



CHAPTER

Technological Disruption in the US Labor Market

by David Deming, Christopher Ong, and Lawrence H. Summers

This paper was produced to provide policy-relevant evidence about current challenges confronting the American economy. Authors are invited to share their views about policy issues, which do not necessarily represent those of the Aspen Institute, members of the Aspen Economic Strategy Group, or their affiliated organizations.

Suggested Citation: Deming, David, Christopher Ong, and Lawrence H. Summers. 2024. "Technological Disruption in the US Labor Market" In *Strengthening America's Economic Dynamism*, edited by Melissa S. Kearney and Luke Pardue. Washington, DC: Aspen Institute.

Technological Disruption in the US Labor Market

AUTHORS

David Deming,* Christopher Ong,** and Lawrence H. Summers***

ABSTRACT

This paper explores past episodes of technological disruption in the US labor market, with the goal of learning lessons about the likely future impact of artificial intelligence (AI). We measure changes in the structure of the US labor market going back over a century. We find, perhaps surprisingly, that the pace of change has slowed over time. The years spanning 1990 to 2017 were less disruptive than any prior period we measure, going back to 1880. This comparative decline is not because the job market is stable today but rather because past changes were so profound. General-purpose technologies (GPTs) like steam power and electricity dramatically disrupted the twentieth-century labor market, but the changes took place over decades. We argue that AI could be a GPT on the scale of prior disruptive innovations, which means it is likely too early to assess its full impacts. Nonetheless, we present four indications that the pace of labor market change has accelerated recently, possibly due to technological change. First, the labor market is no longer polarizing—employment in low- and middle-paid occupations has declined, while highly paid employment has grown. Second, employment growth has stalled in low-paid service jobs. Third, the share of employment in STEM jobs has increased by more than 50 percent since 2010, fueled by growth in software and computer-related occupations. Fourth, retail sales employment has declined by 25 percent in the last decade, likely because of technological improvements in online retail. The post-pandemic labor market is changing very rapidly, and a key question is whether this faster pace of change will persist into the future.

* Isabelle and Scott Black Professor of Political Economy, Harvard Kennedy School

** Harvard University

*** President Emeritus & Charles W. Eliot University Professor, Harvard University

Introduction

In recent decades, the spread of digital technology in the US and around the world has changed the skill requirements of many jobs (Muro et al. 2017). More recently, the global COVID-19 pandemic disrupted labor markets around the world, leading to a rapid increase in remote work and changes in relative wages for in-person service jobs (Barrero, Bloom, and Davis 2023; Autor, Dube, and McGrew 2023). Looking to the future, labor market analysts and pundits have predicted that artificial intelligence (AI) will transform work and that employers will deemphasize traditional educational credentials in favor of “skills-based hiring” (see, for example, World Economic Forum 2023; Ellingrud et al. 2023).

Across ten different advanced and emerging economies, a majority of citizens believe that automation will replace existing jobs, increase inequality, and make it harder to find work (Wike and Stokes 2018). Yet automation anxiety is nothing new. In 1964, President Lyndon Johnson created the Blue Ribbon National Commission on Technology, Automation, and Economic Progress, which concluded (among other things) that the US should have a guaranteed minimum income to support workers through the coming wave of automation. Of course, that wave did not come to pass, and the employment-to-population ratio is substantially higher now than it was in 1964. Is this time different?

In this paper we analyze past episodes of technological disruption in the US labor market, to learn lessons about the future of work with AI. We create a comprehensive measure of labor market “churn,” which allows us to measure and systematically compare changes in the occupation structure of the US labor market over long periods of time.

We find that—contrary to popular imagination—the pace of labor market disruption has slowed in recent decades. The changes in the structure of US employment at the end of the nineteenth century were greater than in any decade of the digital era, including the most recent one. Even more disruptive was the period between 1940 and 1970, when agricultural employment was still disappearing, manual labor was shifting into production and away from railroads, and clerical and administrative work were growing rapidly. The years spanning 1990 to 2017 were the most stable period in the history of the US labor market, going back nearly 150 years.

At the dawn of the twentieth century, 40 percent of US employment was in agriculture, compared to less than 2 percent today. Nearly half of all workers held blue-collar production and manual-labor jobs in 1960, compared to only 20 percent today. While the occupation structure of the US labor market has changed since 1980, that change has been relatively modest in historical perspective.

In the past, labor market disruption was fueled by breakthrough general-purpose technologies (GPTs) like steam power and electricity, which enabled the mechanization of agriculture. Mechanization destroyed farming jobs, but it also created factory jobs by increasing labor productivity in manufacturing. Similarly, computer-based manufacturing techniques developed in the 1970s replaced precision production jobs, while also increasing the availability of digital data and the value of analytical and managerial skills.

Although technological breakthroughs often happen suddenly, technology adoption and the pace of labor market change are often gradual. The share of employment in US agriculture fell steadily by 20 percent per decade between 1880 and 1970, and the decline in blue-collar work after 1970 was equally deliberate.

“The impact of AI is likely to be widespread and long-lasting, fitting the mold of past GPTs.”

Is AI a GPT? And if so, will it create any long-run labor market disruption on the same scale as the technologies of the past? Broadly speaking, AI—and machine learning (ML) in particular—is a technology that improves our ability to analyze and interpret data.

In this sense, it exists downstream of the information technology revolution that began in the 1970s and 1980s. Yet the novelty of AI/ML is the way in which data are used. Rather than following a set of explicit instructions (for example, software code) that are scripted in advance, AI/ML algorithms “learn” about the world by studying and copying actions and implicit rules that are inferred from patterns in the data (Autor 2015; Brynjolfsson, Rock, and Syverson 2019).

AI can predict legal liability from contract language, the likelihood that a medical image indicates a specific pathology, or the next word or phrase in a standard office document, among many other possibilities. In this sense, AI is best understood as a *prediction technology* (Agrawal, Gans, and Goldfarb 2019). Since most jobs require some prediction and decision-making, AI will augment or automate aspects of nearly all jobs in the US economy (see, for example, Deming 2021; Eloundou et al. 2023). Thus, the impact of AI is likely to be widespread and long-lasting, fitting the mold of past GPTs. However, history teaches us that even if AI disrupts the labor market, its impact will unfold over many decades.

For these reasons, it is still too early to forecast the impact of AI on the labor market. Nonetheless, in the second part of the paper we study recent changes in the occupation structure of the US labor market, looking for early signs of technological disruption.

We document four stylized facts about the US employment growth in recent decades. First, we show that job polarization—meaning employment growth at the

bottom and top of the wage distribution and declines in the middle—is no longer an accurate description of what is happening in the US labor market. In recent years, employment has declined in both low- and middle-paid occupations and grown rapidly in high-paying managerial, professional, and technical jobs. *Skill upgrading* describes what is happening in the post-pandemic US labor market better than *polarization* does.

Second, we show that the growth of low-paid service jobs like home health aides, food preparation and service workers, cleaners, barbers, and fitness instructors has stalled completely since 2010, after having grown rapidly in the 1990s and the first decade of the 2000s (Autor and Dorn 2013). Service employment cratered during the COVID-19 pandemic, but employment growth in service occupations slowed in the early 2010s, likely because of rising labor costs (Autor, Dube, and McGrew 2023).

Third, we find that science, technology, engineering, and math (STEM) employment has grown rapidly over the last decade. STEM employment grew from 6.5 percent of all jobs in 2010 to nearly 10 percent in 2024, an increase of more than 50 percent. The growth in STEM work is concentrated in occupations like software developers and programmers, although we also see increases across a broad range of science and engineering occupations. STEM job growth has accelerated especially rapidly since 2017 and is matched by increasing private-capital investment in AI-related technology.

Fourth, we find suggestive evidence of AI-related employment disruption in retail sales. While large language models (LLMs) like ChatGPT are too new for us yet to see any direct impact on the labor market, companies have been using predictive AI to optimize business operations since at least the mid-2010s. Online retailers like Amazon use AI to personalize prices and product recommendations and to manage inventory more efficiently, outcompeting big-box retail (see, for example, Deming 2020).

We find that retail sales jobs have declined by 25 percent in the last decade. There were 850,000 *fewer* retail sales workers in the US in 2023 compared to 2013, even though the US economy added more than 19 million jobs over this period. The decline in retail sales began long before COVID-19 but has accelerated in the last few years. Labor productivity growth in retail sales has also outpaced that in other sectors. Interestingly, online retail has also led to growth of “last-mile” jobs like light-truck delivery drivers and stockers and order fillers.

The decline of retail sales fits a broader pattern of technology-fueled occupational upgrading in white-collar office work. Since 1990, front-office jobs like secretaries and administrative assistants and back-office jobs like billing and financial processing have declined by more than 50 percent as a share of all jobs in the US economy. Yet

managers and business analysis jobs have grown rapidly over this same period. While this change may partly reflect title inflation (for example, manager versus supervisor), it also captures a shift away from routine monitoring and categorizing and toward strategy and decision-making (Deming 2021). Our best guess is that this trend will continue. We expect continued declines in routine office work and retail sales jobs, and a ratcheting up of firms' expectations of managers and business analysts, who will now be valued only to the extent that they can use AI to become more productive.

Taken together, these four facts suggest that we may be entering a period of more pronounced labor market disruption. To illustrate this point, we recalculate our measure of labor market "churn" through 2022. Although the 2010s were very stable pre-pandemic, the post-pandemic labor market has changed dramatically. A key outstanding question is whether the labor market disruption of the past few years is a temporary response to the changes wrought by COVID-19 or an early sign of AI-fueled labor market disruption.

1. Data and measurement of technological change

We measure technological disruption in the US labor market by studying changes in the skill and task content of work over more than a century. Which types of jobs have grown, and which are disappearing? Technology *substitutes* for human labor through automation (for example, Acemoglu and Restrepo 2019). However, technology also *complements* labor by making existing workers more productive and by changing their capabilities in ways that create new types of work (for example, Autor 2015).

We focus on the labor market impacts of general-purpose technologies, rather than on innovations like genome sequencing that mostly affect specific sectors of the economy (Williams 2013). We follow Bresnahan (2010), who defines a GPT as any technology that is widely used and capable of ongoing improvement, and that enables complementary innovation across many different applications. Past examples of GPTs include new forms of energy such as steam power and electricity, as well as process and management innovations such as the encoding and storing of digital information (for example, the computerization of the labor market) or the Toyota Lean Management (TLM) approach in manufacturing.¹

GPTs have remade the US labor market, but their impacts often emerge only gradually. Before the advent of steam power, farmers hand-threshed wheat and oats with a flail or relied on horsepower (for example, Rasmussen 1982). Although steam-

¹ These four GPTs and others are studied in more detail in David 1990; Bresnahan and Trajtenberg 1995; Lipsey and Chrystal 1995; Crafts 2004; and Rosenberg and Trajtenberg 2004.

powered machines were invented in the early 1800s, it was not until the 1860s that farmers began threshing grain with “iron horses”—steam-powered traction engines that mechanized agriculture and increased threshing productivity by a factor of more than a hundred. Yet adoption took place slowly over the next half-century. The mechanization of agriculture was only completed after the adoption of the internal combustion engine by American tractor companies in the late 1910s and early 1920s (Olmstead and Rhode 2001).

Similarly, while Thomas Edison invented the incandescent light bulb in 1879, electricity had barely been adopted twenty years later. Electric lighting was used by only 3 percent of US residences in 1899, and even as late as the early 1920s only half of US factory capacity had been electrified (David 1990). The historical evidence on GPT adoption suggests that productivity growth follows a “J curve,” where growth is initially sluggish because new complementary investments are required (Brynjolfsson, Rock, and Syverson 2021). For example, steam-powered factories had to be organized centrally around inefficient and heavy driveshafts, whereas factories powered by electricity could spread out to accommodate a production line (for example, Du Boff 1979).

To capture the full disruptive impact of GPTs, we study long-run changes in the occupation structure of the US economy from 1900 to the present. It is important to note that many other changes have taken place in the US labor market over the last century, including mass migration from Europe (Abramitzky, Boustan, and Eriksson 2014), rising female labor-force participation during the middle of the twentieth century (Goldin 2006), increasing educational attainment (Goldin and Katz 2007), and declining male labor-force participation (Binder and Bound 2019).

While each of these changes had large effects on the composition of US employment and on wage inequality and social structure, they have typically occurred alongside the changes we document, and we do not explicitly focus on them here. For example, while women are more likely to work in clerical occupations than in blue-collar jobs, the share of women working in both types of jobs increased a lot during the 1950s and 1960s. Of course, changes in labor supply interact with technological disruption. For example, the increasing automation of blue-collar production jobs disproportionately affected less educated men (for example, Abraham and Kearney 2020).

We study the impact of technological change on the labor market by plotting changes over time in the relative frequency of occupations. Specifically, we compute for each year the share of all jobs in the US economy that belong to each occupation group. This approach is superior in our view to studying net job gain or decline, in part because of changes in the size and composition of the labor force over time.

Although the US population has grown by about 20 percent since 2000, labor force participation and hours per worker have declined modestly, so that total hours worked increased by about 16 percent (Bick et al. 2023).

We combine data from the US Census for 1880 to 2000 and the American Community Survey (ACS) for 2001 to 2022 to create the longest possible time series. Wages and labor supply weights are only available from 1940 onward. We supplement the census and ACS data with the Current Population Survey (CPS) monthly data, which is collected through April 2024. We combine the CPS monthly files into annual averages to maximize precision.

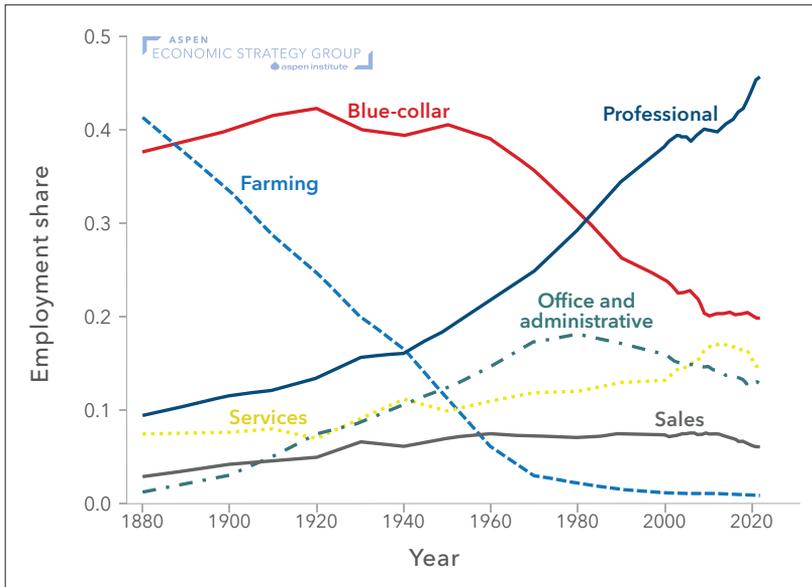
2. Labor market volatility over the last century

With sufficiently broad occupation categories, we can study labor market disruption going back 140 years. Figure 1 presents occupational employment shares from 1880 to the present for six major categories of jobs. In 1880, more than 40 percent of all workers in the US economy were employed as farmers or farm laborers. This share fell steadily by 4 percentage points per decade, and by 1970 only about 3 percent of US employment was in agriculture. Similarly, about 40 percent of workers in 1880 were employed in blue-collar jobs like manual labor, construction, production and manufacturing, transportation, and maintenance and repair. This share stayed relatively constant through 1960, then began a precipitous decline of 4 percentage points per decade, reaching 20 percent by 2010.

Figure 1 offers two key insights. First, even though technological breakthroughs often occur rapidly, technology gets adopted gradually, and the disruption of labor markets takes decades. Second, the productivity growth enabled by GPTs leads to the creation of new occupation categories and makes other occupations more productive.

For example, even though employment in farming fell rapidly during the early twentieth century, agricultural output continued to rise rapidly as it became highly mechanized. The same technology that enabled mechanization of farming also increased the productivity of factory work, enabling new process improvements like assembly lines and creating new jobs.

In some ways this point is obvious. Software developer is one of the fastest-growing occupations in the US, yet it would not exist without the invention of the personal computer. The engineering occupation grew rapidly during the middle of the twentieth century, and “electrical engineer” cannot exist as a job option if a society has no electricity.

Figure 1: Changes in the occupation structure of the US labor market, 1880-2024

Notes: Calculations are based on decadal US census data from 1880 to 2000 (except for 1890) and 2001-2022 American Community Survey (ACS) samples (except for 2020), sourced via the Integrated Public Use Microdata Series (IPUMS) (Ruggles et al. 2024). Occupations are harmonized across decades to two-digit SOC codes using the IPUMS occ1950 encoding and methodology used in Autor and Dorn 2013; a detailed methodology is described in the data appendix. Samples are restricted to workers aged 18 to 64 in noninstitutional quarters who provide nonmilitary occupational responses. See the appendix for exhaustive definitions of each category.

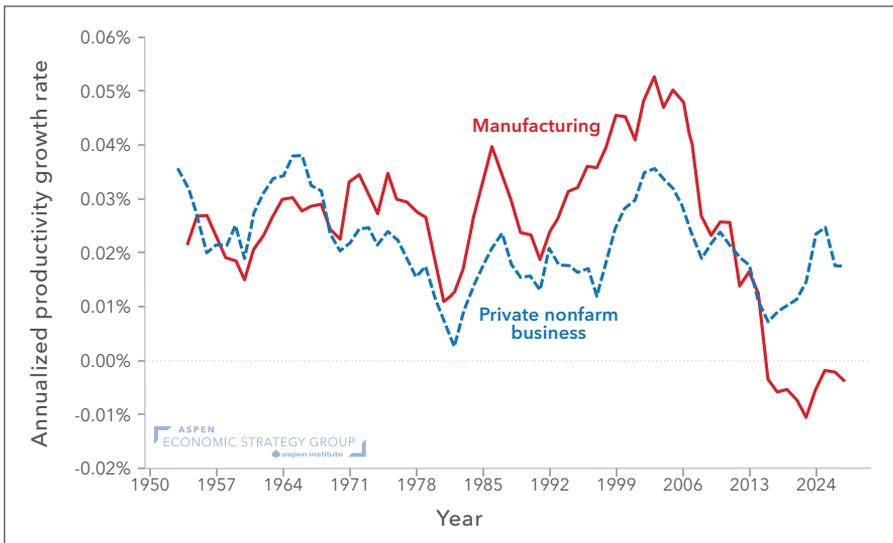
In a groundbreaking new study, Autor et al. (2024) link patent data to the creation of new job titles by census enumerators from 1940 onward to show how technological progress replaces some jobs while complementing the output of others and creating “new work.” They find that some patented innovations directly increase labor productivity, which in turn expands the set of tasks that workers do and leads to net employment growth. They call these changes *augmentation* innovations. Others, which they call *automation* innovations, generate employment declines.

Moreover, Autor et al. (2024) show that during the 1940–1980 period, production occupations were highly exposed to automation innovations, while office clerks and other administrative-support workers were highly exposed to augmentation innovations. This difference is reflected in employment share declines for blue-collar production jobs and growth in clerical, office, and administrative-support jobs.

Figure 2 plots annualized labor productivity growth in manufacturing compared to all other nonfarm sectors of the economy, from 1953 to the present. From 1970 to

the early years of the first decade of the 2000s, labor productivity grew faster in manufacturing than in the economy overall. This growth coincides almost exactly with the period of rapid employment declines in blue-collar work shown in figure 1. Productivity growth in manufacturing was the same or slower than overall economic growth in the 1950s and 1960s, and then more recently since 2010.

Figure 2: US labor productivity growth, manufacturing vs. all other industries, 1950-2024



Notes: Calculations are based on Bureau of Labor Statistics “Historical Productivity and Cost Measures” for manufacturing sectors from 1949 to 2003 and “Labor Productivity and Cost Measures” for major sectors from 2004 to 2024. Displayed yearly growth rates have been smoothed out by annualizing five-year growth rates (US Bureau of Labor Statistics 2024). Labor productivity is measured as output per hour worked.

While technology-enabled productivity growth led to employment declines in manufacturing, it also created new opportunities. Autor et al. (2024) also show how augmentation innovations in businesses processes and information technology led to the rapid rise in the US of a managerial and professional class of workers. To cite just one example, information technology enabled computer-based manufacturing techniques like the computer numerically controlled machine, allowing workers to program routine production tasks like valve cutting with incredible precision (Bartel, Ichniowski, and Shaw 2007). The automation of routine production tasks enabled more customization and process optimization, thereby complementing high-skilled managerial work (Brynjolfsson and Hitt 2003).

We see this process play out in figure 1. Professional occupations—including managers, engineers, lawyers, teachers, and doctors—have grown from less than 10 percent of all jobs in 1880 to about 45 percent in 2024. Professional jobs, which mostly require a college degree and offer high average earnings, have grown especially rapidly in the last 50 years. In both relative and absolute terms, the growth in professional occupations has fully offset the decline of blue-collar work in the US since 1970.

How does the pace of change in the US labor market today compare to that during earlier time periods? The share of employment in farming and agriculture declined steadily by 20 to 25 percent per decade between 1880 and 1980. Employment shares for blue-collar production occupations have declined by an average of 20 percent per decade since 1970. More recently, office and administrative-support occupations declined by 22 percent between 2022 and 2012 after a 12 percent drop the previous decade. Finally, sales occupations have declined by 15 percent in the last ten years, driven by a rapid drop of more than 25 percent in retail sales occupations since 2018. While the decline in office jobs has accelerated in recent years, it does not quite match the magnitude of earlier declines in farming and blue-collar work.

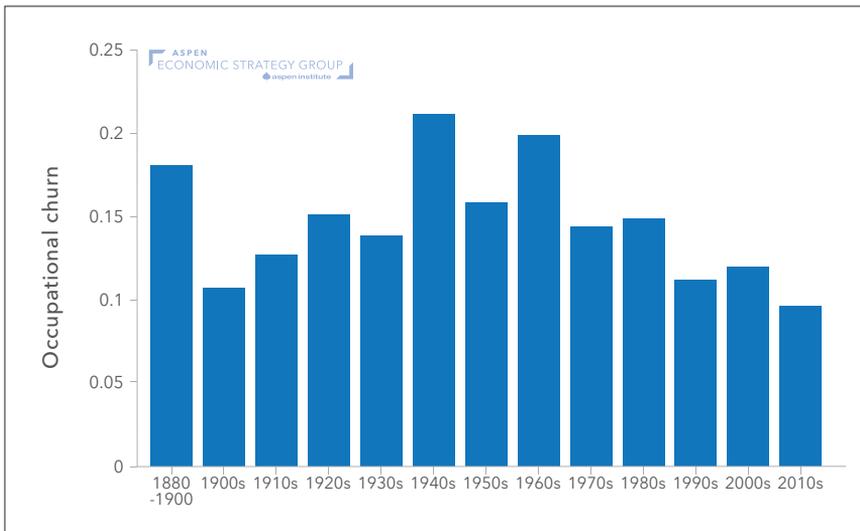
We measure the pace of labor market change more systematically by computing an index of employment disruption that we call “churn.” We define churn for each occupation group as the absolute value of the difference in employment shares over a decade. For example, sales occupations declined from 10.8 percent to 9.1 percent of employment between 2012 and 2022, while business occupations increased from 4.4 to 5.8 percent. By our metric, the churns in sales and business were 1.7 and 1.4 percentage points respectively. We then sum these values up across all occupation categories to get a comprehensive measure of labor market disruption.

While there is no perfect way of capturing labor market change, our churn measure has three advantages relative to other alternatives. First, it is a comprehensive measure that accounts for changes across types of occupations. Second, it is symmetric because employment shares always add up to one and job gains and losses count equally. Third, churn can be interpreted intuitively as a measure of imbalance between periods of time. Suppose there were exactly one hundred occupations. If each had a stable employment share of 1 percent at the beginning and the end of a decade, churn would be equal to zero. If half the shares were at 0.5 percent and half were at 1.5 percent, and the shares flipped completely (so that the occupations at 0.5 percent increased to 1.5 percent, and vice versa), churn would have a value of 1.

Figure 3 plots occupational churn in the US labor market by decade, from 1880 to 2019. To facilitate comparison across many different years, we group occupations into 20 different groups that roughly correspond to the two-digit Standard Occupation Classification (SOC) codes currently used by US federal agencies.²

Figure 3 offers three main lessons. First, the 1880–1900 period and the transition from agriculture to industry were more disruptive than any decade during the computer and digital era. Between 1880 and 1900, the employment share in farming shrank by nearly 8 percentage points. The offsetting gains were distributed evenly across a wide range of occupation groups. On one hand, 1880–1900 is two decades rather than one, so perhaps the lack of data for 1890 overstates the pace of change. On the other hand, shifting employment out of agriculture required mass migration since most other jobs were in urban areas, a disruptive feature that is not captured by our method.

Figure 3: Labor market churn by decade, 1880–2020



Notes: Occupational churn over each period is calculated as the sum of the total changes in absolute value of the employment share attributed to each two-digit SOC code. Employment shares are calculated based on decadal US census data from 1880 to 2000 (except 1890) and ACS samples for 2010 and 2019, sourced via IPUMS (Ruggles et al. 2024). Occupations are harmonized across decades to two-digit SOC codes by extending the methodology used in Autor and Dorn 2013; a detailed methodology is described in the data appendix. Samples are restricted to workers aged 18 to 64 in noninstitutional quarters who provide nonmilitary occupational responses.

² We harmonize occupations across decades by combining two existing crosswalks and a novel crosswalk. We use the occ1950 encoding provided by IPUMS for 1880–1950 and the occ1990dd crosswalk as used in Autor and Dorn 2013 for 1950–2000; we also use a novel crosswalk that harmonizes OCCSOC codes from 2000 to 2022 according to the procedures detailed in IPUMS (Ruggles et al. 2024; Flood et al. 2023). By stitching together crosswalks using overlapping years, we construct a unified crosswalk that allows us to coarsely allocate occupations into groups that correspond to modern SOC codes.

The second lesson from figure 3 is that the three-decade period from 1940 to 1970 was the most volatile period in the history of the US labor market. The transition out of agriculture continued over this period, with the employment share in farming declining by more than 13 percentage points between 1940 and 1970. There was also a significant compositional shift in blue-collar employment. Construction, installation, and repair jobs grew rapidly, but the employment share in transportation occupations fell from 15 percent to 8 percent. This drop was driven by a shift away from railroad transportation and toward the automobile. Professional and administrative-support occupations also grew rapidly.

The third lesson is that after the 1940–1970 period, the US labor market entered a thirty-year period of relative quietude. In fact, the 1990s and the 2010s were among the least volatile periods since 1880, with 2000–2010 ranking not far behind. This comparative calm is particularly striking when juxtaposed against the automation anxiety that captured public imagination during the 2010s. One study—now cited more than 14,000 times—estimated that nearly half of all US employment could be imminently replaced by computerization (Frey and Osborne 2017).

A historical perspective suggests that the labor market has changed relatively slowly over the past three decades. However, if AI meets the conditions of a GPT—widely used, capable of ongoing improvement, and enabling complementary innovation—the pace of change may quicken.

3. Is artificial intelligence (AI) a GPT?

The term “artificial intelligence” refers broadly to methods that enable machines to perceive their environment, learn, and take actions that achieve certain goals. In practice, recent commercial applications of AI are concentrated in machine learning (ML), a branch of AI research that develops complex statistical algorithms that can learn and generalize from data. The foundation of large language models (LLMs) like ChatGPT is a branch of ML called deep learning, which represents text, images, audio/video, and other data as multilayered neural networks in metaphorical relationship to the human brain.

Broadly speaking, AI and ML are technologies for analyzing, interpreting, and responding to data. Thus, one view is that they are a particularly important chapter of the information technology revolution that began in the 1970s and 1980s. Yet the true innovation of AI and ML is the way in which data are used. Standard programming techniques operate from the top down—they write down a set of instructions in advance (for example, code) that can be routinely and rigorously executed. In contrast, machine learning operates from the bottom up. ML algorithms

infer patterns from data, effectively “learning” about the world by studying and copying the actions of others (Autor 2015; Brynjolfsson, Rock, and Syverson 2019).

Robert Solow (1987) famously quipped, “You can see the computer age everywhere but in the productivity statistics.” Indeed, apart from a brief period in the early 1990s, labor productivity growth has been relatively slow over the last several decades. An optimistic view is that we are in the trough of the “J curve” and that the complementary investments in information processing made between 2000 and 2020 will soon lead to an acceleration of productivity growth (Brynjolfsson, Rock, and Syverson 2021).

Agrawal, Gans, and Goldfarb (2019) argue that AI is best understood as a *prediction technology*. They give several examples, such as predicting legal liability based on contract language, surmising the best route between two points given traffic patterns, or supplying text for standard emails and other office documents by predicting the next word or phrase in a sentence string. Better predictions augment human decision-making, allowing some elements of many jobs to be automated. For example, AI can improve diagnostic testing and augment physician accuracy in interpreting medical images (Einav et al. 2018; Agarwal et al. 2023). Even if AI tools do not fully replace physicians (who are still better at understanding the broader context, including patient history), they can increase treatment efficiency, which may lead to other downstream impacts.

Since most jobs involve prediction and decision-making, the impact of AI is likely to be widespread (Deming 2021). Brynjolfsson, Mitchell, and Rock (2018) build an occupation-level measure of “suitability for machine learning” (SML) and find that ML technologies can potentially replace tasks in a broad range of low-, medium-, and

high-paying jobs. Eloundou et al. (2023) find that as many as 80 percent of US workers could have at least 10 percent of their job tasks automated by LLMs like ChatGPT.

Still, the automation of individual job tasks does not necessarily reduce employment and may even lead to job gains in some sectors of the economy. The net impact of AI on employment depends on the balance between the replacement of existing job tasks and the impact of productivity gains from automation (if any) on the *total* demand for labor

in a sector. In principle, being able to automate a previously onerous task could make workers so much more productive that the increased output offsets the fact that some of their work is now being done by a machine.

“The automation of individual job tasks does not necessarily reduce employment and may even lead to job gains in some sectors of the economy.”

In other words, if productivity gains from automation greatly increase the size of the pie, workers may still get a larger slice even if one or two slices now get eaten by robots. Acemoglu and Restrepo (2019) develop a formal model of this tradeoff and argue that the greatest risk for workers is from “so-so” technologies that replace labor without increasing productivity very much, allowing the displacement effect to dominate. One example they give is automated checkout in retail sales.

Three recent studies suggest that generative AI models can significantly increase workers’ productivity in writing, customer service, and programming tasks (Brynjolfsson, Li, and Raymond 2023; Peng et al. 2023; Noy and Zhang 2023). In all three cases, the authors found that generative AI tools had a bigger positive impact on workers who were less productive at baseline, which decreased inequality in performance relative to a workplace without AI (Autor 2024). This finding suggests some reason for optimism about the labor market effects of AI.

A final consideration is whether AI-generated productivity gains will translate into increased demand for the goods and services in a sector. In other words, as we become more prosperous, do productivity gains lead us to want more of what is being produced? Or do we simply spend a smaller share of our income on it over time?

A classic example is agriculture. Farmers became vastly more productive with the adoption of agricultural tools that were powered by GPTs like steam and electricity. Yet because consumers only need so much food, productivity improvements in agriculture caused us to spend less income on food over time, leading to declining employment in farming. A similar dynamic likely holds for manufactured goods, although in both cases US firms are serving a global market, with countries at different stages of economic development. However, it is unlikely that productivity gains in education and healthcare would greatly reduce employment in those sectors, because as we become richer, we want to learn more and live longer. In other words, education and health are more income elastic than food and (perhaps) manufactured goods are.

Our best guess is that AI will end up having broad-based, long-run impacts on the labor market, similar to other GPTs of the past. If AI leads to labor market disruption, how will we know? What will be the early warning signs? The discussion above suggests that we should be looking for two patterns in the data. First, we should expect to see increased investment in new technologies and a J-curve pattern of productivity growth in AI-exposed sectors. Second, we should expect large but steady declines in employment share for AI-exposed jobs, especially jobs in sectors where the income elasticity of demand is low (for example, we don’t want much more as we become richer).

While it is far too early to know whether AI will have a comparable impact to GPTs of the past, we present some early evidence below of recent employment trends in the US economy.

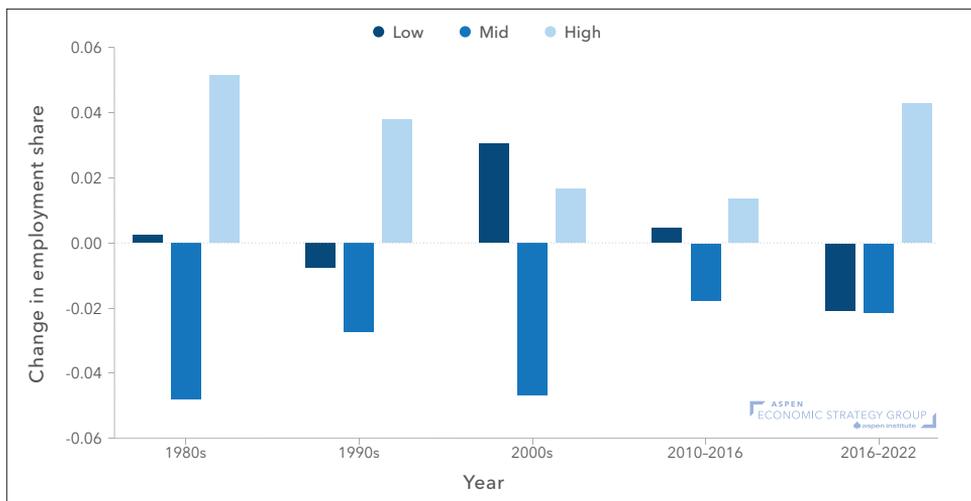
4. Four recent trends in the US labor market

Trend 1: Job polarization has been replaced by general skill upgrading

The first labor market trend is job polarization, with barbell-shaped employment growth at the bottom and the top of the wage distribution and declines in the middle (Autor, Katz, and Kearney 2006; Autor 2019). Job polarization has unfolded over several decades, holds across many different advanced economies, and has been linked to technological changes such as automation in manufacturing and the proliferation of office software.

Figure 4 divides jobs into low-, middle-, and high-paying occupations based on average wages and computes the change in employment share in each category over the last four decades. These results are the same as what is shown in Aspen Economic Strategy Group papers by Autor (2020) and Sacerdote (2020), except those authors only analyzed data through 2016. This figure replicates and extends Autor’s (2019) results, which group 2000–2016 into a single year rather than computing separate results for the 2000–2010 and post-2010 periods.

Figure 4: Employment growth by job skill level, 2000–2022



Notes: Calculations are based on the 1980, 1990, and 2000 US censuses (with a 5 percent state sample taken from each) and ACS data from 2010 and 2022, sourced via IPUMS (Ruggles et al. 2024; Flood et al. 2023). Occupations are harmonized to two-digit SOC codes using the IPUMS occ1950 encoding and methodology used in Autor and Dorn 2013; a detailed methodology is described in the data appendix. Samples are restricted to workers aged 18 to 64 in noninstitutional quarters who provide nonmilitary occupational responses. See the appendix for exhaustive definition of each category.

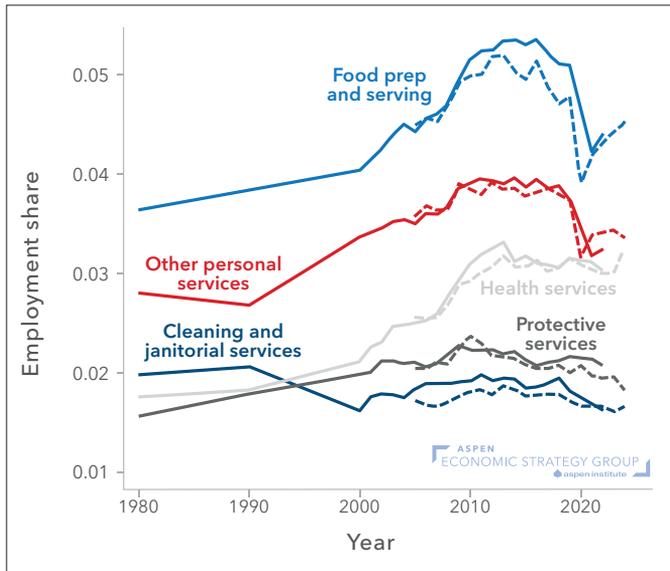
As figure 4 shows, employment growth was highly polarized in the 2000–2010 period, with large gains at the bottom and the top of the wage distribution and declines in the middle. Polarization continued from 2010 to 2016, although to a much lesser degree than in the decade before. However, the labor market has stopped polarizing since 2016. Between 2016 and 2022, low-skilled and middle-skilled jobs both declined by about 2 percentage points, while employment in high-skilled occupations increased by slightly more than 4 percentage points. Thus, employment growth since 2016 looks more like the kind associated with skill upgrading than the kind indicating polarization. We find similar results using the CPS, which collects data through 2024; however, we use the census and ACS in figure 4 because their occupation codes are more consistent over time.

Trend 2: Flat or declining employment in low-paid service occupations

A key explanation for employment polarization in the first decade of the 2000s was the rapid growth of service sector jobs, which replaced middle-skilled (often unionized) production jobs and offered lower wages and fewer employment protections (Autor and Dorn 2013). Figure 5 plots employment in five categories of service sector occupations from 1980 to the present— health support (including home aides), protective services, food preparation and service, cleaning and janitorial services, and personal-care occupations like barbers, fitness instructors, and childcare workers. The solid lines plot data from the census and ACS, while the dashed lines use CPS data through 2024.

Across all five categories, the growth of service jobs stalled in the early 2010s and was flat for most of the rest of the decade. Employment in food service and personal-care occupations fell rapidly in 2020 as a result of the COVID-19 pandemic and had recovered only partly by 2024. With the lone exception of health support occupations, service sector employment is now similar to what it was twenty years ago, early in the first decade of the 2000s. Service occupations have given back nearly all of the rapid job growth they experienced during that decade.

While the decline in service occupations may be indirectly related to technological change, the more proximate cause is rising wages for lower-paid work and increased labor market tightness (Autor, Dube, and McGrew 2023). In theory, rising labor costs could spur automation at the lower end of the job market, as with self-service checkout (for example, Acemoglu and Restrepo 2018).

Figure 5: The rise and fall of employment growth in service occupations, 1980–2024

Notes: Solid lines represent ACS calculations; dashed lines represent CPS calculations. ACS calculations are based on the 1980, 1990, and 2000 US censuses (with a 5 percent state sample taken from each), and ACS data from 2001 to 2022 (except 2021). CPS calculations are based on Current Population Survey (CPS) monthly data, aggregated by year, from 2005 to 2024. Both datasets are sourced via IPUMS (Ruggles et al. 2024; Flood et al. 2023). Occupations are harmonized across decades using the IPUMS occ1990 encoding. Samples are restricted to workers aged 18 to 64 in noninstitutional quarters who provide nonmilitary occupational responses. CPS samples are additionally restricted to respondents who self-report as being employed. See the appendix for exhaustive definitions of each category.

Trend 3: Rapid employment growth in STEM occupations

STEM (science, technology, engineering, and math) jobs shrank as a share of US employment between 2000 and 2012, while employment in non-STEM professional occupations grew rapidly (Deming 2017). Deming (2017) documents this surprising fact and argues that the demand for social skills is rising because social interaction is required for complex work and is not easily automated by current technologies.

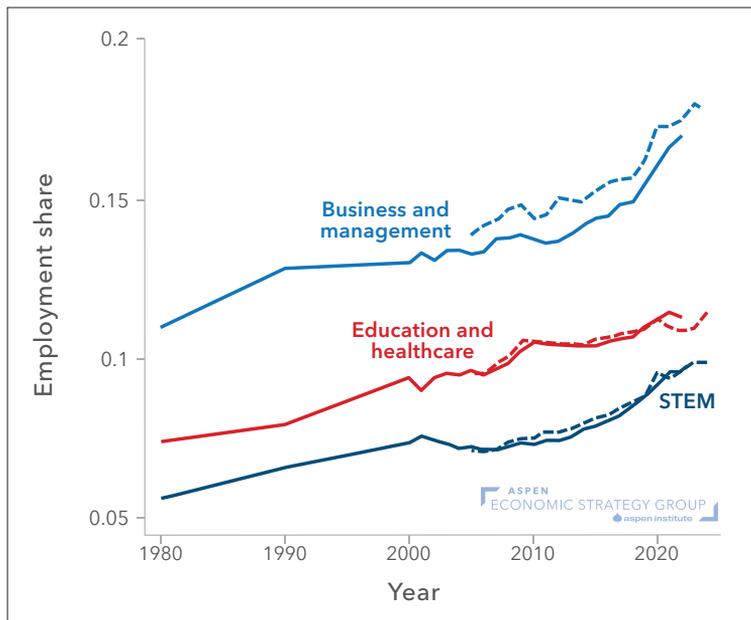
Figure 6 shows that while the growth in social skill-intensive management, business, education, and healthcare jobs has continued, STEM employment is now also increasing rapidly after having declined in the first decade of the 2000s. The share of all employment in STEM grew from 6.5 percent in 2010 to nearly 10 percent in 2024. About 60 percent of this growth is concentrated in computer occupations like

software developers and programmers, although employment has also grown across a wide range of science and engineering occupations as well.

STEM employment growth has accelerated especially quickly in the last five years. Moreover, the rapid employment growth in business and management jobs is concentrated in occupations like science and engineering managers, management analysts, and other business operations specialists.

Increased employment in STEM occupations is also matched by increased capital investment in AI-related technologies. Figure 7 plots private fixed investment in software and information processing equipment and research and development (R&D) as a share of GDP from 1947 to 2023. Investment in software and information processing peaked in 2000 at nearly 4.5 percent and then fell precipitously during the bursting of the dot-com bubble and the 2001 recession.

Figure 6: Employment growth in professional occupations, 1980-2024



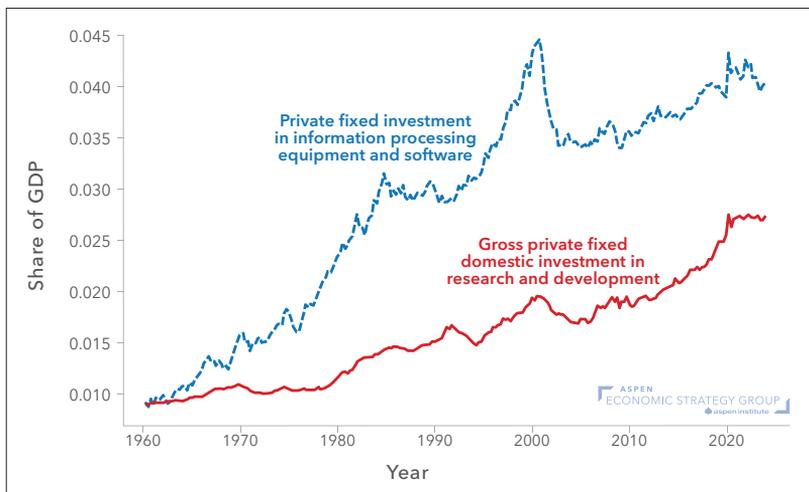
Notes: Solid lines represent ACS calculations; dashed lines represent CPS calculations. ACS calculations are based on the 1980, 1990, and 2000 US censuses (with a 5 percent state sample taken from each), and ACS data from 2001 to 2022 (except 2021). CPS calculations are based on Current Population Survey (CPS) monthly data, aggregated by year, from 2005 to 2024. Both datasets are sourced via IPUMS (Ruggles et al. 2024; Flood et al. 2023). Occupations are harmonized across decades using the IPUMS occ1990 encoding. Samples are restricted to workers aged 18 to 64 in noninstitutional quarters who provide nonmilitary occupational responses. CPS samples are additionally restricted to respondents who self-report as being employed. See the appendix for exhaustive definitions of each category.

However, it has rebounded in recent years and is now back above 4 percent. Private nonresidential R&D spending as a share of GDP is now at an all-time high of nearly 2.9 percent.

Much of this spending surge is driven by the increased computing power (or *compute*) necessary to train frontier AI models. Sevilla et al. (2022) find that the number of FLOPs (floating-point operations, a benchmark for computational complexity) required to train the latest Deep Learning models has increased exponentially in recent years. These models require thousands of high-performance graphical-processing units (GPUs) and can sometimes run for months at a time, requiring massive capital investments.

Figures 6 and 7 show that firms have been hiring and developing technical talent and making large investments in frontier technology. This increased investment in physical capital and technically skilled human capital may be a harbinger of future productivity growth, consistent with the J-curve hypothesis.

Figure 7: Investment growth in information processing, software, and R&D, 1960-2024



Notes: Data is sourced from the Federal Reserve Economic Data series from the Federal Reserve Bank of St. Louis (US Bureau of Economic Analysis 2024a and 2024b).

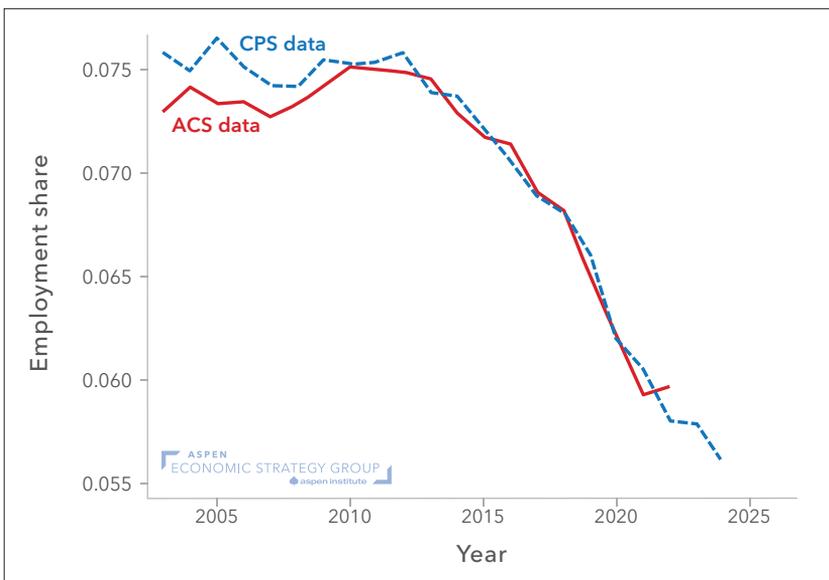
Trend 4: Declining employment in retail sales

ChatGPT was released in November 2022, so it is probably too early to see any direct impact of ChatGPT or other LLMs on the labor market. Still, companies have been using predictive AI models to optimize business operations since at least the mid-2010s. A particularly early adopting sector was online retail. E-commerce has

more than doubled as a share of all retail sales since 2015 (from 7 percent to 15.6 percent; see US Department of Commerce 2024). Online retailers like Amazon use AI to generate personal recommendations (and prices) based on customers' browsing and buying histories, and to predict local product needs and stock warehouses accordingly (Deming 2020).

Figure 8 plots the share of US employment in retail sales occupations from 2003 to 2024. Retail sales jobs held steady at around 7.5 percent of US employment between 2003 and 2013. Between 2013 and 2023, the US economy added more than 19 million jobs. Yet retail sales *declined* by 850,000 jobs over the same period, causing their share of employment to drop from 7.5 to 5.7 percent, a reduction of 25 percent in just a decade. The decline in retail sales jobs began before the pandemic but has accelerated in the last few years.

Figure 8: Employment share in retail sales occupations, 2000-2024



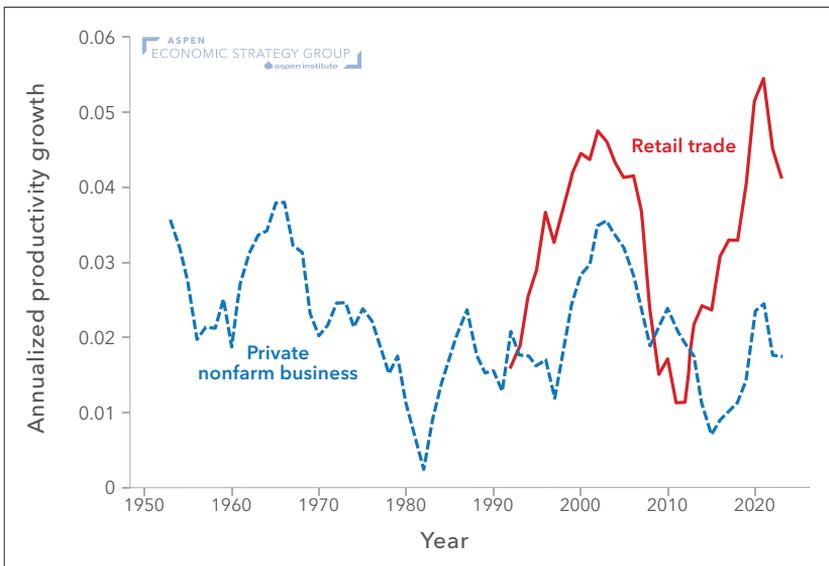
Notes: Calculations are based on ACS data from 2003 to 2022 (except 2021) and the CPS monthly data, aggregated by year, from 2003 to 2024, sourced via IPUMS (Ruggles et al. 2024; Flood et al. 2023). Occupations are harmonized across decades using the IPUMS occ2010 encoding. Samples are restricted to workers aged 18 to 64 in noninstitutional quarters who provide nonmilitary occupational responses and self-report as being employed. See the appendix for an exhaustive list of occupations included.

Labor productivity growth in retail trade has surged over this same period. Figure 9 plots annualized productivity growth in retail trade compared to all other nonfarm business sectors, from 1953 to 2024. Labor productivity growth was slower than average in retail trade during the period from 2007 to 2013 but has surged to between 4 and 5 percent in recent years, compared to the average of 2 percent across all

sectors. This pattern of rapidly declining employment and fast labor-productivity growth is very similar to what occurred with production jobs in manufacturing half a century ago.

Past episodes of technological disruption in the labor market have created new jobs as well, and there is some recent evidence of job growth as a downstream impact of AI-powered online retail. While the total number of jobs in the US economy grew by 14.5 percent between 2013 and 2023, the number of light delivery service truck drivers grew by 29 percent (Bureau of Labor Statistics 2024). This relative surge in short-haul truck driving is likely driven by last-mile package delivery from online retail sales. Similarly, the occupation “stockers and order fillers” grew from 1.8 million to more than 2.8 million jobs, an increase of nearly 60 percent. Much of this increase is fueled by job growth in the large warehouses maintained by online retailers.

Figure 9: US labor productivity growth, retail trade vs. all other industries, 1950-2024



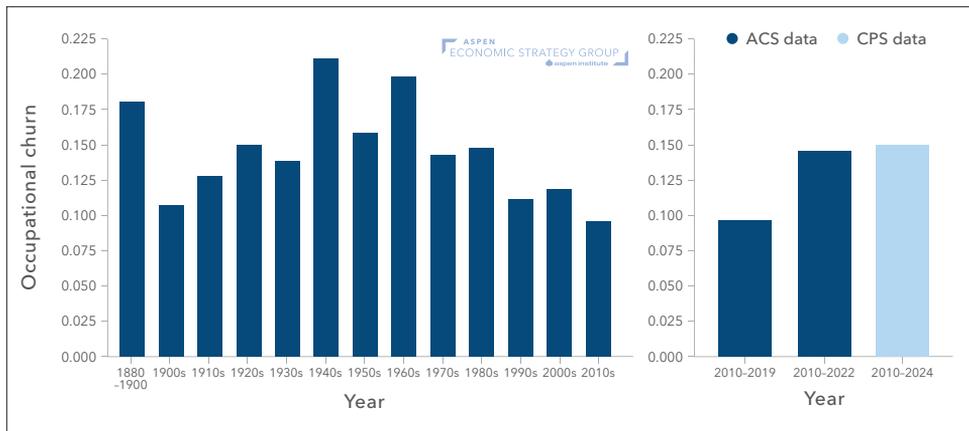
Notes: Calculations are based on Bureau of Labor Statistics “Historical Productivity and Cost Measures” for manufacturing sectors from 1949 to 2003 and “Labor Productivity and Cost Measures” for detailed sectors from 2004 to 2024. Displayed smoothed yearly growth rates are calculated by annualizing five-year growth rates (US Bureau of Labor Statistics 2024).

5. Conclusion

Each of these four trends—the end of polarization, stalled growth of low-paid service jobs, rapidly increasing employment in STEM occupations, and employment declines in retail sales—suggests that the pace of labor market disruption has accelerated in recent years. Figure 10 illustrates this acceleration systematically by recalculating our measure of labor market churn for the 2010–2024 period rather than for 2010–2019 as in figure 3.

According to ACS data, labor market churn between 2010 and 2022 (still only two years after the COVID-19 pandemic) was greater than that during any period since the 1970s. Using more recent data from the CPS shows that churn was slightly greater when measured over the period between 2010 and 2024. This finding is notable since churn from 2010 to 2019 was lower than it had been during almost every decade since the 1880s. This fact suggests that the post-COVID labor market has been especially volatile by historical standards, in part due to the four trends we identified earlier in the paper. A key unresolved question is whether the post-COVID changes to the labor market are here to stay.

Figure 10: Occupational Churn, 2010-2024



Notes: Occupational churn over each period is calculated as the sum of the total changes in absolute value of the employment share attributed to each two-digit SOC code. Employment shares for all bars except the last are calculated based on decadal US census data from 1880 to 2000 (except 1890) and ACS samples for 2010 and 2022, sourced via IPUMS (Ruggles et al. 2024). Employment share for 1900 is calculated using the 1880 and 1900 data. The last bar is calculated by aggregating the CPS monthly samples from 2010 and 2024. Occupations are harmonized across decades to two-digit SOC codes by extending the methodology used in Autor and Dorn 2013; a detailed methodology is described in the data appendix. Samples are restricted to workers aged 18 to 64 in noninstitutional quarters who provide nonmilitary occupational responses.

The recent evidence of labor market disruption should also be understood in the broader context of occupational upgrading in white-collar office work. Back-office administrative jobs such as financial, billing, and information-processing clerks shrank from 2 percent of all jobs in 1990 to 1.1 percent in 2024, a decline of nearly 50 percent. Front-office administration, including secretaries, administrative assistants, typists, and proofreaders accounted for 5.8 percent of all jobs in the US economy in 1990, compared to only 2.2 percent today. There are nearly five hundred thousand fewer secretaries and administrative assistants in the US labor market now than there were a decade ago. At the same time, management and business occupations have grown very rapidly. There were four million more managers and 3.5 million more business and financial operations jobs in the US in 2023 than there were in 2013.

While some of the growth in management occupations probably reflects a relabeling of existing workers (calling someone a manager rather than an office supervisor), it also reflects differences in job function. Notably, the occupation descriptions for office and sales supervisors emphasize monitoring of workers and office processes, whereas the descriptions of managerial roles are more likely to emphasize analysis, strategy, and decision-making (Deming 2021). Notably, about two-thirds of workers in management and business occupations have a four-year college degree, compared to less than one-third of sales and administrative-support workers, and they earn much higher average wages.

History tells us that we should expect the occupational upgrading of office work to continue. In some sense, sales and administrative-support jobs serve intermediary functions. The purpose of a sales worker is to connect firms to consumers who wish to buy their products. The purpose of an administrative-support worker is to reduce communication and coordination frictions between customers and the firm or between workers within the firm.

Personalized pricing algorithms and product recommendations, inventory management, oral and written transcription, and automated scheduling are some of the many sales and administrative-support office innovations made possible by AI. In each case, the purpose of the innovation is to increase productivity by smoothing the transmission of information within firms and between firms and the labor market. As AI technology improves, these innovations may lead to declining employment in sales and administrative-support occupations.

However, the impact of AI on professional and managerial workers is less clear. Early research studies and casual observation suggest that LLMs and other AI tools can replace highly skilled knowledge workers in some job tasks. A reasonable prediction is that the tasks replaced by AI will soon become commodified by the labor market.

These tasks include writing business plans, generating good ideas for article headlines, and writing or translating software code. The remaining tasks—analysis, decision-making, and adjudicating between the conflicting perspectives and desires of co-workers—are likely to become highly valuable as a result. While AI certainly helps with these tasks, the demand for good ideas and cogent analysis of complex counterfactual thought experiments (for example, assessing the likely impact of different business decisions, product strategies, etc.) may be nearly unlimited. At least in the near term, AI is more likely to ratchet up firms' expectations of knowledge workers than it is to replace them. In that case, policy solutions such as increased public investment in STEM education and training and reskilling will be necessary to help workers adapt to and effectively use new technologies.

“At least in the near term, AI is more likely to ratchet up firms' expectations of knowledge workers than it is to replace them.”

References

- Abraham, Katharine G., and Melissa S. Kearney. 2020. "Explaining the Decline in the US Employment-to-Population Ratio: A Review of the Evidence." *Journal of Economic Literature* 58, no. 3 (September): 585–643. <https://doi.org/10.1257/jel.20191480>.
- Abramitzky, Ran, Leah Platt Boustan, and Katherine Eriksson. 2014. "A Nation of Immigrants: Assimilation and Economic Outcomes in the Age of Mass Migration." *Journal of Political Economy* 122, no. 3 (June): 467–506. <https://doi.org/10.1086/675805>.
- Acemoglu, Daron, and Pascual Restrepo. 2018. "Low-Skill and High-Skill Automation." *Journal of Human Capital* 12, no. 2 (Summer): 204–32. <https://doi.org/10.1086/697242>.
- Acemoglu, Daron, and Pascual Restrepo. 2019. "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal of Economic Perspectives* 33, no. 2 (Spring): 3–30. <https://doi.org/10.1257/jep.33.2.3>.
- Agarwal, Nikhil, Alex Moehring, Pranav Rajpurkar, and Tobias Salz. 2023. "Combining Human Expertise with Artificial Intelligence: Experimental Evidence from Radiology." Working paper no. 31422. Cambridge, MA: National Bureau of Economic Research, July. <https://www.nber.org/papers/w31422>.
- Agrawal, Ajay, Joshua S. Gans, and Avi Goldfarb. 2019. "Artificial Intelligence: The Ambiguous Labor Market Impact of Automating Prediction." *Journal of Economic Perspectives* 33, no. 2 (Spring): 31–50. <https://doi.org/10.1257/jep.33.2.31>.
- Autor, David H. 2015. "Why Are There Still So Many Jobs? The History and Future of Workplace Automation." *Journal of Economic Perspectives* 29, no. 3 (Summer): 3–30. <https://doi.org/10.1257/jep.29.3.3>.
- Autor, David H. 2019. "Work of the Past, Work of the Future." *AEA Papers and Proceedings* 109 (May): 1–32. <https://doi.org/10.1257/pandp.20191110>.
- Autor, David H. 2020. *The Faltering Escalator of Urban Opportunity*. Washington, DC: Aspen Economic Strategy Group, July 8. https://live-aspen-economic-strategy-group.pantheonsite.io/wp-content/uploads/2020/12/Autor_Urban-Opportunity.pdf.
- Autor, David H. 2024. "Applying AI to Rebuild Middle Class Jobs." Working paper no. 32140. Cambridge, MA: National Bureau of Economic Research, February. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4722981.
- Autor, David, Caroline Chin, Anna Salomons, and Bryan Seegmiller. 2024. "New Frontiers: The Origins and Content of New Work, 1940–2018." *Quarterly Journal of Economics* advance online publication (March 15). <https://doi.org/10.1093/qje/qjae008>.
- Autor, David H., and David Dorn. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review* 103, no. 5 (August): 1553–97. <https://doi.org/10.1257/aer.103.5.1553>.
- Autor, David, Arindrajit Dube, and Annie McGrew. 2023. "The Unexpected Compression: Competition at Work in the Low Wage Labor Market." Working paper no. 31010. Cambridge, MA: National Bureau for Economic Research, March. <https://doi.org/10.3386/w31010>.

- Autor, David H., Lawrence F. Katz, and Melissa S. Kearney. 2006. "The Polarization of the U.S. Labor Market." *American Economic Review* 96, no. 2 (May): 189–94. <https://doi.org/10.1257/000282806777212620>.
- Barrero, José Maria, Nicholas Bloom, and Steven J. Davis. 2023. "The Evolution of Work from Home." *Journal of Economic Perspectives* 37, no. 4 (Fall): 23–50. <https://doi.org/10.1257/jep.37.4.23>.
- Bartel, Ann, Casey Ichniowski, and Kathryn Shaw. 2007. "How Does Information Technology Affect Productivity? Plant-Level Comparisons of Product Innovation, Process Improvement, and Worker Skills." *Quarterly Journal of Economics* 122, no. 4 (November): 1721–58. <https://doi.org/10.1162/qjec.2007.122.4.1721>.
- Bick, Alexander, and Kevin Bloodworth II. 2023. "Labor Force Participation and Hours Worked Recovery: U.S. vs. Europe." *On the Economy* (Federal Reserve Bank of St. Louis blog), December 21. <https://www.stlouisfed.org/on-the-economy/2023/dec/labor-force-participation-hours-recovery-us-vs-europe>.
- Binder, Ariel J., and John Bound. 2019. "The Declining Labor Market Prospects of Less-Educated Men." *Journal of Economic Perspectives* 33, no. 2 (Spring): 163–90. <https://doi.org/10.1257/jep.33.2.163>.
- Bresnahan, Timothy. 2010. "General Purpose Technologies." In *Handbook of the Economics of Innovation*, vol. 2, edited by Bronwyn H. Hall and Nathan Rosenberg, 761–91. Amsterdam: Elsevier. [https://doi.org/10.1016/S0169-7218\(10\)02002-2](https://doi.org/10.1016/S0169-7218(10)02002-2).
- Bresnahan, Timothy F., and M. Trajtenberg. 1995. "General Purpose Technologies 'Engines of Growth'?" *Journal of Econometrics* 65, no. 1 (January): 83–108. [https://doi.org/10.1016/0304-4076\(94\)01598-T](https://doi.org/10.1016/0304-4076(94)01598-T).
- Brynjolfsson, Erik, and Lorin M. Hitt. 2003. "Computing Productivity: Firm-Level Evidence." *Review of Economics and Statistics* 85, no. 4 (November 1): 793–808. <https://doi.org/10.1162/003465303772815736>.
- Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond. 2023. "Generative AI at Work." Working paper no. 31161. Cambridge, MA: National Bureau of Economic Research, revised November. https://www.nber.org/system/files/working_papers/w31161/w31161.pdf.
- Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock. 2018. "What Can Machines Learn, and What Does It Mean for Occupations and the Economy?" *AEA Papers and Proceedings* 108, no. 5 (May): 43–47. <https://doi.org/10.1257/pandp.20181019>.
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2019. "Artificial Intelligence and the Modern Productivity Paradox." In *The Economics of Artificial Intelligence: An Agenda*, edited by Ajay Agrawal, Joshua Gans and Avi Goldfarb, 23–60. Chicago: University of Chicago Press. <https://doi.org/10.7208/chicago/9780226613475.003.0001>.
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2021. "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies." *American Economic Journal: Macroeconomics* 13, no. 1 (January): 333–72. <https://doi.org/10.1257/mac.20180386>.
- Crafts, Nicholas. 2004. "Steam as a General Purpose Technology: A Growth Accounting Perspective." *Economic Journal* 114, no. 495 (April): 338–51. <https://doi.org/10.1111/j.1468-0297.2003.00200.x>.

- David, Paul A. 1990. "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox." *American Economic Review* 80, no. 2 (January): 355–61.
- Deming, David J. 2017. "The Growing Importance of Social Skills in the Labor Market." *Quarterly Journal of Economics* 132, no. 4 (November): 1593–640. <https://doi.org/10.1093/qje/qjx022>.
- Deming, David J. 2020. "The Robots Are Coming. Prepare for Trouble." *New York Times*, January 30. <https://www.nytimes.com/2020/01/30/business/artificial-intelligence-robots-retail.html>.
- Deming, David J. 2021. "The Growing Importance of Decision-Making on the Job." Working paper no. 28733. Cambridge, MA: National Bureau of Economic Research, April. <https://doi.org/10.3386/w28733>.
- Du Boff, Richard B. 1979. *Electric Power in American Manufacturing, 1889–1958*. New York: Arno Press.
- Einav, Liran, Amy Finkelstein, Sendhil Mullainathan, and Ziad Obermeyer. 2018. "Predictive Modeling of U.S. Health Care Spending in Late Life." *Science* 360, no. 6396 (June 29): 1462–65. <https://doi.org/10.1126/science.aar5045>.
- Ellingrud, Kweilin, Saurabh Sanghvi, Gurneet Singh Dandona, Anu Madgavkar, Michael Chui, Olivia White, and Paige Hasebe. 2023. *Generative AI and the Future of Work in America*. New York: McKinsey Global Institute, July 26. <https://www.mckinsey.com/mgi/our-research/generative-ai-and-the-future-of-work-in-america>.
- Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2023. "GPTs Are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models." *arXiv.org*, last revised August 21. <https://doi.org/10.48550/arxiv.2303.10130>.
- Flood, S., M. King, R. Rodgers, S. Ruggles, J. R. Warren, D. Backman, A. Chen, G. Cooper, S. Richards, M. Schouweiler, and M. Westberry. 2023. *Integrated Public Use Microdata Series, Current Population Survey: Version 11.0 [dataset]*. Minneapolis, MN: IPUMS. <https://doi.org/10.18128/D030.V11.0>.
- Frey, Carl Benedikt, and Michael A. Osborne. 2017. "The Future of Employment: How Susceptible Are Jobs to Computerisation?" *Technological Forecasting and Social Change* 114 (January): 254–80. <https://doi.org/10.1016/j.techfore.2016.08.019>.
- Goldin, Claudia. 2006. "The Quiet Revolution That Transformed Women's Employment, Education, and Family." *American Economic Review* 96, no. 2 (May): 1–21. <https://doi.org/10.1257/000282806777212350>.
- Goldin, Claudia Dale, and Lawrence F. Katz. 2007. "Long-Run Changes in the Wage Structure: Narrowing, Widening, Polarizing." *Brookings Papers on Economic Activity* 2007, no. 2: 135–65. <https://doi.org/10.1353/eca.2008.0000>.
- Lipsey, Richard G., K. Alec Chrystal. 1995. *An Introduction to Positive Economics*. 8th ed. Oxford: Oxford University Press.
- Muro, Mark, Sifan Liu, Jacob Whiton, and Siddharth Kulkarni. 2017. *Digitalization and the American Workforce*. Washington, DC: Brookings Institution, November. <https://www.brookings.edu/articles/digitalization-and-the-american-workforce/>.
- Neal, Derek, and Armin Rick. 2014. "The Prison Boom and the Lack of Black Progress after Smith and Welch." Working paper no. 20283. Cambridge, MA: National Bureau of Economic Research, July. <https://doi.org/10.3386/w20283>.

- Noy, Shakked, and Whitney Zhang. 2023. "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence." *Science (American Association for the Advancement of Science)* 381, no. 6654 (July 13): 187–92. <https://doi.org/10.1126/science.adh2586>.
- Olmstead, Alan L., and Paul W. Rhode. 2001. "Reshaping the Landscape: The Impact and Diffusion of the Tractor in American Agriculture, 1910–1960." *Journal of Economic History* 61, no. 3 (September): 663–98. <https://doi.org/10.1017/S0022050701030042>.
- Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer. 2023. "The Impact of AI on Developer Productivity: Evidence from GitHub Copilot." *arXiv.org*, February 13. <https://doi.org/10.48550/arxiv.2302.06590>.
- Rasmussen, Wayne D. 1982. "The Mechanization of Agriculture." *Scientific American* 247, no. 3 (September): 76.
- Rosenberg, Nathan, and Manuel Trajtenberg. 2004. "A General-Purpose Technology at Work: The Corliss Steam Engine in the Late-Nineteenth-Century United States." *Journal of Economic History* 64, no. 1 (March): 61–99. <https://doi.org/10.1017/S0022050704002608>.
- Steven Ruggles, Sarah Flood, Matthew Sobek, Daniel Backman, Annie Chen, Grace Cooper, Stephanie Richards, Renae Rodgers, and Megan Schouweiler. IPUMS USA: Version 15.0 [dataset]. Minneapolis, MN: IPUMS, 2024. <https://doi.org/10.18128/D010.V15.0>.
- Sacerdote, Bruce. 2020. *Is the Decline of the Middle Class Greatly Exaggerated?* Washington, DC: Aspen Economic Strategy Group, December 10. <https://live-aspen-economic-strategy-group.pantheonsite.io/wp-content/uploads/2020/12/2.-Sacerdote.pdf>.
- Sevilla, Jaime, Lennart Heim, Anson Ho, Tamay Besiroglu, Marius Hobbhahn, and Pablo Villalobos. 2022. "Compute Trends across Three Eras of Machine Learning." In *2022 International Joint Conference on Neural Networks (IJCNN)*. Padua, Italy: IJCNN, 1–8. <https://doi.org/10.1109/IJCNN55064.2022.9891914>.
- Solow, Robert M. 1987. "We'd Better Watch Out [Review of *We'd Better Watch Out*]." *New York Times Book Review*, July 12: 36.
- US Bureau of Economic Analysis. 2024a. "Gross Private Domestic Investment: Fixed Investment: Nonresidential: Intellectual Property Products: Research and Development [Y006RC1Q027SBEA]." FRED, Federal Reserve Bank of St. Louis, updated June 27. <https://fred.stlouisfed.org/series/Y006RC1Q027SBEA>.
- US Bureau of Economic Analysis. 2024b. "Private Fixed Investment in Information Processing Equipment and Software [A679RC1Q027SBEA]." FRED, Federal Reserve Bank of St. Louis, updated June 27. <https://fred.stlouisfed.org/series/A679RC1Q027SBEA>.
- US Bureau of Labor Statistics. 2024a. "Occupational Employment and Wage Statistics." US Department of Labor. https://www.bls.gov/oes/current/oes_data.htm.
- US Bureau of Labor Statistics. 2024b. "Productivity: Tables." US Department of Labor. <https://www.bls.gov/productivity/tables/home.htm>.
- US Department of Commerce. 2024. *Quarterly Retail E-Commerce Sales: 1st Quarter 2024*. Washington, DC: US Census Bureau, May 17. https://www.census.gov/retail/mrts/www/data/pdf/ec_current.pdf.

- Wike, Richard, and Bruce Stokes. 2018. In *Advanced and Emerging Economies Alike, Worries about Job Automation*. Washington, DC: Pew Research Center, September 13. <https://www.pewresearch.org/global/2018/09/13/in-advanced-and-emerging-economies-alike-worries-about-job-automation/>.
- Williams, Heidi L. 2013. "Intellectual Property Rights and Innovation: Evidence from the Human Genome." *Journal of Political Economy* 121, no. 1 (February): 1–27. <https://doi.org/10.1086/669706>.
- World Economic Forum. 2023. *The Future of Jobs Report 2023*. Cologny, Switzerland: World Economic Forum, April 30. <https://www.weforum.org/publications/the-future-of-jobs-report-2023/digest/>.