2025. "Managing a 'Responsibility Vacuum' in AI Monitoring and Governance in Healthcare: A Qualitative Study." *BMC Health Services Research* 25(1): 1217.

Ozkan, Serdar, and Nicholas Sullivan. 2025. "Is AI Contributing to Rising Unemployment? Evidence from Occupational Variation." Federal Reserve Bank of St. Louis—On the Economy Blog. <https://www.stlouisfed.org/on-the-economy/2025/aug/is-ai-contributing-unemployment-evidence-occupational-variation>.

O'Mara, Margaret 2019. *The Code*. New York: Penguin Press.

P1 Security. 2025. 4G LTE Architecture and Security Explained: Protocols, Attack Vectors, and Persistent Vulnerabilities. P1 Security Blog. <https://www.p1sec.com/blog/4g-architecture-and-security>

Pacific Bell. 2009. Pacific Bell Telephone Co. v. linkLine Communications, Inc., No. 07-512. US Supreme Court, February 25. <https://supreme.justia.com/cases/federal/us/555/438>

Page, Larry. 2014. "Lenovo to Acquire Motorola Mobility." Google Blog. <https://blog.google/inside-google/company-announcements/lenovo-to-acquire-motorola-mobility/>

Page, William, and John Lopatka. 2007. *The Microsoft Case: Antitrust, High Technology, and Consumer Welfare*. Chicago: University of Chicago Press.

Palazzolo, Stephanie, and Cory Weinberg. 2025. "The Little-Known Startup That Has Surged Past Scale AI—Without Any Investors." *The Information*. <https://www.theinformation.com/articles/little-known-startup-surged-past-scale-ai-without-investors>

Palazzolo, Stephanie, and Sri Muppidi. 2025. "How Anthropic and OpenAI Are Developing AI 'Co-Workers'." *The Information*. <https://www.theinformation.com/articles/anthropic-openai-developing-ai-co-workers>

Palmer, Annie. 2025. "AI Will Shrink Amazon's Workforce in the Coming Years, CEO Jassy Says." *CNBC*. <https://www.cnbc.com/2025/06/17/ai-amazon-workforce-jassy.html>

Palmer, Maija. 2019. "Does Europe need to fix 'monkey money' stock options?" *Sifted*. <https://sifted.eu/articles/europe-fix-stock-options-notoptional-employee-ownership>

Parbel, Matthias. 2025. Snowflake and Salesforce launch open standard for standardized AI data (Open Semantic Interchange). *Heise Online*. <https://www.heise.de/en/news/Snowflake-and-Salesforce-launch-open-standard-for-standardised-AI-data-10669447.html>

Pardau, Stuart. 2018. "The California Consumer Privacy Act: Towards a European-Style Privacy Regime in the United States?" *Journal of Technology Law and Policy* 23(1): 68-114.

Parker 1978. *Parker, Acting Commissioner of Patents and Trademarks v. Flook*, No. 77-642, U.S. Supreme Court, June 22. [https://scholar.google.com/scholar\\_case?case=12542933152070461616](https://scholar.google.com/scholar_case?case=12542933152070461616)

Parker, Edward, Benjamin Miller, and Colin Levaunt. 2025. *Who Could Fund Future Artificial Intelligence Development? Expert Insights*. RAND Corporation. <https://www.rand.org/pubs/perspectives/PEA3701-1.html>

Parker, Geoffrey, and Marshall Van Alstyne. 2005. "Two-Sided Network Effects: A Theory of Information Product Design." *Management Science* 51(10): 1494-1504.

Parker, Geoffrey, Marshall Van Alstyne, and Sangeet Paul Choudary. 2016. *Platform Revolution*. New York: W.W. Norton.

Partovi, Hadi. 2025. No, AI isn't killing computer science. It's making it essential. *The Information*.  
<https://www.theinformation.com/articles/ai-killing-computer-science-making-essential>

Patel, Dwarkesh. 2025. "Satya Nadella – Microsoft's AGI Plan and Quantum Breakthrough." *The Dwarkesh Podcast*. <https://www.dwarkesh.com/p/satya-nadella>

Patel, Marc Sorel, and Pankaj Sachdeva. 2025. What Is a Data Center? *McKinsey Explainer*. McKinsey & Company. [https://www.mckinsey.com/featured-insights/mckinsey-explainers/what-is-a-data-center#/  
PatentTrack. 2024. Qualcomm patent portfolio overview. <https://patenttrack.com/strategic-patent-integration-qualcomms-portfolio/>](https://www.mckinsey.com/featured-insights/mckinsey-explainers/what-is-a-data-center#/)

Patterson, David and John Hennessy. 2021. *Computer Organization and Design: The Hardware/Software Interface* (6th ed.). San Francisco: Morgan Kaufmann.

Patterson, David, Joseph Gonzalez, Quoc Le, Chen Liang, Lluís-Miquel Munguia, Daniel Rothchild, David So, Maud Texier, and Jeff Dean. 2021. "Carbon Emissions and Large Neural Network Training." *arXiv:2104.10350*. <https://arxiv.org/pdf/2104.10350.pdf>

Paxson, David. 1981. "Potential Impact of Motor Carrier Act of 1980 on Railroad Industry." *Transportation Research Record* 804: 34–39.

PCPIP. 1883. Paris Convention for the Protection of Industrial Property of March 20, 1883, as revised at Brussels on December 14, 1900, at Washington on June 2, 1911, at The Hague on November 6, 1925, at London on June 2, 1934, at Lisbon on October 31, 1958, and at Stockholm on July 14, 1967. WIPO. <https://www.wipo.int/wipolex/en/text/288514>

Peck, Emily. 2025. "AI and College Grads: The Job Crunch Is Real." *Axios*.  
<https://www.axios.com/2025/05/29/ai-college-grads-work-jobs>

Peltzman, Sam. 1976. "Toward a More General Theory of Regulation." *Journal of Law and Economics* 19(2): 211-240.

Perez, Carlota. 2002. *Technological Revolutions and Financial Capital*. Cheltenham, UK: Edward Elgar Publishing.

Peteraf, Margaret. 1993. "The Cornerstones of Competitive Advantage: A Resource-Based View." *Strategic Management Journal* 14(3): 179-191.

Peterson, Kyle. 2011. "Boeing Dreamliner Makes First Delivery after Three-Year Delay." Reuters. <https://www.reuters.com/article/2011/01/20/us-boeing-dreamliner-idUSTRE70J2UX20110120/>

Petzold, Charles. 2022. *Code: The hidden language of computer hardware and software* (2nd ed.). Redmond, WA: Microsoft Press.

Philadelphia National Bank. 1963. *United States v. Philadelphia National Bank*, No. 83. US Supreme Court, June 17. <https://supreme.justia.com/cases/federal/us/374/321/>

Pigou, Arthur. 1920. *The Economics of Welfare*. London: Macmillan.

Pindyck, Robert. 2004. "Mandatory Unbundling and Irreversible Investment in Telecom Networks." *Review of Network Economics* 3(3): 274-298.

Pindyck, Robert. 2007. "Mandatory Unbundling and Irreversible Investment in Telecom Networks." *Review of Network Economics* 6(3): 274–298.

Pineda, Michael. 2021. "Major Suppliers of Apple: Inside Its Supply Chain." Profolus. <https://www.profolus.com/topics/major-suppliers-of-apple-inside-its-supply-chain/>

Piosenka, Gerald Vincent, and Ronald Chandos. 1991. *Unforgeable Personal Identification System*. (US Patent No. 4,993,068). USPTO. <https://patents.google.com/patent/US4993068A/en>

Pitofsky, Robert 2001. *The Essential Facilities Doctrine Under U.S. Antitrust Law*. *Antitrust Law Journal* 70(2): 443-462.

PM Insights. 2025. "Anthropic Approaches \$7B Run Rate in 2025, Outpaces OpenAI." PM Insights Weekly Edition and Stock Spotlight. <https://www.pminsights.com/insights/anthropic-approaches-7b-run-rate-in-2025-outpaces-openai>

Polaroid 1986. *Polaroid Corp. v. Eastman Kodak Co.*, No. 76-1634-Z, U.S. District Court, D. Massachusetts, October 11. <https://law.justia.com/cases/federal/district-courts/FSupp/641/828/1482978>

Pooley, James. 2015. *Secrets*. Palo Alto, CA: Verus Press.

Porter, Michael. 1980. *Competitive Strategy: Techniques for Analyzing Industries and Competitors*. New York: Free Press.

Porter, Michael. 1985. *Competitive Advantage: Creating and Sustaining Superior Performance*. New York: Free Press.

Porter, Michael, and James Heppelmann. 2014. "How Smart, Connected Products Are Transforming Competition." *Harvard Business Review*. <https://hbr.org/2014/11/how-smart-connected-products-are-transforming-competition>

Porter, Michael, and Mariko Sakakibara. 2004. "Competition in Japan." *Journal of Economic Perspectives* 18(1): 27–50.

Posner, Eric, and E. Glen Weyl. 2018. *Radical Markets*. Princeton, NJ: Princeton University Press.

Posner, Richard. 1976. *Antitrust Law: An Economic Perspective*. Chicago: University of Chicago Press.

Posner, Richard A. 1979. The Chicago School of Antitrust Analysis, 127 *University of Pennsylvania Law Review* 925.

Posner, Richard. 2001. *Antitrust Law*. 2nd ed. Chicago: University of Chicago Press.

Prabhakar, Raghu et al. 2024. "SambaNova SN40L: Scaling the AI Memory Wall with Dataflow and Composition of Experts." arXiv preprint arXiv:2405.12345. <https://arxiv.org/abs/2405.12345>

Pregelj, Vladimir. 2001. Most-Favored-Nation Status of the People's Republic of China. CRS Report RL30225. [https://digital.library.unt.edu/ark:/67531/metacrs2000/m2/1/high\\_res\\_d/RL30225\\_2001Jun07.pdf](https://digital.library.unt.edu/ark:/67531/metacrs2000/m2/1/high_res_d/RL30225_2001Jun07.pdf)

Procore Technologies 2024. "Procore Launches Procore AI with New Agents to Boost Construction Management Efficiency." Press release. <https://www.procore.com/press/procore-launches-procore-ai-with-new-agents-to-boost-construction-management-efficiency>

Procore Technologies. n.d. "Procore Insights." <https://www.procore.com/insights>

Provost, Foster, and Tom Fawcett. 2013. *Data Science for Business*. Sebastopol, California: O'Reilly Media.

PwC. 2025. "PwC's 2025 Digital Trends in Operations Survey." PwC.  
<https://www.pwc.com/us/en/services/consulting/business-transformation/digital-supply-chain-survey.html>

Pérez-Cruz, Fernando, Prenio, Jermy, Restoy, Fernando, and Yong, Jeffery. 2025. Managing explanations. FSI Occasional Paper No. 24. Bank for International Settlements. <https://www.bis.org/fsi/fsipapers24.pdf>

Qi, Yuan, Thomas Minka, Rosalind Picard, and Zoubin Ghahramani. 2004. Predictive Automatic Relevance Determination by Expectation Propagation. In *Proceedings of the Twenty-First International Conference on Machine Learning (ICML 2004)*. Banff, Canada, 846–853.

Qualcomm. 2019. *FTC v. Qualcomm Inc.*, 414 F. Supp. 3d 1101 (N.D. Cal.).  
<https://law.justia.com/cases/federal/district-courts/california/candce/3:2017cv01225/284890/218/>

Qualcomm. 2020. *FTC v. Qualcomm, Inc.*, No. 19-16122, 9th Cir., Aug. 11.  
<https://law.justia.com/cases/federal/appellate-courts/ca9/19-16122/19-16122-2020-08-11.html>

Qualcomm Inc. 2006. "Annual Report on Form 10-K for the Fiscal Year Ended September 24, 2006." SEC. <https://www.sec.gov/Archives/edgar/data/804328/000093639206000996/a24679e10vk.htm>

Qualcomm Inc. 2008. "LTE/WiMax Patent Licensing Statement." Qualcomm Public Policy Documents. <https://www.qualcomm.com/media/documents/files/lte-wimax-patent-licensing-statement.pdf>

Qualcomm Inc. 2014. Annual Report (Form 10-K). San Diego, CA: Qualcomm Inc. SEC. <https://www.sec.gov/Archives/edgar/data/804328/000123445214000320/qcom10-k2014.htm>

Qualcomm Inc. 2017. "Qualcomm 5G NR Royalty Terms Statement." Qualcomm Public Releases, November 19. <https://www.qualcomm.com/media/documents/qualcomm-5g-nr-royalty-terms-statement>

Qualcomm Inc. 2018. Annual Report (Form 10-K). San Diego, CA: Qualcomm Inc. SEC. <https://www.sec.gov/Archives/edgar/data/804328/000172894918000095/qcom10-k2018.htm>

Qualcomm Inc. 2019a. *FTC v. Qualcomm Inc.*, Case No. 17-CV-00220-LHK (Findings of Fact and Conclusions of Law). U.S. District Court, Northern District of California, May 21.  
[https://www.ftc.gov/system/files/documents/cases/qualcomm\\_findings\\_of\\_fact\\_and\\_conclusions\\_of\\_law.pdf](https://www.ftc.gov/system/files/documents/cases/qualcomm_findings_of_fact_and_conclusions_of_law.pdf)

Qualcomm Inc. 2019b. "Qualcomm and Apple agree to drop all litigation." Press release.  
<https://www.qualcomm.com/news/releases/2019/04/qualcomm-and-apple-agree-drop-all-litigation>

Qualcomm Inc. 2019c. "Qualcomm Earnings Release Available on Company's Investor Relations Website." Press release, May 1. <https://investor.qualcomm.com/news-events/press-releases/news-details/2019/Qualcomm-Earnings-Release-Available-on-Companys-Investor-Relations-Website-05-01-2019/default.aspx>

Qualcomm Inc. 2020. Annual Report (Form 10-K). San Diego, CA: Qualcomm Inc. SEC.  
<https://www.sec.gov/Archives/edgar/data/804328/000172894920000067/qcom-20200927.htm>

Qualcomm Inc. 2022a. "Qualcomm Incorporated's Response to the Competition and Market Authority's Public Consultation on Retaining the EU Horizontal Block Exemption Regulations."  
<https://assets.publishing.service.gov.uk/media/625009d98fa8f54a841bbe09/Qualcomm.pdf>

Qualcomm Inc. 2022b. "Annual Report on Form 10-K for the Fiscal Year Ended September 25, 2022." SEC. <https://www.sec.gov/Archives/edgar/data/804328/000080432822000021/qcom-20220925.htm>

Qualcomm Inc. 2023. Annual Report (Form 10-K). San Diego, CA: Qualcomm Inc. SEC.  
<https://www.sec.gov/Archives/edgar/data/804328/000080432824000017/qcom-20231224.htm>

Qualcomm Inc. 2024. Annual Report (Form 10-K). San Diego, CA: Qualcomm Inc. SEC.  
<https://www.sec.gov/Archives/edgar/data/804328/000080432824000075/qcom-20240929.htm>

Quinn, Gene. 2014. "Freeman-Walter-Abele: A Tortured History of Software Eligibility." IPWatchdog 2  
<https://ipwatchdog.com/2014/12/02/freeman-walter-abele-a-tortured-history-of-software-eligibility/id=52271/>

Raab, Ben. 2025. "The \$1,999 Liberty Phone Is Made in America. Its Creator Explains How." Wall Street Journal. <https://www.wsj.com/tech/personal-tech/liberty-phone-purism-made-in-america-b4074c89>

Radauskas, Gintaras. 2025. Tech jobs in 2025: it's not looking good, and AI will make it even worse. Cybernews. <https://cybernews.com/editorial/tech-jobs-in-2025-ai-upskilling/>

Radford, Alec, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. "Language Models Are Unsupervised Multitask Learners." OpenAI Blog. [https://cdn.openai.com/better-language-models/language\\_models\\_are\\_unsupervised\\_multitask\\_learners.pdf](https://cdn.openai.com/better-language-models/language_models_are_unsupervised_multitask_learners.pdf)

Raff, Daniel, and Lawrence Summers. 1987. "Did Henry Ford Pay Efficiency Wages?" *Journal of Labor Economics* 5(4): S57-S86.

Raitano, Lucy, and Amanda Cooper. 2025. "Five Debt Hotspots in the AI Data Centre Boom." Reuters. Reprinted in U.S. News and World Report. <https://money.usnews.com/investing/news/articles/2025-11-05/five-debt-hotspots-in-the-ai-data-centre-boom>

Raman, Aneesh. 2025. "I'm a LinkedIn Executive. I See the Bottom Rung of the Career Ladder Breaking." *New York Times*. <https://www.nytimes.com/2025/05/19/opinion/linkedin-ai-entry-level-jobs.html>

Ramaswamy, Sridhar, and Paul Muret. 2018. "Introducing Google Marketing Platform." *Google Marketing Platform Blog*. <https://blog.google/products/marketingplatform/360/introducing-google-marketing-platform>

Rambus 2008. *Rambus Inc. v. Federal Trade Commission*. No. 07-1086. US Court of Appeals, D.C. Circuit, April 22. <https://law.justia.com/cases/federal/appellate-courts/cadc/07-1086/07-1086-1112217-2011-03-24.html>

Ramesh, Aditya, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021. Zero-Shot Text-to-Image Generation. *Proceedings of the 38th International Conference on Machine Learning*, PMLR 139:8821-8831.

Ramkumar, Sudha. 2025. "ISO and ASTM Standards: An Overview." *Microbe Investigations Switzerland*. <https://microbe-investigations.com/understanding-the-differences-between-iso-and-astm-standards/>

Rana, Anurag and Andrew Girard. 2024. Big tech 2025 capex may hit \$200 billion as gen-AI demand booms. *Bloomberg Intelligence*. <https://www.bloomberg.com/professional/insights/technology/big-tech-2025-capex-may-hit-200-billion-as-gen-ai-demand-booms/>

Rao, Rohit, Jason Liu, Robert Verkuil, Joshua Meier, John Canny, Pieter Abbeel, Tom Sercu, and Alexander Rives. 2021. MSA Transformer. Proceedings of the 38th International Conference on Machine Learning, PMLR 139:8844-8856.

Ratner, Alexander, Stephen Bach, Henry Ehrenberg, Jason Fries, Sen Wu, and Christopher Ré. 2017. Snorkel: Rapid training data creation with weak supervision. The VLDB Journal, 29(2): 709-730.

Reagan, Ronald. 1981a. Executive Order 12291: Federal Regulation. 46 FR: 13193.  
<https://www.federalregister.gov/executive-order/12291>

Reagan, Ronald. 1981b. "Statement on U.S. Trade Policy." Public Papers of the Presidents of the United States: Ronald Reagan (614–16). Washington, D.C.: US Government Printing Office.

Reagan, Ronald. 1982. Radio Address to the Nation on International Free Trade. Ronald Reagan Presidential Library. <https://www.presidency.ucsb.edu/documents/radio-address-the-nation-international-free-trade>

Reagan, Ronald. 1986. Radio Address to the Nation on Free and Fair Trade. Ronald Reagan Presidential Library. <https://www.reaganlibrary.gov/archives/speech/radio-address-nation-free-and-fair-trade-1>

Recon Analytics. 2018. How America’s 4G leadership propelled the U.S. economy. Recon Analytics. [https://api.ctia.org/wp-content/uploads/2018/04/Recon-Analytics\\_How-Americas-4G-Leadership-Propelled-US-Economy\\_2018.pdf](https://api.ctia.org/wp-content/uploads/2018/04/Recon-Analytics_How-Americas-4G-Leadership-Propelled-US-Economy_2018.pdf)

Register.com 2004. Register.com, Inc. v. Verio, Inc. 2004. No. 00-9596. U.S. Court of Appeals, Second Circuit, January 23. <https://law.justia.com/cases/federal/appellate-courts/F3/356/393/576556/>

Reiter. 1979. Reiter v. Sonotone Corp., 442 U.S. 330. U.S. Supreme Court, June 18.  
<https://supreme.justia.com/cases/federal/us/442/330>

Research and Markets. 2025. Automated Optical Inspection (AOI) Systems – Global Strategic Business Report. <https://www.researchandmarkets.com/reports/4804449/automated-optical-inspection-aoi-systems>

Reshoring Initiative. 2023. “2022 Reshoring and FDI Data Report.” Reshoring Initiative. [https://reshorenow.org/content/pdf/2022\\_Data\\_Report.pdf](https://reshorenow.org/content/pdf/2022_Data_Report.pdf)

Reuters. 2023. “EXPLAINER—What does Twitter ‘rate limit exceeded’ mean for users?” Reuters.  
<https://finance.yahoo.com/news/explainer-does-twitter-rate-limit-145553424.html>

Reuters. 2025a. “Elon Musk’s xAI buys new property in Memphis amid supercomputer expansion.”  
Reuters. <https://www.reuters.com/technology/artificial-intelligence/elon-musks-xai-buys-new-property-memphis-amid-supercomputer-expansion-2025-03-07/>

Reuters. 2025b. “‘Neocloud’ Crusoe secures \$11.6 billion funding for Texas data center. Reuters.  
<https://www.reuters.com/technology/openais-biggest-data-center-secures-116-billion-funding-wsj-reports-2025-05-21/>

Ricardo, David. 1817. *On the Principles of Political Economy and Taxation*. London: John Murray.

Richards, Porter. 2025. *The E-SIGN Act in the Age of Smart Contracts and AI: Challenges and Opportunities*. University of Cincinnati Intellectual Property and Computer Law Journal.  
<https://uciplj.org/2025/04/16/the-e-sign-act-in-the-age-of-smart-contracts-and-ai-challenges-and-opportunities/>

Richman, Sheldon. 1988. *The Reagan Record on Trade: Rhetoric vs. Reality*. Cato Institute Policy Analysis No. 107. <https://www.cato.org/sites/cato.org/files/pubs/pdf/pa107.pdf>

Richter, Felix. 2023. “Chart: The Trillion-Dollar App Economy.” Statista.  
<https://www.statista.com/chart/30364/sales-and-billings-facilitated-by-the-app-store/>

Riordan, Michael, and Lillian Hoddeson. 1997. *Crystal Fire*. New York: W. W. Norton and Company.

Rives, Alex, et al. 2021. Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *Proceedings of the National Academy of Sciences* 118(15): e2016239118.

Robotics Tomorrow. 2025. *Unitree Showcases Advanced High-Mobility Robots at CES 2025*.  
<https://www.roboticstomorrow.com/news/2025/01/08/unitree-showcases-advanced-high-mobility-robots-at-ces-2025/23818>

Rochet, Jean-Charles, and Jean Tirole. 2003. "Platform Competition in Two-Sided Markets." *Journal of the European Economic Association* 1(4): 990–1029.

Rochet, Jean-Charles, and Jean Tirole. 2006. Two-Sided Markets: A Progress Report. *RAND Journal of Economics* 37(3): 645–667.

Rockoff, Hugh. 1984. *Drastic Measures*. Cambridge: Cambridge University Press.

Rodrik, Dani. 2004. "Industrial Policy for the Twenty-First Century." CEPR Discussion Paper No. 4767. <https://cepr.org/publications/dp4767>.

Rodrik, Dani. 2006. "Goodbye Washington Consensus, Hello Washington Confusion?" *Journal of Economic Literature* 44(4): 973–987.

Rodrik, Dani. 2011. *The Globalization Paradox: Democracy and the Future of the World Economy*. New York: W. W. Norton.

Roessner, David, Jennifer Bond, Sumiye Okubo, and Mark Planting. 2013. "The economic impact of licensed commercialized inventions originating in university research." *Research Policy* 42(1): 23–34.

Rogers, Everett. 1962. *Diffusion of Innovations*. New York: Free Press of Glencoe.

Rogoff, Kenneth. 2025. *Our Dollar, Your Problem: An Insider's View of Seven Turbulent Decades of Global Finance, and the Road Ahead*. New Haven, CT: Yale University Press.

Romer, Paul M. 1990. Endogenous Technological Change. *Journal of Political Economy* 98(5, Part 2): S71–S102.

Romer, Paul. 1993. "Idea Gaps and Object Gaps in Economic Development." *Journal of Monetary Economics* 32(3): 543–573.

Roosevelt, Franklin. 1936. Acceptance Speech for the Renomination for the Presidency, Philadelphia, June 27, 1936. In Samuel I. Rosenman (Ed.), *The Public Papers and Addresses of Franklin D. Roosevelt* (Vol. 5). New York: Random House, 230-236.

Roser, Max, and Hannah Ritchie. 2013. "Technological Progress." *Our World in Data*, 2013. <https://ourworldindata.org/technological-progress>

Rotemberg, Julio, and Michael Woodford. 1992. "Oligopolistic Pricing and the Effects of Aggregate Demand on Economic Activity." *Journal of Political Economy* 100(6): 1153-1207.

Rothstein, Mark, and Stacey Tovino. 2019. "California Takes the Lead on Data Privacy Law." *Hastings Center Report* 49(5): 4–5.

Rouffet. 1998. In re Rouffet. No. 97-1404. US Federal Circuit Court, July 15.  
<https://law.justia.com/cases/federal/appellate-courts/F3/149/1350/560701/>

Rowe, Frederick. 1962. *Price Discrimination Under the Robinson-Patman Act*. Boston: Little, Brown and Company.

RSNA. 2025. "Integrating Imaging Tools Helps Radiology AI Deliver Real Value." *RSNA News*.  
<https://www.rsna.org/news/2025/june/interoperability-helps-radiology-ai-deliver-value>

Rubinfeld, Daniel. 2001. Antitrust Enforcement in Dynamic Network Industries. *Antitrust Bulletin* 46(3): 859-882.

Ruggie, John. 1982. International regimes, transactions, and change: Embedded liberalism in the postwar economic order. *International Organization* 36(2): 379–415.

Russell, Stuart, and Peter Norvig. 2021. *Artificial Intelligence: A Modern Approach*. 4th ed. Hoboken, NJ: Pearson.

Rysman, Marc. 2009. "The Economics of Two-Sided Markets." *Journal of Economic Perspectives* 23(3): 125–143.

Rysman, Marc, and Julian Wright. 2014. The Economics of Payment Card Fee Structure: What Drives Payment Card Rewards? *Review of Network Economics* 13(3): 229-259.

S&P Dow Jones Indices 2025. S&P 500 Market Capitalization Concentration Report.  
<https://www.spglobal.com/spdji/en/indices/equity/sp-500>

Sabre 2020. *United States v. Sabre Corp. et al.*, No. 1:19-cv-01548. U.S. District Court, D. Delaware, April 7.  
[https://appliedantitrust.com/14\\_merger\\_litigation/cases\\_doj/sabre2019/2\\_d\\_del/sabre\\_ddel\\_opinion2020\\_04\\_08public\\_version.pdf](https://appliedantitrust.com/14_merger_litigation/cases_doj/sabre2019/2_d_del/sabre_ddel_opinion2020_04_08public_version.pdf)

Sacra. 2025a. xAI revenue, valuation and funding. Sacra. <https://sacra.com/c/xai/>

Sacra. 2025b. Cohere revenue, valuation and funding. Sacra. <https://sacra.com/c/cohere/>

Saint-Martin, Léo. 2024. Economic Analysis of the EU and International Semiconductor Ecosystem. European Commission.  
<https://ec.europa.eu/research/participants/documents/downloadPublic?appId=PPGMS&documentIds=080166e50f159e79>

Sakagami, Ryo, Florian Lay, Andreas Dömel, Martin Schuster, Alin Albu-Schäffer, and Freck Stulp. 2023. Robotic world models—conceptualization, review, and engineering best practices. *Frontiers in Robotics and AI* 10:1253049.

Salemi, Alireza, Sheshera Mysore, Michael Bendersky, and Hamed Zamani. 2024. LaMP: When Large Language Models Meet Personalization. *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 7370–7392.

Salop, Steven. 2010. Question: What Is the Real and Proper Antitrust Welfare Standard—Answer: The Consumer Welfare Standard. *Loyola University Chicago Law Journal* 38:336.

Sampat, Bhaven. 2006. “Patenting and US Academic Research in the 20th Century.” *Research Policy* 35(6): 772–789.

Samuelson, Pamela, and Suzanne Scotchmer. 2002. "The Law and Economics of Reverse Engineering." *Yale Law Journal* 111(7): 1575–1663.

Sathish, Madhumitha, and Peter Rutten. 2024. Worldwide Artificial Intelligence Infrastructure Forecast, 2024–2028. IDC #US52597124. <https://my.idc.com/getdoc.jsp?containerId=US52597124>

Sato, Kaz, and Cliff Young. 2017. An in-depth look at Google’s first Tensor Processing Unit (TPU). Google Cloud Blog. <https://cloud.google.com/blog/products/ai-machine-learning/an-in-depth-look-at-googles-first-tensor-processing-unit-tpu>

Savino, Vito. 2025. Edge computing and data centers in the age of AI. *Data Center Dynamics*.  
<https://www.datacenterdynamics.com/en/opinions/edge-computing-and-data-centers-in-the-age-of-ai/>

Scale AI. 2023. Scale Data Engine: Data Annotation, Collection, and Curation Platform. ScaleAI.  
<https://scale.com/data-engine>

ScaleAI. 2025. Scale AI Announces Next Phase of Company's Evolution. ScaleAI.  
<https://scale.com/blog/scale-ai-announces-next-phase-of-company-evolution>

Schaake, Marietje. 2021. "Regulating Technology with Marietje Schaake." WIRED.  
<https://www.wired.com/story/regulating-technology>

Schacht, Wendy H., and John Thomas. 2006. Patent Reform: Innovation Issues. CRS Report  
RL32996. [https://www.everycrsreport.com/files/20061207\\_RL32996\\_3ee4011c8beb27223894b69fefc42a4c9b39377e.pdf](https://www.everycrsreport.com/files/20061207_RL32996_3ee4011c8beb27223894b69fefc42a4c9b39377e.pdf)

Schaeffer, Rylan, Brando Miranda, and Sanmi Koyejo. 2023. Are Emergent Abilities of Large Language Models a Mirage? *Advances in Neural Information Processing Systems* 36: 27159–27182.

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

Schmalensee, Richard. 2000. Antitrust Issues in Schumpeterian Industries. *American Economic Review* 90(2): 192–196.

Schmalensee, Richard. 2002. Payment Systems and Interchange Fees. *Journal of Industrial Economics* 50(2): 103–122.

Schmalensee, Richard. 1982. "Product Differentiation Advantages of Pioneering Brands." *American Economic Review* 72(3): 349-365.

Schmidt, Douglas. 2018. Google Data Collection. Nashville, TN: Vanderbilt University.  
<https://www.dre.vanderbilt.edu/~schmidt/PDF/google-data-collection.pdf>

Schneider, Bruno. 2026. Catch-up with the US or prosper below the tech frontier? An EU artificial intelligence strategy. Bruegel Policy Brief. <https://www.bruegel.org/policy-brief/catch-us-or-prosper-below-tech-frontier-eu-artificial-intelligence-strategy>

Schrage, Michael, and David Kiron. 2025. Philosophy eats AI. MIT Sloan Management Review.  
<https://sloanreview.mit.edu/article/philosophy-eats-ai/>

Schumpeter, Joseph. 1942. *Capitalism, Socialism and Democracy*. New York: Harper & Brothers.

Schwartz, Alan, and Robert Scott. 2002. "Contract Theory and the Limits of Contract Law." *Yale Law Journal* 113(3): 541-619.

Schwartz, Paul, and Edward Janger. 2002. The Gramm-Leach-Bliley Act, Information Privacy, and the Limits of Default Rules. *Minnesota Law Review* 86: 1217-1335.

Schwinn 1967. United States v. Arnold, Schwinn and Co., 388 U.S. 365.  
<https://supreme.justia.com/cases/federal/us/388/365>

Scoble, Robert. 2010. Walking through computer history with Apple cofounder Woz at the Computer History Museum. YouTube. <https://www.youtube.com/watch?v=hsB8Hxnb52o>

SCoJ. 1980. Patent and Trademark Law Amendments Act of 1980. Senate Report No. 96-480. 96th Congress, 2nd Session. Washington, DC: U.S. Government Printing Office.

Scribner, Marc. 2013. Slow Train Coming? Misguided Economic Regulation of U.S. Railroads, Then and Now. Issue Analysis No. 1. Competitive Enterprise Institute.  
<https://cei.org/sites/default/files/Marc%20Scribner%20-%20Slow%20Train%20Coming.pdf>

Sculley, D., Gary Holt, Daniel Golovin, Eugene Davydov, Todd Phillips, Dietmar Ebner, Vinay Chaudhary, Michael Young, Jean-Francois Crespo, and Dan Dennison. 2015. Hidden Technical Debt in Machine Learning Systems. *Advances in Neural Information Processing Systems* 28:2503-2511.

Seetharaman, Deepa, and Sriram Akash. 2025. “Nvidia to Invest Up to \$100 Billion in OpenAI, Linking Two Artificial Intelligence Giants.” Reuters. <https://www.reuters.com/business/nvidia-invest-100-billion-openai-2025-09-22/>

Seila Law. 2020. Seila Law LLC v. Consumer Financial Protection Bureau. 2020. No. 19-7. U.S. Supreme Court, June 29. <https://supreme.justia.com/cases/federal/us/591/19-7/>

Seliger, Yuri. 2025. Big Tech Financing Capacity and Balance Sheet Analysis. Bank of America Global Research, November 2025. <https://investor.bankofamerica.com>

Sell, Susan. 2003. Private Power, Public Law. Cambridge: Cambridge University Press.

SIA. 2020. 2020 State of the U.S. Semiconductor Industry Report. Washington, DC: Semiconductor Industry Association. <https://www.semiconductors.org/wp-content/uploads/2020/07/2020-SIA-State-of-the-Industry-Report-FINAL-1.pdf>

SIA. 2021. 2021 State of the U.S. Semiconductor Industry Report. Washington, DC: Semiconductor Industry Association. <https://www.semiconductors.org/wp-content/uploads/2021/09/2021-SIA-State-of-the-Industry-Report.pdf>

Serant, Claire. 2000. "Why Motorola Chose Outsourcing Route." EE Times. <https://www.eetimes.com/why-motorola-chose-outsourcing-route/>

Setser, Brad. 2020. "When the Services Trade Data Tells You More About Tax Avoidance Than About Actual Trade." Council on Foreign Relations Blog. <https://www.cfr.org/blog/when-services-trade-data-tells-you-more-about-tax-avoidance-about-actual-trade>

Sevilla, Jaime, Lennart Heim, Anson Ho, Tamay Besiroglu, Marius Hobbhahn and Pablo Villalobos. 2022. Compute Trends Across Three Eras of Machine Learning. arXiv preprint arXiv:2202.05924. <https://arxiv.org/abs/2202.05924>

Shakudo. 2025. "12 MLOps Best Practices Every Enterprise Needs in 2025." Shakudo Blog. <https://www.shakudo.io/blog/mlops-best-practices-enterprise-2025>

Shamsuddoha, Mohammad, Tasnuba Nasir, and Mohammad Saifuddoha Fawaaz. 2025. "Humanoid Robots like Tesla Optimus and the Future of Supply Chains: Enhancing Efficiency, Sustainability, and Workforce Dynamics." *Automation* 6(1): 9.

Shankland, Stephen. 2004. "Intel Shifts 64-bit Emphasis." CNET. <https://www.cnet.com/tech/tech-industry/intel-shifts-64-bit-emphasis>

Shapira, Philip. 2001. "US Manufacturing Extension Partnerships." *Research Policy* 30(6): 977–992.

Shapiro, Carl, and Hal Varian. 1999. *Information Rules: A Strategic Guide to the Network Economy*. Boston, MA: Harvard Business School Press.

Shapiro, Carl. 1989. "Theories of Oligopoly Behavior." In *Handbook of Industrial Organization*, Vol. 1, edited by Richard Schmalensee and Robert Willig. Amsterdam: Elsevier, 329-414.

Shapiro, Carl. 1996. "Mergers with Differentiated Products." *Antitrust* 10(2): 23–30.

Shapiro, Carl. 2010a. "The Role of Antitrust in Preventing Patent Holdup." *Antitrust Law Journal* 82(2): 201–275.

Shapiro, Carl. 2010b. "The 2010 Horizontal Merger Guidelines: From Hedgehog to Fox in Forty Years." *Antitrust Law Journal* 77: 49–168.

Shapiro, Carl. 2018. "Competition and the Dynamics of Industrial Structure." In *Innovation Policy and the Economy*, Vol. 18, edited by Josh Lerner and Scott Stern. Chicago: University of Chicago Press, 23–63.

Shapiro, Steven. 1993. "Preliminary Injunction Motions in Patent Litigation." *IDEA: The Journal of Law and Technology* 33: 323–348.

Sharma, Debendra Das and Ravi Mahajan. 2024. Advanced packaging of chiplets for future computing needs. *Nature Electronics* 7(3): 181–192.

Shelanski, Howard. 2011. Enforcing Competition During an Economic Emergency. *Minnesota Law Review* 95:1459.

Shelanski, Howard, and Gregory Sidak. 2001. "Antitrust Divestiture in Network Industries." *University of Chicago Law Review* 68(1): 1–99.

Shepherd, George. 1996. "Fierce Compromise." *Northwestern University Law Review* 90(4): 1557–1683.

Shojaee, Parshin, Iman Mirzadeh, Keivan Alizadeh, Maxwell Horton, Samy Bengio, and Mehrdad Farajtabar. 2025. "The Illusion of Thinking: Understanding the Strengths and Limitations of Reasoning Models via the Lens of Problem Complexity." *airXiv preprint* airXiv: 2506.06941.

<https://arxiv.org/pdf/2506.06941>

Shumailov, Iliia, Yarin Gal, Vitaly Kurin, Nicholas Lane, and Ross Anderson. 2023. The Curse of Recursion: Training on Generated Data Makes Models Forget. *arXiv preprint*: arXiv.2305.17493.

<https://doi.org/10.48550/arXiv.2305.17493>

Sidak, Gregory and David Teece. 2009. "Dynamic Competition in Antitrust Law." *Journal of Competition Law and Economics* 5(4): 581–631.

Sidak, Gregory. 2013. "The Meaning of FRAND, Part I: Royalties." *Journal of Competition Law and Economics* 9(4): 931–1055.

Singhal, Karan, et al. 2023. Large language models encode clinical knowledge. *Nature* 620: 172-180.

Sites, Richard, and Richard Witek. 1993. Alpha AXP Architecture Reference Manual. 2nd ed. Burlington, MA: Digital Press.

Skibba, Ramin, Subhasish Mitra, and Philip Wong. 2015. "Skyscraper-Style Chip Design Boosts Performance by Factor of a Thousand." Stanford Report. <https://biox.stanford.edu/highlight/skyscraper-style-chip-design-boosts-performance-factor-thousand>

Skidmore et al. 1944. Skidmore et al. v. Swift and Co., No. 12, U.S. Supreme Court, December 4. <https://tile.loc.gov/storage-services/service/ll/usrep/usrep323/usrep323134/usrep323134.pdf>

Slack Developers. 2025. Rate limit changes for non-Marketplace apps. Slack Developers. <https://docs.slack.dev/changelog/2025/05/29/rate-limit-changes-for-non-marketplace-apps/>

Smith, Brad 2025. The Golden Opportunity for American AI. Microsoft Blog. <https://blogs.microsoft.com/on-the-issues/2025/01/03/the-golden-opportunity-for-american-ai/>

Smith, Ben. 2018. "Project Strobe: Protecting Your Data, Improving Our Third-Party APIs, and Sunsetting Consumer Google+." Google Blog. <https://blog.google/technology/safety-security/project-strobe/>

Smith, Bob. 2024. Semiconductor Engineering. "The State of the EDA Industry in 2024." <https://semiengineering.com/the-state-of-the-eda-industry-in-2024/>

Smith, Steven, Mike Hanlon, and Robert Bailey. 1992. Power Management for a Laptop Computer with Slow and Sleep Modes. (US Patent No. 5,167,024). USPTO. <https://patents.google.com/patent/US5167024A/en>

Snowflake.help. 2025. How Snowflake's Open Semantic Interchange is revolutionizing AI data readiness. Snowflake.help. <https://snowflake.help/how-snowflakes-open-semantic-interchange-is-revolutionizing-ai-data-readiness/>

Socony-Vacuum Oil. 1940. United States v. Socony-Vacuum Oil Co., 310 U.S. 150. <https://supreme.justia.com/cases/federal/us/310/150>

SoftBank Group Corp. 2025. SoftBank Group Report 2025 (Annual Report). <https://group.softbank/en/news/info/20250728>

Sokol, Daniel, and Roisin Comerford. 2020. Does Antitrust Have a Role to Play in Regulating Big Data? In *The Cambridge Handbook of Antitrust, Intellectual Property, and High Tech*, edited by Roger Blair and D. Daniel Sokol. Cambridge: Cambridge University Press, 271–291.

Solove, Daniel, and Woodrow Hartzog. 2014. "The FTC and the New Common Law of Privacy." *Columbia Law Review* 114(3): 583–676.

Solow, Robert. 1988. Growth Theory and After. *The American Economic Review* 78(3): 307–317.

Solow, Robert. 1957. "Technical Change and the Aggregate Production Function." *Review of Economics and Statistics* 39(3): 312-320.

Solow, Robert. 1987. "We'd Better Watch Out." *The New York Times Book Review*, July 12, 1987, p. 36.

Solsman, Joan. 2018. "YouTube's AI Is the Puppet Master over Most of What You Watch." CNET. <https://www.cnet.com/tech/services-and-software/youtube-ces-2018-neal-mohan/>

Son, Hugh. 2025. "Goldman Sachs Launches AI Assistant as the Tech Sweeps Banking." CNBC. <https://www.cnbc.com/2025/01/21/goldman-sachs-launches-ai-assistant.html>

Song, Seagull and Mo Wang. 2024. "China's First Generative AI Copyright Infringement Case: Ultraman." King and Wood Mallesons. <https://www.kwm.com/cn/en/insights/latest-thinking/china-s-first-case-on-aigc-output-infringement-ultraman.html>

Soper, Taylor 2024. Databricks, set to raise massive \$10B funding round, looks to grow Seattle-area footprint. GeekWire. <https://www.geekwire.com/2024/databricks-set-to-raise-massive-10b-funding-round-looks-to-grow-seattle-area-footprint>

Specht et al. 2022. Christopher Specht et al. v. Netscape Communications Corporation and America Online, Inc., Nos. 01-7870, 01-7860, U.S. Court of Appeals, Second Circuit, October 1. [https://scholar.google.com/scholar\\_case?case=9587085159184835436](https://scholar.google.com/scholar_case?case=9587085159184835436)

Spectrum Sports. 1993. *Spectrum Sports, Inc. v. McQuillan*, 506 U.S. 447, U.S. Supreme Court, January 13. <https://supreme.justia.com/cases/federal/us/506/447>

Spegele, Brian. 2025. "The Fortress That China Built for Its Battle With America," *Wall Street Journal*. <https://www.wsj.com/world/china/china-us-technology-economy-advancements-bb8d7439>

Spence, Michael. 2021. "Some Thoughts on the Washington Consensus and Subsequent Global Development Experience." *Journal of Economic Perspectives* 35(3): 67–82.

Spengler, Joseph. 1950. "Vertical Integration and Antitrust Policy." *Journal of Political Economy* 58(4): 347–352.

Spulber, Daniel. 2015. "How Patents Provide the Foundation of the Market for Inventions." *Journal of Competition Law and Economics* 11(2): 271-316.

Srinivasan, Dina. 2019. "The Antitrust Case Against Facebook." *Berkeley Business Law Journal* 16(1): 39–98.

Srivathsan, Bhargos, Marc Sorel, and Pankaj Sachdeva, with Arjita Bhan, Haripreet Batra, Raman Sharma, Rishi Gupta, and Surbhi Choudhary. 2024. "Power to the People: The Gigawatt-Scale Challenge of AI Data Centers." *McKinsey & Company Insights*. <https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/ai-power-expanding-data-center-capacity-to-meet-growing-demand#/>

Stainless. 2025. Production-ready MCP servers optimized for agentic coding and context limits. <https://www.stainless.com/products/mcp>

Stalk, George, Philip Evans, and Lawrence Shulman. 1992. "Competing on Capabilities: The New Rules of Corporate Strategy." *Harvard Business Review* 70(2): 57-69.

Standard Oil. 1911. *Standard Oil Co. of New Jersey v. United States*, 221 U.S. 1, U.S. Supreme Court, May 15. <https://supreme.justia.com/cases/federal/us/221/1/>

Standard Oil. 1949. *Standard Oil Co. of California v. United States*, No. 279. US Supreme Court, June 13. <https://supreme.justia.com/cases/federal/us/337/293>

Staples. 1997. *FTC v. Staples, Inc.*, 970 F. Supp. 1066 (D.D.C.). <https://law.justia.com/cases/federal/district-courts/FSupp/970/1066/1639260>

Stasik, Eric. 2010. "Royalty Rates and Licensing Strategies for Essential Patents on LTE (4G) Telecommunication Standards." *Les Nouvelles* 45(3): 114–119.

StatCounter. 2024. Mobile operating system market share worldwide. StatCounter Global Stats.  
<https://gs.statcounter.com/os-market-share/mobile/worldwide>

StatCounter. 2025. "Browser Market Share Worldwide." StatCounter Global Stats.  
<https://gs.statcounter.com/browser-market-share/>

State Street Bank. 1998. State Street Bank and Trust Co. v. Signature Financial Group, Inc. No. 96-1327. US Federal Circuit Court, July 23. <https://law.justia.com/cases/federal/appellate-courts/F3/149/1368/560460>

State Street Bank. 1998. State Street Bank and Trust Co. v. Signature Financial Group, Inc. No. 96-1327. US Federal Circuit Court, July 23. <https://law.justia.com/cases/federal/appellate-courts/F3/149/1368/560460/>

Statista. 2024a. "Volume of Data/Information Created, Captured, Copied, and Consumed Worldwide from 2010 to 2020, with Forecasts from 2021 to 2025." Statista.  
<https://www.statista.com/statistics/871513/worldwide-data-created>

Statista. 2024b. Leading ad-supported digital video platforms in the United States, by share of viewing time. Statista. <https://www.statista.com/statistics/276623/share-of-video-streaming-usage-in-the-us-by-service/>

Stein, David, Kaitlin Johns, Maryam Fatima, and Shari Allen. 2024. U.S. International Services: Trade in Services in 2023 and Services Supplied Through Affiliates in 2022. U.S. Bureau of Economic Analysis.  
<https://apps.bea.gov/scb/issues/2024/10-october/1024-international-services.htm>

Stern, Jonathan. 2001. The Electronic Signatures in Global and National Commerce Act. Berkeley Technology Law Journal 16(1): 391-414.

Stevens, Ashley. 2004. "The Enactment of Bayh-Dole." Journal of Technology Transfer 29(1): 93–99.

Stigler, George. 1968. The Organization of Industry. Homewood, IL: Richard D. Irwin.

Stigler, George. 1971. "The Theory of Economic Regulation." Bell Journal of Economics and Management Science 2(1): 3-21.

Stiglitz, Joseph. 2002. Globalization and Its Discontents. New York: W.W. Norton and Company.

Stiglitz, Joseph. 2006. *Making Globalization Work*. New York: W.W. Norton and Company.

Stocking, George, and Willard Mueller. 1955. The Cellophane Case and the New Competition. *The American Economic Review* 45(1): 29-63.

Stranger, Greg, and Shane Greenstein. 2007. "Pricing at the On-Ramp to the Internet: Price Indexes for ISPs during the 1990s." In *Hard-to-Measure Goods and Services: Essays in Honor of Zvi Griliches*, edited by Ernst Berndt and Charles Hulten, Chicago: University of Chicago Press, 197–234.

Strauss, Robert. 1979. "Symposium—The Tokyo Round: It's Meaning and Effect, Introduction." *Georgia Journal of International and Comparative Law* 9(2): 151–175.

Strubell, Emma, Ananya Ganesh, and Andrew McCallum. 2019. Energy and Policy Considerations for Deep Learning in NLP. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3645–3650.

Studiengesellschaft Kohle 1981. *U.S. v. Studiengesellschaft Kohle et al.*, No. 79-1634, U.S. Court of Appeals, District of Columbia Circuit, December 11.

[https://scholar.google.com/scholar\\_case?case=867784977804525218](https://scholar.google.com/scholar_case?case=867784977804525218)

Stuecke, Jan. 2025. Top 7 Open-Source Frameworks for Federated Learning. Apheris.

<https://www.apheris.com/resources/blog/top-7-open-source-frameworks-for-federated-learning>

Sukumar, Vinesh. 2025. "Shifting AI inference from the cloud to your phone can reduce AI costs." Qualcomm OnQ Blog. <https://www.qualcomm.com/news/onq/2025/09/shifting-ai-inference-from-the-cloud-to-your-phone-can-reduce-ai-costs>

Suleyman, Mustafa. 2023. *The Coming Wave: Technology, Power, and the Twenty-first Century's Greatest Dilemma*. New York: Crown. <https://the-coming-wave.com>

Sullivan and Cromwell LLP. 2020. "Ninth Circuit Holds That Qualcomm's Patent Licensing Program Does Not Violate U.S. Antitrust Law." Sullivan and Cromwell.

[https://www.sullcrom.com/SullivanCromwell/\\_Assets/PDFs/Memos/sc-publication-ninth-circuit-holds-qualcomm-patent-licensing-program-does-not-violate-us-antitrust-law.pdf](https://www.sullcrom.com/SullivanCromwell/_Assets/PDFs/Memos/sc-publication-ninth-circuit-holds-qualcomm-patent-licensing-program-does-not-violate-us-antitrust-law.pdf)

Sullivan, Lawrence, and Warren Grimes. 2006. *The Law of Antitrust*. 2nd ed. St. Paul: West Academic Publishing.

Sullivan, Nolan. 2025. In Depth: Speakeasy vs Stainless. Speakeasy Blog.  
<https://www.speakeasy.com/blog/speakeasy-vs-stainless>

Sunstein, Cass. 2002. *The Cost-Benefit State: The Future of Regulatory Protection*. Chicago, IL: American Bar Association.

Sunstein, Cass. 2013. "The Office of Information and Regulatory Affairs." *Harvard Law Review* 126(7): 1838–1878.

Surapaneni, Rao, Miku Jha, Michael Vakoc, and Todd Segal. 2025. Announcing the Agent-to-Agent (A2A) protocol: A new era of agent interoperability. Google Developers Blog.  
<https://developers.googleblog.com/en/a2a-a-new-era-of-agent-interoperability/>

Svensson, Philip. 2024. Qualcomm: Quality communication and semiconductor innovation. Quatr.  
<https://quatr.com/insights/company-research/qualcomm-quality-communication-and-semiconductor-innovation>

SWIFT. 2025. Momentum builds as industry advances ISO 20022 adoption. <https://www.swift.com/news-events/news/momentum-builds-industry-advances-iso-20022-adoption>

Synergy Research Group. 2024. "Cloud Infrastructure Market Share: AWS, Microsoft and Google Control 67%." Synergy Research Group. <https://www.prnewswire.com/news-releases/cloud-market-gets-its-mojo-back-ai-helps-push-q4-increase-in-cloud-spending-to-new-highs-302051366.html>

Syufy Enterprises. 1990. *United States v. Syufy Enterprises*, 903 F.2d 659, U.S. Court of Appeals for the Ninth Circuit, May 9. <https://law.justia.com/cases/federal/appellate-courts/F2/903/659/435777>

Syverson, Chad. 2011. What Determines Productivity? *Journal of Economic Literature* 49(2): 326–365.

Sze, Simon, and Kwok Ng. 2006. *Physics of Semiconductor Devices*. 3rd ed. Hoboken, NJ: John Wiley and Sons.

Tan, Lisa, and Ken Huang. 2025. "The AI Agent Economy." In *Progress in Information Systems*, edited by Jörg Becker and Ralf Knackstedt. Cham: Springer Nature: 99–134.

Tao, Fei, Qinglin Qi, Ang Liu, and Andrew Kusiak. 2018. Data-driven smart manufacturing. *Journal of Manufacturing Systems* 48(Part C): 157–169.

Tassey, Gregory. 2000. "Standardization in Technology-Based Markets." *Research Policy* 29(4-5): 587–602.

Teare, Gené. 2025. Startup Funding Regained Its Footing In 2024 As AI Became the Star of the Show. *Crunchbase News*, January 7. <https://news.crunchbase.com/venture/global-funding-data-analysis-ai-eoy-2024/>

Techlevated. 2023. FinFET vs MOSFET (Planar) Transistors in Chip Design. *Techlevated*. <https://techlevated.com/finfet-vs-mosfet-planar-transistor/>

Teece, David. 1982. "Towards an Economic Theory of the Multiproduct Firm." *Journal of Economic Behavior and Organization* 3(1): 39-63.

Teece, David. 1984. "Economic Analysis and Strategic Management." *California Management Review* 26(3): 87–110.

Tellex, Stefanie, Nakul Gopalan, Hadas Kress-Gazit, and Cynthia Matuszek. 2020. Robots that use language. *Annual Review of Control, Robotics, and Autonomous Systems* 3: 25–55.

Temin, Peter and Louis Galambos. 1987. *The Fall of the Bell System*. New York: Cambridge University Press.

Temkin, Marina. 2025. Perplexity reportedly raised \$200M at \$20B valuation. *TechCrunch*. <https://techcrunch.com/2025/09/10/perplexity-reportedly-raised-200m-at-20b-valuation/>

Tempus. 2025. "Clinical trial matching (Tempus TIME)." *Tempus*. <https://www.tempus.com/oncology/clinical-trial-matching/>

Tene, Omer, Bethany Withers, and Tayjus Surampudi. 2025. California Moves to Regulate Frontier AI With a Focus on Catastrophic Risk. *Goodwin Law Alert*. <https://www.goodwinlaw.com/en/insights/publications/2025/11/alerts-technology-aiml-california-moves-to-regulate-frontier-ai-with-a-focus-on-catastrophic-risk>

Teti, Feodora. 2020. 30 Years of Trade Policy: Evidence from 5.7 Billion Tariffs. ifo Working Papers No. 334. <https://www.ifo.de/DocDL/wp-2020-334-teti-trade-policy-tariffs.pdf>

Texas Instruments. 2021. Developing the First ICs to Orbit Earth. TI Blog. <https://www.ti.com/about-ti/newsroom/company-blog/developing-the-first-ics-to-orbit-earth.html>

The Economist. 2002 Innovation's golden goose. The Economist. December 14: 3.

The Economist. 2015a. A tightening grip. <https://www.economist.com/briefing/2015/03/12/a-tightening-grip>

The Economist. 2015b. Made in China? <https://www.economist.com/leaders/2015/03/12/made-in-china>

The Economist. 2017. "The World's Most Valuable Resource Is No Longer Oil, but Data." The Economist. <https://www.economist.com/leaders/2017/05/06/the-worlds-most-valuable-resource-is-no-longer-oil-but-data>

The Economist. 2020. Are data more like oil or sunlight? <https://www.economist.com/special-report/2020/02/20/are-data-more-like-oil-or-sunlight>

The Information Staff. 2024. 7 Charts That Explain 2024. The Information. <https://www.theinformation.com/articles/7-charts-that-explain-2024>

Thompson, Alan. 2024. Models: Gemini. LifeArchitect.ai. <https://lifearchitect.ai/gemini/>

Tierno, Paul. 2024. Artificial intelligence and machine learning in financial services. CRS Report R47997. <https://www.congress.gov/crs-product/R47997>

Tilley, Aaron. 2024. "How Apple Fell Behind in the AI Arms Race." The Wall Street Journal. <https://www.wsj.com/tech/ai/apple-ai-siri-development-behind-9ea65ee8>

Tirole, Jean. 1988. The Theory of Industrial Organization. Cambridge, MA: MIT Press.

Topco Associates. 1972. United States v. Topco Associates, Inc., 405 U.S. 596. <https://supreme.justia.com/cases/federal/us/405/596>

Torchia, Marcus and Shirer, Matt 2024. New IDC Spending Guide Forecasts Edge Computing Investments to Reach \$232 Billion in 2024. IDC.

<https://www.businesswire.com/news/home/20240314632794/en/New-IDC-Spending-Guide-Forecasts-Edge-Computing-Investments-Will-Reach-%24232-Billion-in-2024>

Touvron, Hugo, et al. 2023. LLaMA: Open and Efficient Foundation Language Models. arXiv preprint arXiv:2302.13971. <https://arxiv.org/abs/2302.13971>

Train, Kenneth. 2009. Discrete Choice Methods with Simulation. 2nd ed. Cambridge: Cambridge University Press.

Trenton Potteries. 1927. United States v. Trenton Potteries Co., 273 U.S. 392, U.S. Supreme Court, February 21. <https://supreme.justia.com/cases/federal/us/273/392>

Trump, Donald. 2017. " Executive Order 13771—Reducing Regulation and Controlling Regulatory Costs. 82 FR: 9339. <https://www.govinfo.gov/app/details/FR-2017-02-03/2017-02451>

Trump, Donald. 2025a. America’s AI Action Plan—Winning the Race. White House. <https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf>

Trump, Donald. 2025b. Executive Order 14365. "Ensuring a National Policy Framework for Artificial Intelligence." Federal Register 90, no. 240: 58499.

<https://www.federalregister.gov/documents/2025/12/16/2025-23092/ensuring-a-national-policy-framework-for-artificial-intelligence>

Trump, Donald. 2025c. "Executive Order 14257—Regulating Imports With a Reciprocal Tariff to Rectify Trade Practices That Contribute to Large and Persistent Annual United States Goods Trade Deficits."

<https://www.whitehouse.gov/presidential-actions/2025/04/regulating-imports-with-a-reciprocal-tariff-to-rectify-trade-practices-that-contribute-to-large-and-persistent-annual-united-states-goods-trade-deficits/>

Tschand, Arya, et al. 2024. “MLPerf Power: Benchmarking the Energy Efficiency of Machine Learning Systems from Microwatts to Megawatts for Sustainable AI.” arXiv preprint arXiv: 2410.12032.

<https://arxiv.org/abs/2410.12032>

TSMC. 2023a. Annual Report. Hsinchu, Taiwan: TSMC (Investor Relations).

[https://www.tsmc.com/english/investorRelations/annual\\_reports](https://www.tsmc.com/english/investorRelations/annual_reports)

TSMC. 2023b. Sustainability Report. Hsinchu, Taiwan: TSMC (ESG Reporting). <https://esg.tsmc.com/en/>

TSMC. 2024. TSMC Annual Report 2024. TSMC. <https://investor.tsmc.com/sites/ir/annual-report/2024/2024%20Annual%20Report.E.pdf>

Tsui, David, and Andrew Chang. 2025. Research Update: NVIDIA Corp. Outlook Revised to Positive. S&P Global Ratings. <https://www.spglobal.com/ratings/en/regulatory/article/-/view/type/HTML/id/3463382>

Turner, Donald. 1949. The Definition of Agreement Under the Sherman Act: Conscious Parallelism and Refusals to Deal. *Harvard Law Review* 62(4): 661-701.

Tyman, Annette 2024. Mobley v. Workday: Court holds AI service providers could be directly liable for employment discrimination. Seyfarth Shaw LLP. <https://www.seyfarth.com/news-insights/mobley-v-workday-court-holds-ai-service-providers-could-be-directly-liable-for-employment-discrimination-under-agent-theory.html>

US Census Bureau. n.d. Exports from Manufacturing Establishments: Historical Data. DoC. [https://www.census.gov/manufacturing/exports/historical\\_data/index.html](https://www.census.gov/manufacturing/exports/historical_data/index.html)

US Congress. 1890. Sherman Act. Ch. 647, 26 Stat. 209. <https://www.govinfo.gov/content/pkg/COMPS-3055/pdf/COMPS-3055.pdf>

US Congress. 1914a. Clayton Act. Ch. 323, 38 Stat. 730. <https://www.govinfo.gov/content/pkg/COMPS-3049/pdf/COMPS-3049.pdf>

US Congress. 1914b. Federal Trade Commission Act. Ch. 311, 38 Stat. 717. <https://www.ftc.gov/legal-library/browse/statutes/federal-trade-commission-act>

US Congress. 1934. Communications Act of 1934, Pub. L. No. 73–416. <https://www.govinfo.gov/content/pkg/STATUTE-48/pdf/STATUTE-48-Pg1064.pdf>

US Congress. 1936. Robinson-Patman Antidiscrimination Act. Ch. 592, 49 Stat. 1526. <https://www.govinfo.gov/content/pkg/COMPS-12152/pdf/COMPS-12152.pdf>

US Congress. 1946. Administrative Procedure Act. Pub. L. No. 79-404. <https://www.govinfo.gov/content/pkg/STATUTE-60/pdf/STATUTE-60-Pg237.pdf>

US Congress. 1950. Celler-Kefauver Act. Pub. L. No. 81-899, 64 Stat. 1125.  
<https://www.govinfo.gov/content/pkg/STATUTE-64/pdf/STATUTE-64-Pg1125.pdf>

US Congress. 1952. Patent Act of 1952. Pub. L. No. 82-593.  
<https://www.govinfo.gov/content/pkg/STATUTE-66/pdf/STATUTE-66-Pg792.pdf>

US Congress. 1965. Food and Agriculture Act of 1965. Pub. L. No. 89-321.  
<https://nationalaglawcenter.org/wp-content/uploads/assets/farmbills/1965.pdf>

US Congress. 1976a. Hart-Scott-Rodino Antitrust Improvements Act. Pub. L. No. 94-435, 90 Stat. 1383.  
<https://www.govinfo.gov/content/pkg/STATUTE-90/pdf/STATUTE-90-Pg1383.pdf>

US Congress. 1976b. US Copyright Act. Pub. L. 94-553. <https://www.copyright.gov/history/pl94-553.pdf>

US Congress. 1978. Public Utility Regulatory Policies Act of 1978. Pub. L. No. 95-617.  
<https://www.congress.gov/bill/95th-congress/house-bill/4018>

US Congress. 1979a. The 1979 Joint Economic Report. 96th Congress, 1st Session.  
[https://www.jec.senate.gov/reports/96th%20Congress/The%201979%20Joint%20Economic%20Report%20\(930\).pdf](https://www.jec.senate.gov/reports/96th%20Congress/The%201979%20Joint%20Economic%20Report%20(930).pdf)

US Congress. 1979b. Trade Agreements Act of 1979. Pub. L. No. 96-39.  
<https://www.congress.gov/bill/96th-congress/house-bill/4537>

US Congress. 1980a. Stevenson-Wydler Technology Innovation Act of 1980, Pub. L. No. 96-480.  
<https://uscode.house.gov/statutes/pl/96/480.pdf>

US Congress. 1980b. Patent and Trademark Law Amendments Act (Bayh-Dole Act), Pub. L. No. 96-517.  
<https://www.govinfo.gov/content/pkg/STATUTE-94/pdf/STATUTE-94-Pg3015.pdf>

US Congress. 1981. H.R. Rep. No. 97-312, 97th Cong., 1st Sess. (Federal Courts Improvement Act of 1981). <https://www.congress.gov/congressional-report/97th-congress/house-report/312>

US Congress. 1982. Federal Courts Improvement Act of 1982. Pub. L. No. 97-164.  
<https://www.congress.gov/bill/97th-congress/house-bill/4482>

US Congress. 1984a. Semiconductor Chip Protection Act of 1984, Pub. L. No. 98-620.  
<https://www.congress.gov/98/statute/STATUTE-98/STATUTE-98-Pg3335.pdf>

US Congress. 1984b. National Cooperative Research Act of 1984. Pub. L. No. 98-462.  
<https://www.congress.gov/bill/98th-congress/senate-bill/1841>

U.S. Congress. 1986a. Federal Technology Transfer Act of 1986. Pub. L. No. 99-502.  
<https://www.congress.gov/bill/99th-congress/house-bill/3773>

US Congress. 1986b. Computer Fraud and Abuse Act of 1986. Pub. L. No. 99-474.  
<https://www.congress.gov/bill/99th-congress/house-bill/4718>

US Congress. 1988. Omnibus Trade and Competitiveness Act of 1988. Pub. L. No. 100-418.  
<https://www.congress.gov/bill/100th-congress/house-bill/4848>

US Congress. 1996. Telecommunications Act of 1996. Pub. L. 104-104.  
<https://www.congress.gov/104/plaws/publ104/PLAW-104publ104.pdf>

US Congress. 1998. Digital Millennium Copyright Act. Pub. L. 105-304.  
<https://www.congress.gov/105/plaws/publ304/PLAW-105publ304.pdf>

US Congress. 1999. Gramm-Leach-Bliley Act. Pub. L. No. 106.  
<https://www.govinfo.gov/content/pkg/PLAW-106publ102/pdf/PLAW-106publ102.pdf>

US Congress. 2000. Electronic Signatures in Global and National Commerce Act (E-SIGN Act). Pub. L. 106-229. <https://www.congress.gov/bill/106th-congress/house-bill/1714>

US Congress. 2002. E-Government Act of 2002. Pub. L. No. 107-347.  
<https://www.congress.gov/bill/107th-congress/house-bill/2458>

US Congress. 2004. Standards Development Organization Advancement Act of 2004. Pub. L. No. 108-237. <https://www.congress.gov/bill/108th-congress/house-bill/1086>

US Congress. 2016. Defend Trade Secrets Act of 2016. Pub. L. No. 114-153.  
<https://www.congress.gov/114/plaws/publ153/PLAW-114publ153.pdf>

US Congress. 2018. Allow States and Victims to Fight Online Sex Trafficking Act of 2017. Pub. L. No. 115-164. <https://www.congress.gov/bill/115th-congress/house-bill/1865>

US Congress. 2022. CHIPS and Science Act of 2022. Pub. L. 117-167.  
<https://uscode.house.gov/download/bills/117-2/117-167.pdf>

USPTO. 2013. Policy Statement on Remedies for Standards-Essential Patents Subject to Voluntary F/RAND Commitments. Washington, DC.

USPTO. 2021. U.S. Patent Statistics Chart Calendar Years 1963-2020.  
[https://www.uspto.gov/web/offices/ac/ido/oeip/taf/us\\_stat.htm](https://www.uspto.gov/web/offices/ac/ido/oeip/taf/us_stat.htm)

USPTO. 2024. MPEP. 9th ed. DoC. <https://www.uspto.gov/web/offices/pac/mpep/index.html>

Ugwi, Monica. 2025. “The Future of Manufacturing: AI for Data Standardization.” Microsoft Industry Blogs. <https://www.microsoft.com/en-us/industry/blog/manufacturing-and-mobility/manufacturing/2025/01/29/the-future-of-manufacturing-ai-for-data-standardization/>

UK CaMA. 2020. Online Platforms and Digital Advertising. UK CaMA.  
[https://assets.publishing.service.gov.uk/media/5fa557668fa8f5788db46efc/Final\\_report\\_Digital\\_ALT\\_TEKT.pdf](https://assets.publishing.service.gov.uk/media/5fa557668fa8f5788db46efc/Final_report_Digital_ALT_TEKT.pdf)

CMA. 2020. Online Platforms and Digital Advertising Market Study. London: Competition and Markets Authority.

ULC. 1999. Uniform Electronic Transactions Act (UETA). National Conference of Commissioners on Uniform State Laws. <https://www.uniformlaws.org/viewdocument/final-act-21?CommunityKey=2c04b76c-2b7d-4399-977e-d5876ba7e034>.

UN (United Nations). 2019. Globally Harmonized System of Classification and Labelling of Chemicals (GHS). UN. <https://unece.org/ghs-rev8-2019>

UNESCO. 2021. Recommendation on the Ethics of Artificial Intelligence.  
<https://unesdoc.unesco.org/ark:/48223/pf0000381137>

UNESCO. 2025. Guidance for Generative AI in Education and Research (Updated from the 2023 version). Paris: UNESCO. <https://www.unesco.org/en/articles/guidance-generative-ai-education-and-research>

Unify. 2025a. Employee data and trends for OpenAI. <https://www.unifygtm.com/insights-headcount/openai>

Unify. 2025b. Employee data and trends for Cohere. <https://www.unifygtm.com/insights-headcount/cohere>

Urbain, Thomas, and Luna Lin. 2025. As US battles China on AI, some companies choose Chinese. TechXplore. <https://techxplore.com/news/2025-12-china-ai-companies-chinese.html>

Urofsky, Melvin I. 2009. Louis D. Brandeis: A Life. New York: Pantheon Books.

US Congress. 1890. Sherman Act. Ch. 647, 26 Stat. 209. <https://www.govinfo.gov/content/pkg/COMPS-3055/pdf/COMPS-3055.pdf>

US Congress. 1914a. Clayton Act. Ch. 323, 38 Stat. 730. <https://www.govinfo.gov/content/pkg/COMPS-3049/pdf/COMPS-3049.pdf>

US Congress. 1914b. Federal Trade Commission Act. Ch. 311, 38 Stat. 717. <https://www.ftc.gov/legal-library/browse/statutes/federal-trade-commission-act>

US Congress. 1934. Communications Act of 1934, Pub. L. No. 73–416. <https://www.govinfo.gov/content/pkg/STATUTE-48/pdf/STATUTE-48-Pg1064.pdf>

US Congress. 1936. Robinson-Patman Antidiscrimination Act. Ch. 592, 49 Stat. 1526. <https://www.govinfo.gov/content/pkg/COMPS-12152/pdf/COMPS-12152.pdf>

US Congress. 1946. Administrative Procedure Act. Pub. L. No. 79-404. <https://www.govinfo.gov/content/pkg/STATUTE-60/pdf/STATUTE-60-Pg237.pdf>

US Congress. 1950. Celler-Kefauver Act. Pub. L. No. 81-899, 64 Stat. 1125. <https://www.govinfo.gov/content/pkg/STATUTE-64/pdf/STATUTE-64-Pg1125.pdf>

US Congress. 1952. Patent Act of 1952. Pub. L. No. 82-593. <https://www.govinfo.gov/content/pkg/STATUTE-66/pdf/STATUTE-66-Pg792.pdf>

US Congress. 1965b. Food and Agriculture Act of 1965. Pub. L. No. 89-321. <https://nationalaglawcenter.org/wp-content/uploads/assets/farmbills/1965.pdf>

US Congress. 1976a. Hart-Scott-Rodino Antitrust Improvements Act. Pub. L. No. 94-435, 90 Stat. 1383. <https://www.govinfo.gov/content/pkg/STATUTE-90/pdf/STATUTE-90-Pg1383.pdf>

US Congress. 1976b. US Copyright Act. Pub. L. 94–553. <https://www.copyright.gov/history/pl94-553.pdf>

US Congress. 1978. Public Utility Regulatory Policies Act of 1978. Pub. L. No. 95-617.  
<https://www.congress.gov/bill/95th-congress/house-bill/4018>

US Congress. 1979a. The 1979 Joint Economic Report. 96th Congress, 1st Session.  
[https://www.jec.senate.gov/reports/96th%20Congress/The%201979%20Joint%20Economic%20Report%20\(930\).pdf](https://www.jec.senate.gov/reports/96th%20Congress/The%201979%20Joint%20Economic%20Report%20(930).pdf)

US Congress. 1979b. Trade Agreements Act of 1979. Pub. L. No. 96-39.  
<https://www.congress.gov/bill/96th-congress/house-bill/4537>

US Congress. 1980a. Patent and Trademark Law Amendments Act (Bayh-Dole Act), Pub. L. No. 96-517.  
<https://www.govinfo.gov/content/pkg/STATUTE-94/pdf/STATUTE-94-Pg3015.pdf>

US Congress. 1980b. Stevenson-Wydler Technology Innovation Act of 1980, Pub. L. No. 96-480.  
<https://uscode.house.gov/statutes/pl/96/480.pdf>

US Congress. 1981. H.R. Rep. No. 97-312, 97th Cong., 1st Sess. (Federal Courts Improvement Act of 1981). <https://www.congress.gov/congressional-report/97th-congress/house-report/312>

US Congress. 1982. Federal Courts Improvement Act of 1982. Pub. L. No. 97-164.  
<https://www.congress.gov/bill/97th-congress/house-bill/4482>

US Congress. 1984a. Semiconductor Chip Protection Act of 1984, Pub. L. No. 98-620.  
<https://www.congress.gov/98/statute/STATUTE-98/STATUTE-98-Pg3335.pdf>

US Congress. 1984b. National Cooperative Research Act of 1984. Pub. L. No. 98-462.  
<https://www.congress.gov/bill/98th-congress/senate-bill/1841>

US Congress. 1986b. Computer Fraud and Abuse Act of 1986. Pub. L. No. 99-474.  
<https://www.congress.gov/bill/99th-congress/house-bill/4718>

US Congress. 1988. Omnibus Trade and Competitiveness Act of 1988. Pub. L. No. 100-418.  
<https://www.congress.gov/bill/100th-congress/house-bill/4848>

US Congress. 1996. Telecommunications Act of 1996. Pub. L. 104-104.  
<https://www.congress.gov/104/plaws/publ104/PLAW-104publ104.pdf>

US Congress. 1998. Digital Millennium Copyright Act. Pub. L. 105–304.  
<https://www.congress.gov/105/plaws/publ304/PLAW-105publ304.pdf>

US Congress. 1999. Gramm–Leach–Bliley Act. Pub. L. No. 106.  
<https://www.govinfo.gov/content/pkg/PLAW-106publ102/pdf/PLAW-106publ102.pdf>.

US Congress. 2000. Electronic Signatures in Global and National Commerce Act (E-SIGN Act). Pub. L. 106-229. <https://www.congress.gov/bill/106th-congress/house-bill/1714>

US Congress. 2002. E-Government Act of 2002. Pub. L. No. 107-347.  
<https://www.congress.gov/bill/107th-congress/house-bill/2458>

US Congress. 2004. Standards Development Organization Advancement Act of 2004. Pub. L. No. 108-237. <https://www.congress.gov/bill/108th-congress/house-bill/1086>

US Congress. 2016. Defend Trade Secrets Act of 2016. Pub. L. No. 114-153.  
<https://www.congress.gov/114/plaws/publ153/PLAW-114publ153.pdf>

US Congress. 2018. Allow States and Victims to Fight Online Sex Trafficking Act of 2017. Pub. L. No. 115-164. <https://www.congress.gov/bill/115th-congress/house-bill/1865>

US Congress. 2022. CHIPS and Science Act of 2022. Pub. L. 117–167.  
<https://uscode.house.gov/download/bills/117-2/117-167.pdf>

USPTO. 1995. "Request for Comments on Proposed Examination Guidelines for Computer-Implemented Inventions." Federal Register 60 (106): 28778–28780.

USPTO. 1996. "Examination Guidelines for Computer-Related Inventions." Federal Register 61(40): 7478–7492.

USPTO. 2020. "Public Views on Artificial Intelligence and Intellectual Property Policy." U.S. Patent and Trademark Office, October 6. <https://www.uspto.gov/subscription-center/2020/uspto-releases-report-artificial-intelligence-and-intellectual-property>

USPTO. 2024. MPEP. 9th ed. DoC. <https://www.uspto.gov/web/offices/pac/mpep/index.html>

USTR. 1985. Israel Free Trade Agreement. <https://ustr.gov/trade-agreements/free-trade-agreements/israel-fta>

USTR. 1989. 1989 Special 301 Report. Office of the United States Trade Representative.  
<https://ustr.gov/sites/default/files/1989%20Special%20301%20Report.pdf>

USTR. 1993. The North American Free Trade Agreement: Summary of the Provisions.  
<https://ustr.gov/sites/default/files/uploads/Countries%20Regions/africa/agreements/nafta/NAFTA%20Chapter%20Summaries.pdf>

USTR. 2025. History of the United States Trade Representative. <https://ustr.gov/about-us/history>

Utility Dive. 2024. US data center electricity demand could double by 2030, driven by artificial intelligence: EPRI. <https://www.utilitydive.com/news/artificial-intelligence-doubles-data-center-demand-2030-EPRI/717467/>

Utterback, James, and William Abernathy. 1975. "A Dynamic Model of Process and Product Innovation." *Omega* 3(6): 639–656.

Vahdat, Amin. 2025. Ironwood: The first Google TPU for the age of inference. Google Cloud Blog.  
<https://blog.google/products/google-cloud/ironwood-tpu-age-of-inference/>

Valentino-DeVries, Jennifer, Natasha Singer, Michael Keller, and Aaron Krolik. 2018. "Your Apps Know Where You Were Last Night, and They're Not Keeping It Secret." *New York Times*.  
<https://www.nytimes.com/interactive/2018/12/10/business/location-data-privacy-apps.html>

Valletti, Tommaso. 2018. "How EU Markets Became More Competitive Than U.S. Markets: A Study of Institutional Drift." NBER Working Paper No. 24700. Cambridge, MA: National Bureau of Economic Research.

Van Alstyne, Marshall, Geoffrey Parker, and Sangeet Paul Choudary. 2016. "Pipelines, Platforms, and the New Rules of Strategy." *Harvard Business Review* 94(4): 54-62.

Varian, Hal. 2014. "Big Data: New Tricks for Econometrics." *Journal of Economic Perspectives* 28(2): 3-28.

Varian, Hal. 2007. "Position Auctions." *International Journal of Industrial Organization* 25(6): 1163–78.

Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention Is All You Need. *Advances in Neural Information Processing Systems* 30. <https://papers.neurips.cc/paper/7181-attention-is-all-you-need>

Vedder, Richard, and Lowell Gallaway. 1993. *Out of Work*. New York: Holmes and Meier.

Verizon. 2002. *Verizon Communications, Inc. v. FCC*. No. 00-511. US Supreme Court, May 13. <https://supreme.justia.com/cases/federal/us/535/467/>

Verizon. 2004. *Verizon Communications, Inc. v. Law Offices of Curtis V. Trinko, LLP*. No. 02-682. US Supreme Court, January 13. <https://supreme.justia.com/cases/federal/us/540/398/>

Vermont Yankee Nuclear Power. 1978. *Vermont Yankee Nuclear Power Corp. v. Natural Resources Defense Council, Inc., et al.*, No. 76-419, U.S. Supreme Court, April 3. [https://scholar.google.com/scholar\\_case?case=4079290121887020113](https://scholar.google.com/scholar_case?case=4079290121887020113)

Vertiv. 2025. High-density cooling: A guide to advanced thermal solutions for AI and ML workloads in data centers. LinkedIn Post. [https://www.linkedin.com/posts/vertiv\\_high-density-cooling-a-guide-to-advanced-activity-7343548574292426752-5r4P](https://www.linkedin.com/posts/vertiv_high-density-cooling-a-guide-to-advanced-activity-7343548574292426752-5r4P)

Veugelers, Reinhilde. 2018. Are European firms falling behind in the global corporate research race? *Bruegel Policy Contribution* 2018/06. <https://www.bruegel.org/policy-brief/are-european-firms-falling-behind-global-corporate-research-race>

Via Licensing Corporation. 2018. Via Expands LTE Patent Pool License Coverage to Include LTE-Advanced Pro. Press Release. [www.via-corp.com/licensing/lte/index.html](http://www.via-corp.com/licensing/lte/index.html)

Vietor, Richard. 1994. *Contrived Competition*. Cambridge, MA: Harvard University Press.

Villalobos, Pablo, Jaime Sevilla, Lennart Heim, Tamay Besiroglu, Marius Hobbahn, and Anson Ho. 2022. "Will We Run Out of Data? An Analysis of the Limits of Scaling Datasets in Machine Learning." arXiv preprint [arXiv:2211.04325](https://arxiv.org/abs/2211.04325). <https://arxiv.org/abs/2211.04325>.

Villareal, Angeles. 2017. "NAFTA and the Mexican Economy." CRS Report for Congress R42965. Washington, DC: Congressional Research Service.

Visa and Mastercard. 2001. U.S. v. Visa U.S.A. Inc. and Mastercard International, Inc., No. 98 CIV. 7076(BSJ). US District Court, S.D. New York, October 9. <https://www.justice.gov/atr/case-document/file/515596/dl>

Vogel, David. 1995. *Trading Up*. Cambridge, MA: Harvard University Press.

Von's Grocery. 1966. *United States v. Von's Grocery Co.*, No. 303. US Supreme Court, May 31. <https://supreme.justia.com/cases/federal/us/384/270/>

W3Techs. n.d. "Usage Statistics and Market Share of Google Analytics for Websites." W3Techs—Web Technology Surveys. <https://w3techs.com/technologies/details/ta-googleanalytics>.

Wallace, James, and Jim Erickson. 1992. *Hard Drive: Bill Gates and the Making of the Microsoft Empire*. New York: HarperBusiness.

WaMC. 2015. *Yes, Reagan Used Trade Promotion Authority*. US House of Representatives. <https://waysandmeans.house.gov/2015/03/13/yes-reagan-used-trade-promotion-authority/>

Ward, John William. 2017. "Digital Markets and Antitrust." *Antitrust Law Journal* 81(3): 919-946.

Warren, Kenneth. 2020. *Administrative Law in the Political System*. New York: Routledge.

Wartell, Sarah. 2009. National Economic Council. In *Change for America: A Progressive Blueprint for the 44th President*, edited by Mark Green and Michele Jolin. Washington, DC: Center for American Progress, 81–87.

Washburn, Brian. 2023. *Omdia's new AI network traffic forecast expects surge, pressure on telcos to meet future growth*. Omdia Analyst Opinion. <https://omdia.tech.informa.com/om119985/omdias-new-ai-network-traffic-forecast-expects-surge-pressure-on-telcos-to-meet-future-g>

Washburn, Brian. 2024. *Road to 2030: AI and the Future of Network Services – Traffic Outlook and Implications*. Omdia Report. <https://omdia.tech.informa.com/om121972/road-to-2030-ai-and>

Wassel, Bryan. 2024. "Verizon CEO says it will 'lead AI revolution,' improve customer service at scale." Customer Experience Dive. <https://www.customerexperiencedive.com/news/verizon-ai-revolution-customer-service-myplan-personalization/714226/>

Waymo. 2025. 20 million. That's how many times you've trusted Waymo to get you where you're going. X (formerly Twitter). <https://x.com/Waymo/status/2001351692240015403>

Wayt, Theo. 2025. "Elon Musk's xAI Buys Building for Third Supersized Data Center." The Information. <https://www.theinformation.com/articles/elon-musks-xai-buys-building-third-supersized-data-center>

WCMP. 2025. "2025 Wisconsin Manufacturing Report." <https://www.wicmp.org/report-conclusions/2025-wisconsin-manufacturing-report/>

Webster, Graham, Rogier Creemers, Elsa Kania, and Paul Triolo. 2017. Full Translation: China's 'New Generation Artificial Intelligence Development Plan'. DigiChina (Stanford Cyber Policy Center). <https://digichina.stanford.edu/work/full-translation-chinas-new-generation-artificial-intelligence-development-plan-2017/>

Weights and Biases. 2024. Experiment Tracking and Model Management. Weights and Biases Documentation. <https://docs.wandb.ai/>

Weil, Jonathan. 2025. "Is the Flurry of Circular AI Deals a Win-Win—or Sign of a Bubble?" The Wall Street Journal. <https://www.wsj.com/tech/ai/is-the-flurry-of-circular-ai-deals-a-win-win-or-sign-of-a-bubble-8a2d70c>

Weiss, Geoff and Katherine Tangalakis-Lippert. 2024. "Trump says Google has 'a lot of power' and he would do 'something' about it — but stops short of favoring a break-up." Business Insider. <https://www.businessinsider.com/trump-google-break-up-china-comments-2024-10>

Welle, Elissa. 2025. "OpenAI reportedly signs \$300 billion cloud deal with Oracle." The Verge. <https://www.theverge.com/ai-artificial-intelligence/776170/oracle-openai-300-billion-contract-project-stargate>

Werden, Gregory. 2009. Next Steps in the Evolution of Antitrust Law: What to Expect from the Roberts Court. *Journal of Competition Law and Economics* 5(1): 49-74.

Werden, Gregory. 2015. Identifying Exclusionary Conduct Under Section 2: The "No Economic Sense" Test. *Antitrust Law Journal* 73(2) 413-433.

Werden, Gregory. 1992. "The History of Antitrust Market Delineation." *Marquette Law Review* 76(1): 123–159.

Werden, Gregory. 1998. "Demand Elasticities in Antitrust Analysis." *Antitrust Law Journal* 66(2): 363–414.

Werden, Gregory. 2003. "The 1982 Merger Guidelines and the Ascent of the Hypothetical Monopolist Paradigm." *Antitrust Law Journal* 71(1): 253–75.

Wessels, Walter. 2001. The Effect of Minimum Wages on the Labor Force Participation Rates of Teenagers. Employment Policies Institute. [https://epionline.org/app/studies/wessels\\_06-2001.pdf](https://epionline.org/app/studies/wessels_06-2001.pdf)

West Virginia et al. 2022. *West Virginia et al. v. United States Environmental Protection Agency and Michael S. Regan (Administrator of the United States Environmental Protection Agency)*, No. 20-1530, US Supreme Court, June 30. [https://www.supremecourt.gov/opinions/21pdf/20-1530\\_n758.pdf](https://www.supremecourt.gov/opinions/21pdf/20-1530_n758.pdf)

Weyerhaeuser. 2007. *Weyerhaeuser Co. v. Ross-Simmons Hardwood Lumber Co.*, No. 05-381. US Supreme Court, February 20. <https://supreme.justia.com/cases/federal/us/549/312>

Weyl, Eric Glen. 2010. "A Price Theory of Multi-Sided Platforms." *American Economic Review* 100(4): 1642–1672.

White and Case LLP. 2025. "AI Watch: Global Regulatory Tracker – China." White and Case Insights. <https://www.whitecase.com/insight-our-thinking/ai-watch-global-regulatory-tracker-china>

White House Council of Economic Advisers. 2024. *Economic Report of the President*. Washington, DC: U.S. Government Publishing Office.

Wiggers, Kyle. 2024a. "Anthropic raises another \$4B from Amazon, makes AWS its 'primary' training partner." TechCrunch. <https://techcrunch.com/2024/11/22/anthropic-raises-an-additional-4b-from-amazon-makes-aws-its-primary-cloud-partner/>

Wiggers, Kyle. 2024b. "Elon Musk's xAI Lands \$6B in New Cash to Fuel AI Ambitions." TechCrunch. <https://techcrunch.com/2024/12/25/elon-musks-xai-lands-billions-in-new-cash-to-fuel-ai-ambitions/>

Wikipedia. 2025a. "3G." Wikipedia. <https://en.wikipedia.org/wiki/3G>

Wilander, John. 2017. "Intelligent Tracking Prevention." WebKit Blog. <https://webkit.org/blog/7675/intelligent-tracking-prevention/>

Williams, Tom. 2025. "What We Know About OpenAI's First Consumer Devices." Information Age ACS. <https://ia.acs.org.au/article/2025/what-we-know-about-openais-first-consumer-devices.html>

Williamson, John. 1990. "What Washington Means by Policy Reform." In *Latin American Adjustment*. Edited by John Williamson. Washington, D.C.: Institute for International Economics, 5-20.

Williamson, Oliver. 1975. *Markets and Hierarchies: Analysis and Antitrust Implications*. New York: Free Press.

Williamson, Oliver. 1985. *The Economic Institutions of Capitalism*. New York: Free Press.

Winkler, Rolfe, and Nate Rattner. 2025. "The Driver of Apple's Exploding Valuation Is Under Threat. See What's at Stake." *The Wall Street Journal*. <https://www.wsj.com/tech/apple-stock-service-value-growth-analysis-f265a5e7>

Winston, Clifford and Steven Morrison. 1997. *The Fare Skies: Air Transportation and Middle America*. Brookings Institution.

Wiseman, Bill, Henry Marciel, and Marc de Jong. 2025. "Semiconductors have a big opportunity—but barriers to scale remain." McKinsey & Company. <https://www.mckinsey.com/industries/semiconductors/our-insights/semiconductors-have-a-big-opportunity-but-barriers-to-scale-remain>

Wolla, David. 2020. Examining the "Lump of Labor" Fallacy Using a Simple Economic Model. Federal Reserve Bank of St. Louis. <https://www.stlouisfed.org/publications/page-one-economics/2020/11/02/examining-the-lump-of-labor-fallacy-using-a-simple-economic-model>

Wong, Kenny. 2024. What is RAG? 4 analogies for this powerful AI approach. Coda Blog. <https://coda.io/blog/ai/what-is-rag-analogies>

Woo, Seongjin, and Raffaele Huang. 2025. How China's DeepSeek Outsmarted America. *The Wall Street Journal*. [https://www.wsj.com/tech/ai/china-deepseek-ai-nvidia-openai-02bdbbce?mod=hp\\_lead\\_pos1](https://www.wsj.com/tech/ai/china-deepseek-ai-nvidia-openai-02bdbbce?mod=hp_lead_pos1)

Workday. 2023. Responsible AI practices. Workday. <https://www.workday.com/en-us/artificial-intelligence/responsible-ai-practices.html>

World Bank. 2023. “Charges for the Use of Intellectual Property, Receipts (BoP, current US\$)” (indicator BX.GSR.ROYL.CD) and “Charges for the Use of Intellectual Property, Payments (BoP, current US\$)” (indicator BM.GSR.ROYL.CD). World Development Indicators.

<https://data.worldbank.org/indicator/BX.GSR.ROYL.CD> and

<https://data.worldbank.org/indicator/BM.GSR.ROYL.CD>

WEF. 2025. “Towards Equitable AI: New Report Charts Path to AI Competitiveness.” WEF.

<https://www.weforum.org/press/2025/01/towards-equitable-ai-new-report-charts-path-to-ai-competitiveness/>

Wright, Joshua David. 2011. Antitrust, Multi-Dimensional Competition, and Innovation: Do We Have an Antitrust-Relevant Theory of Competition Now? In *Competition Policy and Patent Law under Uncertainty: Regulating Innovation*, edited by Geoffrey Manne and Joshua Wright. Cambridge: Cambridge University Press, 228–251.

Wright, Joshua. 2012. Defining and Measuring Search Bias: Some Preliminary Evidence. George Mason Law and Economics Research Paper No. 12-14.

[https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2004649](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2004649)

Wright, Joshua. 2014. Evidence-Based Antitrust Enforcement in the Technology Sector. *CPI Antitrust Chronicle* 3(1): 1-9.

Wright, Joshua, and John Yun. 2018. "Ohio v. American Express: Implications for Non-Transaction Multisided Platforms." *CPI Antitrust Chronicle* 2(1): 1–8.

WTO. 2021. World Tariff Profiles 2021. WTO [https://www.wto.org/english/docs\\_e/legal\\_e/marag\\_e.htm](https://www.wto.org/english/docs_e/legal_e/marag_e.htm)

WTO. 2008. World Trade Report 2008: Trade in a Globalizing World. WTO.

[https://www.wto.org/English/res\\_e/booksp\\_e/anrep\\_e/world\\_trade\\_report08\\_e.pdf](https://www.wto.org/English/res_e/booksp_e/anrep_e/world_trade_report08_e.pdf)

WTO. n.d.a. Overview: the TRIPS Agreement. WTO

[https://www.wto.org/english/tratop\\_e/trips\\_e/intel2\\_e.htm](https://www.wto.org/english/tratop_e/trips_e/intel2_e.htm)

WTO. n.d.b. Intellectual Property: Protection and Enforcement. WTO  
[https://www.wto.org/english/thewto\\_e/whatis\\_e/tif\\_e/agrm7\\_e.htm](https://www.wto.org/english/thewto_e/whatis_e/tif_e/agrm7_e.htm)

Wu, Philipp, Alejandro Escontrela, Danijar Hafner, Pieter Abbeel, and Ken Goldberg. 2022. "Daydreamer: World Models for Physical Robot Learning." Proceedings of the 6th Conference on Robot Learning (CoRL). <https://openreview.net/forum?id=3RBY8fKjHeu>.

Wu, Tim. 2018. The Curse of Bigness. New York: Columbia Global Reports.

Wu, Yonghui, et al. 2016. "Google's Neural Machine Translation System: Bridging the Gap between Human and Machine Translation." arXiv preprint arXiv:1609.08144. <https://arxiv.org/abs/1609.08144>

Xu, Kevin. 2025. "DeepSeek's Three Idiosyncratic Advantages." Interconnected. <https://interconnected.blog/deepseeks-three-idiosyncratic-advantages/>

Yahoo Finance. 2025. AI Data Center Global Research Report 2025–2030. Yahoo Finance. <https://uk.finance.yahoo.com/news/ai-data-center-global-research-091100624.html>

Yale Budget Lab. 2025. "State of U.S. Tariffs: November 17, 2025." New Haven, CT: The Budget Lab at Yale. <https://budgetlab.yale.edu/research/state-us-tariffs-november-17-2025>.

Yan, Peifeng, Xinyuan Bai, and Zehao Shen. 2022. "Operations and Supply Chain Analysis of the Smartphone Industry: Comparing Apple and Huawei." Advances in Economics, Business and Management Research 215: 1523-1525.

Yuan, Dave, and Jair Verçosa. 2025. Does AI remove the incumbent data gravity advantage? Tidemark Capital. <https://www.tidemarkcap.com/post/does-ai-remove-the-incumbent-data-gravity-advantage>

Zaitsev, Peter. 2024. Vector Search in Modern Databases. Percona University (OSACon presentation). <https://percona.university/wp-content/uploads/2025/03/Vector-Search-in-Modern-Databases-by-Peter-Zaitsev.pdf>

Zeran. 1997. Zeran v. America Online, Inc. No. 96-433. US Court of Appeals, Fourth Circuit, November 12. <https://law.justia.com/cases/federal/appellate-courts/F3/129/327>

Zhu, Qiuling, Berkin Akin, H. Ekin Sumbul, Fazle Sadi, James C. Hoe, Larry Pileggi, and Franz Franchetti. 2013. A 3D-Stacked Logic-in-Memory Accelerator for Application-Specific Data Intensive Computing. Proceedings of the 2013 IEEE International 3D Systems Integration Conference (3DIC): 1–7.

Zhu, Yukun, Ryan Kiros, Richard S. Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. “Aligning Books and Movies: Towards Story-Like Visual Explanations by Watching Movies and Reading Books.” Proceedings of the IEEE International Conference on Computer Vision (ICCV 2015): 9–27.

Zilbershot, Dorit. 2025. Unlock Productivity: AI Agent Solutions. ServiceNow Blog. <https://www.servicenow.com/blogs/2025/unlock-productivity-ai-agent-solutions>

Zittrain, Jonathan. 2008. The Future of the Internet—And How to Stop It. New Haven, CT: Yale University Press.

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

Zubrow, Keith. 2021. “Facebook Whistleblower Says Company Incentivizes ‘Angry, Polarizing, Divisive Content.’” CBS News. <https://www.cbsnews.com/news/facebook-whistleblower-frances-haugen-60-minutes-polarizing-divisive-content/>

Zuckerberg, Mark. 2024. "Open Source AI Is the Path Forward." Meta Newsroom. <https://about.fb.com/news/2024/07/open-source-ai-is-the-path-forward/>

Zulli, Diana, and David James Zulli. 2022. "Extending the Internet Meme." *New Media and Society* 24(8): 1872–1890.