

Labor Markets and the Future of Work

By: Josh Rowe, CFA, PhD, Managing Director of Research & Family Office Investment Strategy, HB Wealth

This is the third installment in our four-part series examining artificial intelligence (AI) through a wide-angle economic lens. In the [first essay](#), we explored the historical development of artificial intelligence and the recurring cycles of enthusiasm and disappointment that have shaped the field's development. [The second](#) examined how major technological innovations diffuse through the economy and why productivity gains often take years or decades to fully materialize. If you have not yet read those pieces, we encourage you to explore them as well.

Across this series, we examine four related questions: the origins of modern AI technologies, the economic dynamics that shape technological adoption, the implications for labor markets and the future of work, and the potential consequences for productivity and long-term economic growth.

Executive Summary

The third installment of this series examines how generative artificial intelligence may reshape labor markets and the future of work.

Historically, technological revolutions have triggered fears of mass unemployment that have ultimately always proved exaggerated. While individual occupations were displaced, new industries and roles emerged, and the growth generated by new technology lifted living standards across the board. Economists often point to this historical pattern as evidence that automation ultimately complements human labor rather than replacing it.

Generative AI may challenge that assumption. Unlike previous waves of automation that targeted routine manual work, modern AI systems are increasingly capable of performing cognitive tasks once dominated with highly educated professionals. Lawyers, analysts, programmers, marketers, and other knowledge workers now face potential automation of the key workflows that define their careers.

Early research suggests the effects may be uneven. In routine tasks, AI tools appear to help less experienced, junior workers improve productivity. In more complex, judgment-based work, however, the technology may amplify differences in expertise, benefiting highly skilled professionals while displacing junior workers whose roles

traditionally served as training grounds. From drug discovery to scientific research to legal and financial services, AI increasingly doesn't just automate; it breaks new ground. Benefits are likely to accrue to the highest-paid, most senior talent and to the owners of capital

At the same time, corporate incentives are increasingly aligned with labor substitution. Firms are adopting AI tools to reduce headcount, slow hiring, and improve margins. This shift is occurring alongside structural labor market trends such as declining mobility, rising education costs, and slower workforce adaptation.

Whether generative AI will obviate most human skills, necessitating a one-day workweek and universal basic income, or will simply shift the frontiers of what humans and computers can do together cannot be known with certainty. The trajectory will depend not only on technological capabilities but also on policy choices, organizational strategies, and how institutions adapt to the changing nature of work.

For investors, these dynamics matter because changes in productivity and labor costs ultimately shape corporate profitability, economic growth, and long-term market returns.

About the Author

Josh Rowe, Managing Director of Research at HB Wealth, wrote a PhD thesis in the history and economics of technology, focusing on computer automation of office work in the 20th century. He has studied the history of AI, venture capital's funding of technological innovation, and the impact of technological change on financial markets—both as a resident of the ivory tower and as an investor. This surprising moment in history is the first time that he can say with any confidence that the years he spent in libraries and databases working on a doctoral dissertation might be of any practical use. He used AI in organizing and editing these essays, but the ideas (right and wrong) here are his own.

Introduction: Technological Substitution or Human-Computer Complementarity

The previous section traced a familiar arc of history: a potentially transformative invention captures the public imagination, and initial expectations of revolutionary change are disappointed by the messy reality of implementation. Over time, organizations learn to restructure around the new technology, novel uses, methods, and expertise are developed, legacy systems are gradually updated or replaced, and new generations of workers come of age. Only then do we see the efflorescence of productivity. Stock markets are, of course, discounting machines and (especially when interest rates are low) are apt to look boldly to the future. While the ChatGPT

moment of 2022 may not yet have altered day-to-day life for most office workers, precocious narratives about an Eden of material abundance just over the horizon are quite in keeping with the historical pattern. Generative AI has still to enter the typical prolonged adolescence of diffusion lags, trial-and-error, and creative destruction. We should also look to historical precedent to understand the darker side of technological golden ages; fears about worker displacement have reemerged right on cue, often trumpeted most loudly by the technologists themselves.

At base, productivity is an accounting ratio: output divided by hours worked. When economists discuss AI's potential to boost productivity, they typically emphasize the numerator—more output, better quality, faster turnaround. The techno-optimist case assumes that human labor adapts in tandem, shifting toward higher-value activities as routine tasks are automated. In the economist's telling, capital and labor may substitute for one another, but they are necessarily combined to grow output—and as output grows, there is more for everyone to do. Certain types of jobs, even certain industries, are eliminated of course. But a dynamic labor economy reallocates people to where labor is scarce. Factory workers in Indiana move to Nevada to build houses, and so on. Meanwhile, the benefits of greater output can be redirected by farsighted policymakers to help workers level up. Construction workers get an associate's degree in 3D printing, for example. Everyone learns to code. Workers and machines become complements; the pie grows; everyone benefits eventually.

But ratios always represent tradeoffs—how much you get versus how much you put in. Gains in a ratio can also come from the denominator. If AI enables the same output with fewer human hours, the ratio improves even as labor demand contracts. In this scenario, the economics of AI are just as compelling, but the spoils are shared differently. They accrue to capital, to the owners of the technology, and to the shrinking cadre of workers whose skills remain complementary. Everyone else faces displacement, wage compression, and the erosion of bargaining power that comes from competing against machines.

This distinction—numerator versus denominator, augmentation versus substitution—is more than academic. It shapes who benefits from technological change, the incentives of businesses, and whether or to what degree the social fabric is torn. It also informs the policy choices society must make: prohibition, regulation, redistribution, or so-called “directed technological change.”

The history of past technological revolutions counsels against too much alarmism. Those who battled against automation, who [rioted or broke the power looms](#) of Victorian factories, did not halt innovation. Social reformers of the 19th century, such as Marx, Engels, and Luxembourg, who viewed technology as an instrument of class

warfare, have not thus far been vindicated. Luddites, however justified their immediate grievances, proved to be little more than a narrow interest group. Technology and automation did not liquidate labor; they raised living standards for workers and managers over a century and a half of technological progress. This was the [hope of liberal economists like David Ricardo](#) and Adam Smith – technology was a general good that could cheapen output and improve real incomes for all classes.¹ Their views have become standard orthodoxy.

And yet history is not destiny. There are reasons to wonder whether generative AI might break the pattern. Perhaps we are only now experiencing the effects of an economy in thrall to the technologists' labor-saving vision. For four decades, economists have documented a persistent decline in labor's share of national income across the developed world. In the United States, the [share of GDP flowing to workers](#) in the form of wages and benefits fell from roughly 65 percent in the early 1980s to around 58 percent by the late 2010s. Loukas Karabarbounis and Brent Neiman, in [influential work](#) published in 2014, traced the decline across more than fifty countries and linked it to the falling price of capital goods, particularly information technology, which made it increasingly attractive for firms to substitute machines for workers. MIT economist David Autor and collaborators have documented a related phenomenon: the [rise of "superstar firms"](#) with high productivity, high profits, and relatively low labor intensity. These juggernauts' growing market share mechanically reduces the aggregate share of labor even when most individual firms' practices remain unchanged.

This trend predates generative AI, but large language models and other more specialized models threaten to accelerate it. Without necessarily lowering total employment (or even, in many cases, real wages), each previous wave of automation, industrial machinery, computing, robotics, pushed the labor share lower by enabling capital to perform tasks that previously required human effort. Owners of capital reaped the productivity benefits. Even as the pie grew, the distribution of wages also became more uneven. What economists call "skill-biased" technical change meant that technologies like computers made highly educated workers more valuable, allowing them to command a premium. Less skilled workers were increasingly commoditized and saw many of their jobs eliminated.

¹ Ricardo, notably, revised his views as the Industrial Revolution progressed. He saw that automation could hit certain industries or even regions disproportionately and society might fail to reallocate labor to complementary sectors. This story is told in Daron Acemoglu and Simon Johnson's [recent article](#) in the *Annual Review of Economics*, and in their 2023 book *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity* (PublicAffairs).

What is distinctive about generative AI is how it is poised to upend the long-standing pattern of skill-biased technological change. For the first time, automation is extending its reach into skilled cognitive work, the domain where labor had largely retained its advantage. The entire labor force is now exposed. Owners of capital stand to gain all the more.

Notwithstanding our worries, recent data complicates this supposedly ineluctable trajectory. Real wages in the United States have grown robustly since 2022, with the largest gains accruing to workers at the bottom of the income distribution—a reversal of the pattern that prevailed for decades. Tight post-pandemic labor markets gave workers bargaining power they had not enjoyed in a generation; quit rates rose, employers competed for talent, and wage compression narrowed income inequality. Whether this represents a durable shift or a cyclical blip remains unclear. The labor share ticked upward during the recovery, though it remains well below its mid-century peak. If AI deployment accelerates and labor markets soften, the window of worker leverage may close as quickly as it opened.

Recurring Fears of Technological Unemployment

Popular anxieties about technological unemployment are as old as technological progress itself. In 1589, William Lee invented the stocking frame knitting machine and petitioned Queen Elizabeth I for a patent. She refused, reportedly saying: “Consider thou what the invention could do to my poor subjects. It would assuredly bring them ruin by depriving them of employment, thus making them beggars.” The machine spread anyway, and the knitting industry eventually employed far more workers than before—but Lee died in poverty, his fears having outlived his hopes.

The Luddites of early nineteenth-century England, often caricatured as irrational technophobia, were in fact skilled craftsmen facing genuine economic displacement. The textile workers who smashed machines in 1811 were not superstitious primitives afraid of progress; they were artisans whose livelihoods were being destroyed by factory mass production. Their protests were ultimately futile, looms spread, factories expanded, and the textile industry grew tremendously, but for those caught in it, the transition was brutal. New factory jobs that eventually bore no resemblance to the skilled craft work that had been obliterated. While enabling an industrial economy, the shift to factories reduced many workers’ livelihoods, lowered life expectancy, and gutted worker bargaining power. In the long run, of course, economists were correct. The gains from the industrial revolution were repaid handsomely, including to ordinary citizens through cheaper, better products and, ultimately, higher-value jobs in manufacturing and services.

Similar panics attended each subsequent wave of mechanization. John Maynard Keynes, writing in 1930, was somewhat less sanguine about technological unemployment than his classical brethren, but he, too, hoped that technology would enable such economic growth to pay social dividends. In his famous essay “Economic Possibilities for Our Grandchildren,” he foresaw a future in which machines would satisfy most human needs and the “great economic problem” would be solved—leaving humanity to grapple with the unfamiliar challenge of abundant leisure. Keynes may have been characteristically prescient about the long run but mistaken about the timing: here we are, ninety-five years later, and most adults still spend the majority of their waking hours at work. To date, three-day workweeks, “universal basic income” and other salves to displaced workers have been [political losers](#).

By the 1960s, fears of machines replacing workers had acquired a more modern idiom. Postwar factory production was increasingly incorporating sophisticated robots. Computers, personified by Big Blue (IBM) and the “giant brain” thinking machines of Control Data Corporation and Cray, dominated the news. In a 1962 address, President John F. Kennedy identified “maintaining full employment at a time when automation is replacing men” as the “major domestic challenge of the sixties.” A presidential commission warned that the new technologies “can be used to displace workers from productive activity, and this displacement is already taking place.” Secretary of Labor Willard Wirtz testified before Congress that the “threat of automation” to employment was “one of the most serious problems we face.” Computers brought, for the first time, significant automation to the office. Wall Street Journal headlines announced that computers ([“electronic wizards”](#)) had assumed many tasks previously performed by clerical workers, leading to “a sharp slowdown in white-collar hiring.”

The complete computerization that many predicted of the office never arrived. Nor did any pickup in measured productivity materialize on schedule. Studies from the late 1980s revealed that the service sector, which economist Stephen Roach described as “the most heavily endowed with high-tech capital,” registered the worst productivity performance of any major sector during the computer age. Factory robots did more with fewer human hands. But knowledge work, and other service sectors like healthcare, education, finance, and retail services, remained stuck behind the technology curve, suffering from what economist William Baumol termed [“cost disease.”](#) As Solovian paradoxes tend to do, Gen AI’s productivity impact will also trail the technology’s narrative ubiquity.

Why Economists Are Still (Mostly) Sanguine

Labor economists have developed a standard repertoire of responses to recurring fears about automation. The arguments are worth understanding because they form a benchmark against which to evaluate the effects of contemporary AI in the workplace and will help us calibrate how much alarmism is warranted.

First is the *lump of labor fallacy*—the mistaken belief that there exists only a fixed amount of work to be done, such that labor-saving technology must then mechanically reduce total employment. This error reflects short-term, linear thinking: at any given moment, in any particular industry, automation can indeed eliminate specific jobs. But the economy as a whole is not a zero-sum system. When technology reduces the cost of producing goods or services, the savings flow somewhere—to lower prices for consumers, higher profits for firms, or both. Those gains generate new spending, which creates new demand, which creates new employment. Furthermore, technology opens new product markets and new job roles, including ones scarcely imagined by advocates of automation. Textile workers displaced by power looms found new work in the expanded factories producing cheaper cloth for expanded markets. The typists made redundant by word processors found employment in the vastly larger information economy that computers enabled.

The second argument concerns the distinction between *tasks* and *jobs*. A job is a bundle of tasks; technology rarely automates entire bundles at once. Instead, it automates particular tasks within jobs, changing the composition of work without eliminating the work itself. The economist David Autor of MIT, perhaps the most influential contemporary researcher on automation and labor markets, has documented this pattern across multiple technological transitions. Sixty percent of employment in the United States in 2018 consisted of job categories that did not exist in 1940. “Airplane designer” entered the census in the 1950s; “conference planner” arrived in the 1990s. Each wave of automation destroyed old tasks while generating conditions for new ones to emerge.

The third argument emphasizes *complementarity*—ways in which human labor and technological tools enhance each other’s productivity rather than substituting for one another. Air traffic controllers, for instance, work with sophisticated computer systems that process flight data and flag potential conflicts. The technology augments human judgment rather than replacing it; controllers make more decisions, and better ones, than they could without the tools. Their wages remain high because the technology increases their productivity without commoditizing their skills. By contrast, self-checkout systems at grocery stores automate the mechanical aspects

of a cashier's job, scanning items, calculating change, while leaving behind only the residual tasks of handling exceptions and deterring theft. The skill requirements drop; job openings and wages stagnate.

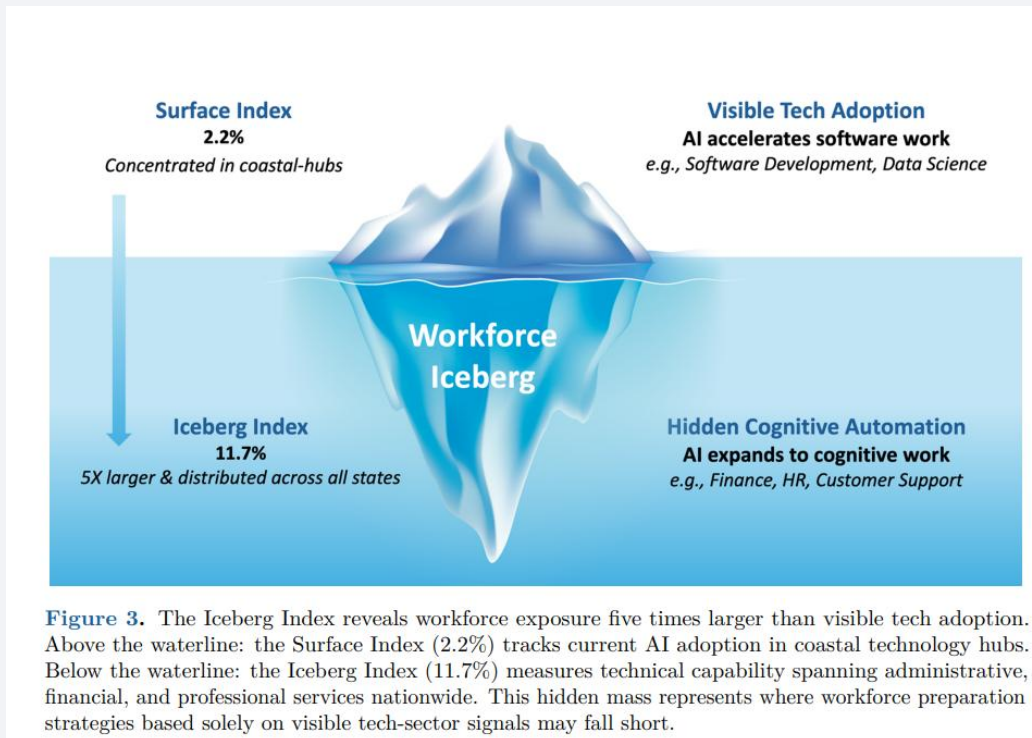
These arguments command substantial empirical support. Across the long sweep of economic history, labor markets have proven remarkably adaptable. Technological unemployment has been a transitional phenomenon—painful for those caught in the transition, but transitional nonetheless. New industries have emerged to absorb workers displaced from old ones. Living standards have risen across generations. The Luddites were wrong in the long run, even if they were right to be worried in the short run.

Why This Time Might be Different

There are grounds to believe that generative AI poses challenges that differ in kind, not merely degree, from previous waves of automation.

One critical difference is scope. Previous automation technologies targeted specific categories of work. Mechanical power displaced human muscle; assembly-line robots replaced repetitive manual tasks; early computing automated routine administrative work like bookkeeping and calculation. Generative AI threatens a far broader swath of economic activity. Research suggests that somewhere between 60 and 80 percent of global jobs involve tasks that could be partially or fully automated by current AI systems.² That is a far broader potential impact zone than factory robots or even personal computers. It does not mean 40 to 80 percent of jobs will disappear, but it does mean that virtually no white-collar occupation is entirely immune. Lawyers, doctors, software engineers, graphic designers, financial analysts, journalists, marketers, customer service representatives—the exposure is nearly universal across knowledge work.

² For reviews of economists' recent findings on AI disruption exposure among the current job landscape, see research [from the Brookings Institution](#) and [the IMF](#).



MIT researchers' ["Iceberg Index"](#)

The breadth of exposure can have a chilling effect on hiring even before substantial automation occurs. Employers uncertain about which roles AI will transform in the coming years become reluctant to make long-term commitments, especially those with healthcare and pension benefits attached. Entry-level positions, where training costs are high and the payoff period extends years into an uncertain future, face particular pressure. Why invest in developing a junior associate if the work that the associate would do in five years might be handled by software?

A second difference from prior regime changes concerns speed. Historical analogies to electricity or computing emphasize the long lags between technological invention and productivity impact—the decades required for complementary infrastructure, organizational redesign, and human capital formation. Generative AI has the potential to significantly compress those timelines as we detailed earlier. Unlike factory electrification, which required rebuilding physical plants, or enterprise computing, which required custom software development and systems integration, large language models can be accessed through natural-language interfaces and integrate with existing office tools. A knowledge worker with a browser and an API key can deploy frontier AI capabilities today, without waiting for IT departments to implement new systems or managers to redesign workflows.

Current models' accessibility is double-edged. Ease-of-use accelerates the diffusion of productivity gains but also accelerates displacement pressures. Historical transitions allowed long intervals for workers to retrain, for educational institutions to adapt, for new industries to emerge, and absorb displaced labor. If generative AI transforms knowledge work over years rather than decades, that adjustment time compresses accordingly. It is also worth noting that the U.S. labor market [may have become significantly more rigid](#) in recent decades; whereas intercity mobility and expanding access to higher education helped create a dynamic, fluid labor force through much of the second half of the 20th century, today U.S. workers face high housing costs (especially in the fastest job-creating markets), higher costs and declining enrollments in higher and professional education, and, in some cases, less flexible rules about hiring and firing. Economists' ideal of a frictionless labor market enabling dynamic reallocation of talent has likely become less, not more realistic.

The third difference can be found in the direction of technology's impact. Previous automation waves followed a rough pattern: routine work was automated first, with the highest-skilled workers benefiting most from new tools. This led to the skill-biased change documented by labor economists: rising returns to education, growing wage premiums for cognitive work, a hollowing out of middle-skill employment. Knowledge workers were largely on the winning side of these transitions.

With today's AI tools, the mechanism works differently. Automating technology excels precisely the tasks that constitute the bulk of mid-level knowledge work: synthesizing information, drafting documents, analyzing data, and writing code. At A&O Shearman, a major law firm, [AI tools now handle](#) much of the routine work once performed by associates and paralegals. The software can analyze contracts, compare them with past deals, and suggest revisions in under thirty seconds. At ServiceNow, a business software company, [executives estimate](#) that more than 80 percent of customer service cases for some clients no longer require human involvement. The grunt work of white-collar industries, the work that trains the next generation of professionals, is being automated first.

Software engineering, ironically, may be experiencing this dynamic most acutely today. The profession that created the AI boom has also been the first to feel its labor market effects. Technology sector [layoffs have been extensive](#); hiring of entry-level developers has slowed dramatically; bootcamps and computer science programs report declining placement rates. 2025 saw high-profile RIFs (reductions in force) exceeding 50,000–100,000, with AI cited as a direct factor by firms like Microsoft, Amazon, Salesforce, and Dropbox. The workers who built the tools are discovering that the tools work on them, too.

The recent swoon in software stocks, triggered by the realization that agentic AI can now write serviceable code faster and cheaper than human developers, carries echoes of older upheavals. During the Industrial Revolution, the rise of textile factories did not just eliminate handweaving but hollowed out an entire social class of skilled artisans whose craft had earned them a good living and social esteem for centuries. The cloth that emerged from the factories was not quite as fine, but it was vastly cheaper and produced in quantities the old guild system could never match. Software engineers—long the aristocrats of the knowledge economy, commanding salaries and stock options that placed them among the most prosperous workers in history—may be experiencing the [early stages of a similar displacement](#). The question now extending to lawyers, accountants, analysts, and consultants is not whether AI will eliminate these professions entirely, but whether it will compress them into a smaller, more precarious class, much like the manufacturing workers a generation ago.

Replacing skilled labor creates a further human capital dilemma. Professions like law, engineering, medicine, consulting, and investment banking have traditionally functioned as apprenticeship systems. Junior employees do tedious work (i.e. document review, financial modeling, code debugging) while learning the judgment and expertise required for senior roles. If the tedious work is automated, what happens to the talent pipeline? How do associates become partners if the bottom of the pyramid is hollowed out? And what happens to the elite narrative about automation's benefits when high-skill workers, politically influential and overrepresented in the press, in lobby groups, and academia, see that automation is, this time, coming for them or their children?

Who Benefits, Who Loses

Early research on AI and productivity generated optimism about equalizing effects across income distribution. A 2023 study by Erik Brynjolfsson of Stanford and collaborators at MIT found that generative AI tools boosted productivity by 14 percent for novice customer support workers, helping them resolve queries faster and more effectively. Experienced workers, by contrast, saw little benefit—the AI was reinforcing approaches they already used. Similar results were seen in studies of writing and legal work: lower-skilled workers showed the greatest improvements when using AI assistance.

These findings implied that AI might narrow skill gaps by transferring best practices from high performers to lower performers. If a junior lawyer could produce partner-quality memos with AI assistance, the technology might be democratizing rather than polarizing. In a world where technology has been a lightning rod at the center of debates around economic inequality, these early hints are reassuring.

More recent research complicates this picture. Early studies examined relatively simple, well-defined tasks. When researchers look at more complex, judgment-intensive work, the distributional effects reverse. In a study of materials science research, Aidan Toner-Rodgers of MIT found that AI tools nearly doubled the productivity of top researchers while having no measurable impact on the bottom third. Elite scientists could identify promising AI suggestions and discard poor ones; less effective researchers struggled to discern signal from noise. In finance, [sophisticated investors using AI](#) to analyze earnings calls achieved nearly 10 percent higher returns than less sophisticated investors with the same tools. In entrepreneurship, high-performing Kenyan business owners raised profits by over 15 percent with AI assistance, while struggling business owners saw profits fall.

The story connecting these data points is troubling. For routine tasks with clear right answers such as customer service queries, basic legal documents, and standard business writing, AI can elevate lower performers toward a common baseline. But for complex work requiring judgment, creativity, and domain expertise, AI amplifies existing differences. High performers use the tools to become more productive; lower performers use them badly or are simply displaced by them. As [one investment manager](#) put it: "Really smart people who may be bored with analyzing routine earnings releases will benefit the most. The skill that is going to be rewarded most in the short run is imagination in finding creative ways to use AI."

The distributional implications cut multiple ways. By skill level: workers with strong foundations and good judgment are positioned to gain leverage from AI models, while those with less tacit knowledge and training may find their tasks automated without gaining compensating abilities. By age and experience: senior professionals with deep expertise can evaluate AI outputs and integrate them into workflows; junior professionals who would normally be building that expertise through routine work may never develop it. By industry: perversely, the sectors that seemed most insulated from previous automation waves, professional services, creative industries, software development, now appear most exposed.

Labor Force Headwinds and Competitive Strategy

Theoretical concerns about labor displacement are one thing; corporate behavior is another. Boardrooms today are signaling the beginning of a competitive arms race to drive down labor costs through AI substitution. In previous eras, executives discussed layoffs with well-chosen euphemisms—"restructuring," "right-sizing," "realigning resources." There was something faintly embarrassing about workforce reductions; they signified failure or retrenchment. That norm appears to be changing. Executives in the age of workplace AI rather herald shrinking headcounts as achievements to be celebrated.

Wells Fargo's CEO Charlie Scharf [recently boasted](#) that the bank had reduced headcount for twenty consecutive quarters—five years of continuous workforce shrinkage—totaling a 23 percent reduction over the period. “Attrition is our friend,” he told investors. Bank of America's Brian Moynihan reminded shareholders that headcount had fallen from 300,000 to 212,000 since he became CEO in 2010. “We just got to keep working that down,” he added. Union Pacific reported that labor productivity had reached record quarterly highs as staff shrank by three percent. Verizon's CEO told investors the company's headcount is “going down all the time.” Intel announced plans to cut 15 percent of its workforce and eliminate layers of middle management. Microsoft cut 15,000 workers in two months. Amazon told staff that AI would “likely lead to a smaller workforce.”

The shift reflects both a cooling labor market with employers regaining bargaining power after the pandemic's tight conditions and a new managerial ideology. [As one communications consultant put it](#), there is now “almost a moral neutrality to headcount reductions. Being honest about cost and headcount isn't just allowed—it's rewarded.” Executives see large workforces less as assets and more as impediments. Even firms, like Google, Danaher, or Hilton that have long cultivated reputations as desirable places to work and compete with best-in-class perks to attract and retain human capital, may be caught up in the competitive contest to run leaner: if you don't automate your labor force and improve margins, your competitor will and you will lose share.

It's not just big companies. Startups are increasingly finding ways to develop, ship, and sell products with many fewer workers. As engines of the new, AI-native economy, younger businesses are where economists might hope redundant labor might be upskilled and redeployed. That's [not happening](#). Garry Tan of Y Combinator, reported that the biggest change in the startup ecosystem is “how effective companies can be with hardly any employees, using AI to grow. They can go from zero to \$10 million a year in revenue, sometimes in the course of less than 12 months, and they can do it with less than 10 people.”

Even in mundane, mature and labor-intensive service industries, venture capitalists are hoping to consolidate fragmented markets onto AI-native platforms. In areas like accounting and tax, property and HOA management, and medical billing, investors are explicitly deploying the language and logic of technological automation: “headcount,” “gross margin” and “efficiency.” “If 80% of the work will be done by AI, it can have an 80% to 90% gross margin,” [says Navin Chaddha](#), managing director of the VC firm Mayfield's managing director. “You could have blended margins of 60% to 70% and produce 20% to 30% net income.”

Large-scale layoffs are not, in fact, the primary mechanism by which companies are shrinking. More common are hiring freezes, position eliminations when workers leave, and the quiet non-replacement of roles that were once considered essential. The cumulative effect is the same: fewer workers doing more work, with the gains flowing to shareholders and senior management.

Molly Kinder, a senior fellow at Brookings who studies AI's effects on workers, has been struck by how little public criticism companies receive when executives telegraph AI-related staffing reductions. In industries where AI is displacing white-collar functions, unionization is typically scant, and employees may be powerless to push back. "I'm worried it's happening in plain sight with no blowback, no pushback, and it's going to become the norm," she said. "I don't think that's good news for the American worker." To some extent, Silicon Valley-style technofuturism and AI maximalism have captured the zeitgeist, at least in America and on Wall Street and Capitol Hill.

Everyone a CEO

At a glamorous AI Summit held in Paris in early 2025, technology executives competed to issue the grandest predictions for AI's role of rewiring the global economy. Sundar Pichai, Alphabet's CEO, called AI "the most profound shift of our lifetimes." Dario Amodei of Anthropic predicted "the largest change to the global labor market in human history." Sam Altman of OpenAI wrote that "in a decade perhaps everyone on earth will be capable of accomplishing more than the most impactful person can today."

Altman's prediction echoes a hopeful vision that accompanied the technology's early commercial rollout: that of AI as Great Equalizer. Models can handle protein folding and poetry writing; it ought to be able to put such powers within the grasp of ordinary people. Jensen Huang of NVIDIA envisioned a future in which all workers "are going to be CEOs of AI agents."

Evidence increasingly suggests otherwise. History's lessons teach that technological upheavals favor the skilled. In the Industrial Revolution, engineers who mastered new machinery saw their wages soar while routine laborers lost ground. The computer age rewarded software engineers and rendered typists obsolete. AI may yet follow that same playbook, even while lowering skill barriers to use. Today, its augmentation effects are felt unequally, benefiting those with the judgment, agility, and expertise to navigate complex, information-rich environments. Employers feel pressure not only to adopt the technology, but to be judicious in hiring layers of junior talent. Whether the sectoral effects (more widespread adoption and labor pressure in cognitively intense industries) or

the hierarchical effects (disproportionate impact on less experienced workers) dominate, may determine the path of inequality.

In Huang's vision, every worker will manage AI systems the way executives manage human organizations—delegating tasks, evaluating outputs, and making strategic decisions. But management is itself a skill. Not everyone is equally good at it; not everyone can learn to be equally good at it. If AI makes every worker a potential CEO, the most talented workers will simply become the best CEOs. Hierarchy reasserts itself in a new form.

What We Are Not Discussing

This analysis has deliberately avoided a set of questions that dominate much AI discourse: whether and when artificial general intelligence will arrive, when machines might exceed human capabilities across all domains, and what happens when (or if) we build systems that recursively improve themselves. These are fascinating questions for futurists and philosophers. They are not the essential questions imminently facing investors or policymakers.

The maximalist case—associated with figures like Ray Kurzweil and, more recently, some AI lab executives—holds that we are approaching a singularity after which any predictions become unreliable. If that scenario proves correct, the labor market questions discussed here become moot, along with most other questions about economic and social organization. We would be dealing with something genuinely unprecedented.

Gary Marcus and other AI skeptics make the opposite argument: that current systems, however impressive, are bound by fundamental limitations that prevent generalization beyond their training distributions. Large language models excel at pattern-matching over text; they do not reason, plan, or understand in ways that would enable true generality. On this view, the labor market impacts will be real but bounded—particular tasks automated, whole professions disrupted, but not the wholesale transformation that maximalists envision.

Interestingly, two of the foremost innovators of “deep learning,” the approach that underlies current AI models, are [Geoffrey Hinton](#) and [Yann LeCun](#). They find themselves in different camps with regard to the timeline of AGI's arrival. If they can't agree, investors and managers would be wise to hedge their bets.

We do not know where the capability curve flattens; we do not know which jobs AI will struggle with and which it will master; we do not know how quickly organizational and regulatory complements will develop. What we can observe is the present direction of change, the patterns in early adoption data, and the historical precedents that inform reasonable expectations.

And what those observations suggest is that the labor market impacts are likely to be substantial, concentrated in knowledge work, unequally distributed across skill levels and career stages, and faster-moving than historical transitions—even if they fall far short of the maximalist predictions. That is enough to warrant serious concern, and enough to raise questions about how policy might shape the trajectory.

Policy Choices and Path Dependence

One of the most important insights from the economics of technical change is that outcomes are not solely determined by technology itself. The same technologies can be deployed in ways that augment human capabilities or substitute for them, that concentrate gains among elites or distribute them broadly, that create new categories of good jobs or eliminate existing ones. Policy choices, tax structures, regulatory frameworks, educational investments, and labor protections shape which path predominates.

Daron Acemoglu and Simon Johnson have been particularly insistent on this point. The current tax code, they argue, subsidizes capital investment relative to labor—encouraging automation even where human workers remain cost-effective. The bias is not inevitable; it reflects political choices about depreciation schedules, investment credits, and the relative taxation of labor versus capital income. A tax code that treated human and machine inputs more neutrally would shift the calculus toward augmentation rather than replacement.

Similarly, the allocation of research funding shapes the direction of technological development. Public and private R&D could prioritize technologies that enhance human productivity, developing tools that make workers more effective rather than rendering them unnecessary. Or it could prioritize technologies that minimize human involvement, treating labor as a cost to be eliminated. The choice is not foreordained.

Educational institutions face their own adaptation challenges. If entry-level knowledge work is being automated, the traditional pathway through professional training, law school to associate to partner, medical school to residency to practice, computer science degree to junior developer to senior engineer, may need fundamental rethinking. What should legal education look like if document review is fully automated? How should medical training change if most diagnostics are handled by AI? These are not questions that universities can answer overnight, but neither are they questions that can be deferred indefinitely.

Policy conversations are complicated both by uncertainty and the pace of current change, and by a lack of consensus on the endpoint. Should technology develop more slowly, with long lags before labor market impacts

materialize, aggressive intervention may be premature and hamper productivity. If adoption accelerates and the tools improvement arc continues at current rates, waiting for clear evidence of displacement will cause governments to act too late for a whole generation of workers.

There is one further question that the current productivity data cannot yet answer: whether this building is creating more stuff of value or is merely creating more stuff. The same tools that enable a product manager to prototype an app in an afternoon also enable the proliferation of what one observer refers to as “[busyware](#)”—the minor website update no one would miss, the custom dashboard built for an audience of one, the executive’s half-baked demo that engineers must then make real. Each may serve a purpose in the moment, but much of it will inevitably end up on the scrap heap. In a future of disposable code, the biggest productivity hack may turn out to be the restraint to know what not to build at all.

Business writers have observed a proliferation of what others are calling “workslop.” When the cost of producing content drops toward zero, the volume of content can explode. If AI can generate unlimited drafts, summaries, analyses, and communications at zero marginal cost, organizations may produce more of all of them—much of it mediocre, redundant, or actively counterproductive. Email volumes increase; meeting agendas proliferate; reports multiply. Knowledge workers find themselves not freed from drudgery but drowning in AI-generated material that must be read, evaluated, and responded to.

This dynamic is familiar from previous information technologies. Solow noted that some of the expected gains were offset by the costs of managing information abundance. Computers “pour out 7,300 miles of paper a day,” the *Wall Street Journal* reported in 1966, burdening executives with excessively detailed reports. The corporate equivalent of TMI, too much information, absorbed the productivity gains that computers were supposed to deliver. Especially when some significant proportion of that information is of low value.

Generative AI may produce similar phenomena at gigantic scale. Imagine a future where every meeting spawns an AI-generated summary, every project an AI-drafted status report, every question an AI-produced analysis. Aggregate output may increase while high-quality work is drowned in the soup of “deliverables”, or even actively discouraged. “Workslop” is a catchy name for this kind of pollution. Expect lots of TED talks about how to get more value out of your AI workflows. But they wouldn’t change incentives or the supply dynamics of low-effort, low-cost production.

The Denominator Revisited

We began with productivity as a ratio: output divided by hours worked. The question for labor markets is which term generative AI will primarily affect, and for whom. Optimists assume that AI will operate disproportionately on the numerator—enabling more and better output while human work adapts to higher-value activities. In this scenario, historical cycles repeat: technology destroys some jobs while creating others, transition is painful but temporary, and long-run complementarity between humans and machines delivers rising living standards across the income distribution, even if with some skew.

Pessimists worry that AI will work mainly through the denominator—reducing the human hours required to produce any given output, with gains accruing to capital and highly skilled labor while everyone else faces displacement, wage compression, and the erosion of bargaining power. In this scenario, something genuinely new is happening: the breadth of exposure, the speed of adoption, and the inversion of skill-bias toward knowledge work combine to produce labor market disruptions more severe and less readily absorbed than in any previous technological transition.

There is not yet any satisfying resolution to this debate. What we can say is that the early evidence is markedly less positive than in recent tech adoption cycles including personal computers, the Internet, and mobile computing. Distributional bias appears to run against junior knowledge workers and toward senior experts. This technology is diffusing faster than previous GPTs, while the adjustment mechanisms—retraining, reallocation, new job creation—operate at the same historical speed or slower. And much of corporate America is cheerleading a disruptive transition.

For investors, the labor market dynamics create a peculiar tension. The same forces that threaten widespread displacement—the breadth and rapid adoption of AI agents—are precisely what make the technology valuable to the companies deploying it. The headcount reductions that CEOs boast about are sources of earnings growth. The junior associates not hired are a cost savings. The productivity gains from automation accrued, at least initially, to shareholders. Benefits will largely concentrate among big companies with high-powered CTOs and capex budgets.

But markets exist within societies, and societies that experience widespread labor market disruption tend to birth political responses. If the labor share continues to decline, if entry-level knowledge workers face a generation of diminished opportunity, if the gains from AI flow overwhelmingly to capital and a narrow stratum of highly skilled labor, political economy will turn. One unresolved question is whether societies have the potential for productive

adaptation. Will we make investments in education, redesign professional training, enact holistic tax and regulatory reform to encourage inclusive augmentation, or will continue along the recent path of destructive populist backlash against science and technology.

History offers precedents for both scenarios. And for investors trying to assess AI's long-term trajectory, the labor market dynamics deserve as much attention as the technology itself.

This is the third article in our AI series, examining the implications of artificial intelligence for labor markets and the future of work. In the final installment, we will turn to the broader question of productivity and long-term economic growth.

If you have any questions, please reach out to your client service team, visit us at hbwealth.com, or call 404.264.1400.

Notes

1. Kennedy's remarks on technological unemployment appear in his address to the Economic Club of New York, December 14, 1962, and the context is explored in detailed contemporaneous coverage.
2. David Autor's research on job creation and task composition is central to modern labor economics. The finding that 60 percent of employment in 2018 consisted of job categories nonexistent in 1940 appears in his work with Anna Salomons and Bryan Seegmiller; see "New Frontiers: The Origins and Content of New Work, 1940–2018," *Quarterly Journal of Economics* (2024).
3. The Brynjolfsson, Li, and Raymond study of customer support workers was published as "Generative AI at Work," NBER Working Paper No. 31161 (2023), and represents some of the earliest rigorous evidence on AI productivity effects.
4. Erik Brynjolfsson's broader work on technology and productivity includes *The Second Machine Age* (with Andrew McAfee, 2014) and numerous academic papers spanning several decades.
5. Daron Acemoglu's estimates of AI's potential impact on employment structure derive from several [working papers](#) with David Autor and Pascual Restrepo.
6. The Toner-Rodgers study on AI and materials science research was circulated as an MIT working paper in 2024 and generated considerable attention for its finding that AI benefits primarily accrue to top performers.
7. Gary Marcus's critiques of large language models appear in numerous essays and in his book *Rebooting AI: Building Artificial Intelligence We Can Trust* (with Ernest Davis, 2019). His skepticism about naive extrapolation from current capabilities remains influential among AI critics.
8. Paul David's foundational work on technology diffusion and productivity lags remains "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox," *American Economic Review Papers and Proceedings* 80, no. 2 (1990): 355–361.

9. Stephen Mihm's discussion of Solow's Paradox and its application to AI appeared in Bloomberg Opinion, June 2025, drawing on his broader work on the history of technological expectations.
10. The Brookings Institution analysis by Molly Kinder and colleagues examines AI's labor market effects with particular attention to worker voice and bargaining power; ongoing research is available through the Brookings Metropolitan Policy Program.
11. Corporate statements from Wells Fargo, Bank of America, Intel, Verizon, and other companies cited here derive from earnings calls, investor presentations, and contemporaneous Wall Street Journal reporting, particularly Chip Cutter's coverage in July 2025.
12. Garry Tan's observations about AI-enabled startups appeared in podcast interviews in spring 2025 and reflect Y Combinator's assessment of changes in the startup ecosystem.
13. The concept of "workslop" as AI-generated busywork draws on emerging concerns in organizational studies; the term and concept are explored in Harvard Business Review's September 2025 coverage of AI-induced productivity paradoxes.

Important Disclosures

This article may not be copied, reproduced, or distributed without HB Wealth's prior written consent.

All information is as of the date above unless otherwise disclosed. The information is provided for informational purposes only and should not be considered a recommendation to purchase or sell any financial instrument, product, or service sponsored by HB Wealth or its affiliates or agents. The information does not represent legal, tax, accounting, or investment advice; recipients should consult their respective advisors regarding such matters. This material may not be suitable for all investors. Neither HB Wealth nor any affiliates make any representation or warranty as to the accuracy or merit of this analysis for individual use. This information contains forward-looking statements, predictions, and forecasts ("forward-looking statements") concerning our belief and opinions in respect to the future. Forward-looking statements involve risks and uncertainties, and undue reliance should not be placed on them. There can be no assurance that forward-looking statements will prove to be accurate, and actual results and future events could differ materially from those anticipated in such statements. Certain information herein is based on third-party sources believed to be reliable, but which have not been independently verified. Past performance is not a guarantee or indicator of future results; inherent in any investment is the risk of loss. Specific investments described herein do not represent all investment decisions made by the above date. The reader should not assume that investment decisions identified and discussed were or will be profitable. Specific investment advice references provided herein are for illustrative purposes only and are not necessarily representative of investments that will be made in the future. Investors are advised to consult with their investment professional about their specific financial needs and goals before making any investment decision. HB may hold positions (long or short) in the companies mentioned in this paper.