Economic
Policy
Institute

Unbalanced labor market power is what makes technology—including AI—threatening to workers

The best “AI policy” to protect workers is boosting their bargaining position

Report • By Josh Bivens and Ben Zipperer • March 28, 2024

Unbalanced labor market power is what makes technology—including AI—threatening to workers

The best “AI policy” to protect workers is boosting their bargaining position

Summary: The root causes of sluggish wage growth for most workers are intentional policy decisions that have led to an extreme imbalance of power between employers and typical workers–technological advances, like AI, have little to do with this and are too frequently invoked as a distraction from these deeper problems.

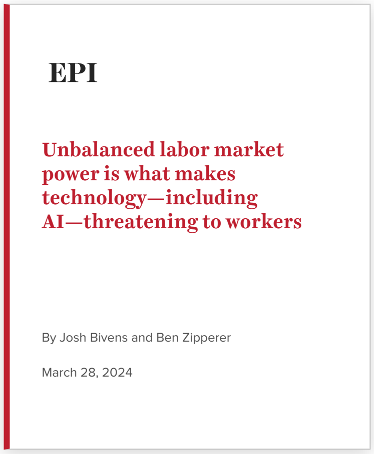
EPI

Unbalanced labor market power is what makes technology—including AI—threatening to workers

By Josh Bivens and Ben Zipperer

March 28, 2024

Read the full report
epi.org/278901

Key findings

- In recent decades, it is not technology, but institutional changes (like the decline of unions, the erosion of the federal minimum wage, and a change in macroeconomic policy priorities) that undercut typical workers’ leverage and bargaining power in labor markets.
- Economists typically measure technology as an increase in productivity. Yet productivity growth has not historically been associated with higher unemployment or higher inequality.
- Like any other technology, AI can be used as a zero-sum tool for increased employer control of work intensity and wages. However, it is the unbalanced power that is the root of this problem—not technological change per se, which could easily boost workers’ wages or make jobs easier in more balanced labor markets.

Why this matters

Outcomes for workers will not improve unless policymakers get the story right on technology, inequality, and labor market dysfunction. Efforts to blame inequality and unemployment on “technology” conveniently divert attention from the real cause of rising inequality and weak wage growth: excess employer power.

How to fix it

Each individual workplace will have its own challenges in integrating AI in ways that help rather than harm workers. Policymakers have almost no serious ability to address these workplace-specific challenges around AI. The best “AI policy” that they can provide is boosting workers’ power by improving social insurance systems, removing barriers to organizing unions, and sustaining lower rates of unemployment.

> “Bolstering workers’ bargaining power—not the newest AI development—should preoccupy policymakers.

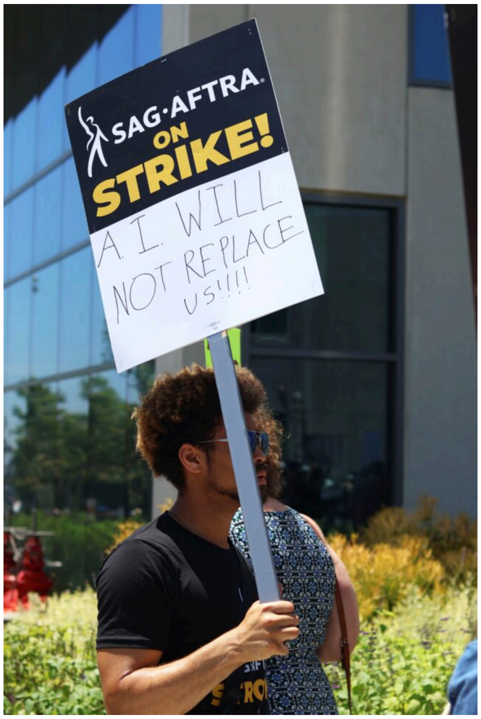

Many of the concerns raised recently about advances in artificial intelligence (AI)—for example, its implications for national security or media disinformation—are outside our areas of expertise. An area we *do* have considerable expertise to draw on is AI's potential effect on labor markets and our outlook might surprise some who have followed recent public debates: AI, like most technological advances, is unlikely to be a direct threat to the wages and employment of U.S. workers. Instead, it has the potential to raise these workers' living standards. Realizing this potential does not hinge on the specifics of AI policy, but instead on restoring the balance of economic power in key markets—especially the labor market.

Being relatively sanguine about the effect of technology and AI on labor markets does not imply that we think labor markets have been working well for U.S. workers. On the contrary, unemployment has been too high and wage growth too slow for decades. But the roots of labor market dysfunction—both past and future—have *very* little to do with technological changes. Instead, this dysfunction is driven by the concerted policy push to exacerbate the extreme imbalance of power between typical workers and the corporate managers and capital-owners who hire them.

It is important to get the facts and analysis right on the questions of why labor markets have not delivered enough jobs or acceptable wage growth, and what the real threats are to decent jobs in the future. Faddish debates about AI distract attention away from the more fundamental problem of imbalanced power in labor markets, pulling policy in less useful directions.

More specifically, we argue:

- Interpretations of past episodes of rising wage inequality—whether they were driven by changes in technology or changes in policy, institutions, and norms—differ enormously based on one's assessment of employers' ability to exercise power in labor markets. If this power is great, then policy, institutions, and norms have great scope to influence wage inequality. If instead employer power is limited, technological change becomes the major force driving inequality.

SECTIONS

1. Background on past waves of concern regarding technology and labor markets • 2
2. Competing models of the labor market • 6
3. How economists typically measure technological progress: productivity growth • 9
4. Can accelerating technological progress cause mass joblessness? • 11
5. Does technology reduce demand for workers without the right credentials or skills? • 19
6. Technology is a tool—the balance of labor market power determines who it helps • 22
7. Conclusion and policy implications—looking through the latest technological fad to see the real threats to workers' well-being • 26

Notes • 27
References • 29

- Technology manifests most directly in measured economic statistics as an increase in *productivity*—the amount of output generated in an average hour of work in the economy. Productivity growth has not historically been associated with higher unemployment or higher inequality, meaning that worries that technological change could be driving a jobless future have yet to materialize.
- Economic research claiming that the very rapid rise in wage inequality in the 1980s through the mid-2000s was caused by the rapid introduction of new technologies (mostly the spread of personal computers and other information and communications technologies) has not stood the test of time; few economists today would highlight the impact of technology alone as a driver of this inequality.
- While it is possible that technology can reduce the demand for specific jobs, these job losses can be more than counterbalanced by expanding employment in other sectors, as long as we maintain aggregate demand.
- In labor market models that allow for employer power, technological change in and of itself is largely neutral in its effect on the distribution of economic growth. But when employers exercise unbalanced power in wage-setting, they are often able to use new forms of technology to claim more of a firm's output at the expense of typical workers. However, it is the unbalanced power that is the root of this problem—not technological change per se, which could easily boost workers' wages if deployed in more balanced labor markets.
- Given this history of technology and labor markets, there is very little AI-specific *labor market policy* that will do much to help workers. Instead, policymakers should focus on broader policy levers to boost workers' leverage in wage bargaining that will aid workers in claiming the potential gains spurred by AI in the future and reclaiming lost ground from past periods of economic growth. AI-specific provisions in *workplace negotiations and collective bargaining agreements*, of course, make lots of sense. How AI—or any technological tool—can be deployed to raise productivity and foster broad-based wage growth instead of increasing employer control will be a crucial question for many workplaces. But the best *policy support* for this process that can be given by national policymakers is strengthening worker voice and power, not trying to micromanage how AI is used in specific workplaces.

Background on past waves of concern regarding technology and labor markets

Concerns that technological changes can cause labor market distress for workers has a long history. The term "Luddite," for example, has its origins in a movement of British textile workers in the early 19th century who opposed the introduction of new machinery they feared would displace their jobs.

In more recent decades, there have been waves of popular concern regarding technological advances as either a direct threat to workers' well-being or an enabler of other threats (like globalization). In the early 1980s, for example, the rise of personal computers raised fears of "technological unemployment."[1] In the early 2000s, IT-enabled growth of "white-collar offshoring" was cast as a major threat to U.S. workers.[2] In the 2010s, the introduction (real or imagined) of robots and autonomous vehicles was argued to imminently threaten huge swathes of the workforce.[3] And, of course, in the last year or two, advances in AI have spurred a multifaceted debate about its impact—including its potential labor market effects.

Much of this concern over technological changes in recent decades has coincided with undeniably bad outcomes for most workers in the U.S. labor market. The post-1979 period has seen unemployment at excessively high levels for extended periods and wage growth for typical workers slow dramatically relative to what prevailed in the first three decades following World War II. Wage growth has slowed even relative to the *much* slower pace of economywide productivity growth that has characterized the post-1979 period. Slow wage growth for most workers has led to sharply higher levels of wage inequality, along with a shift of income away from labor compensation and toward business incomes (particularly corporate profits).[4]

However, despite the concern about the effect of technological change on labor markets—and even despite the objectively poor performance of labor markets for most U.S. workers in recent decades—the effect of technological change has been generally positive when looked at from the perspective of the U.S. working class writ large. The anemic wage growth since 1979 for the typical worker would have been far smaller, and perhaps even negative, had there been no technological advances and no corresponding increase in labor productivity since that time.

But the full *potential* boost to living standards that technology could have provided has been more than swamped by the declining leverage and bargaining power of typical workers over this same period. The shift in labor market power away from typical workers and toward employers is the result of intentional changes in public policies, institutions, and norms. Key examples include failures to protect workers' right to organize unions from growing employer hostility, raise the federal minimum wage for long periods of time, and the maintain extended periods of very low unemployment. It is these intentional policy decisions, not technological progress itself, that have redistributed so much income away from typical workers and toward corporate profits and those at the very top of the pay scale (CEOS and other corporate managers, for example).[5]

Before walking through the economics and data supporting this statement, it is important to note one powerful piece of anecdotal evidence regarding the technological dog that *didn't* bark. We highlighted waves of concern about technological advancements in the 1980s, early 2000s, 2010s, and today. We could not find any serious wave of concern from the mid- to late 1990s. This should be strange. The 1990s saw the technological advance with by far the greatest effect on the economy in several decades—the introduction of the Internet and the rise of e-commerce—often at the expense of brick-and-mortar retailers. Unlike the other waves of popular concern surrounding technological changes, the rise of the Internet in U.S. economic life really did show up in key statistics (**Figure A** clearly shows a sharp uptick in productivity growth in the 1990s business cycle, for example). Yet the late 1990s (even in real time) was generally seen as a period of broad-based prosperity and healthy labor markets.[6] The explanation for this perception is simple: for the first time in decades, unemployment was driven low enough to generate opportunities for many who had been shut out of job markets and spur genuinely healthy wage growth for most workers. In short, technology—even the significant acceleration of technological advance in the late 1990s—was never really a headwind to decent labor market performance. Instead, the headwinds were all poor *policy* choices and changing some important ones (like allowing an extended period of very low unemployment) improved labor market performance radically, even in the midst of the most rapid technological change in decades.

Getting the story right on technology, inequality, and labor market dysfunction is crucially important for making the right policy decisions. Efforts to blame inequality and unemployment on bloodless, apolitical forces like "technology" constitute a convenient alibi for those social forces supporting the concrete policy changes that actually drove these outcomes. This technology alibi has been *extraordinarily* effective in distracting attention away from the major causes of rising inequality and anemic wage growth. As the debate over AI's potential effect on labor markets begins, this history of technology-as-alibi needs to be kept front and center in the minds of analysts and policymakers alike.

A concrete example illustrates how myopic focus on new technological trends can divert attention away from the root causes of labor market dysfunction. In the mid-2010s, long-term unemployment (unemployment spells exceeding six months) was particularly high and had been for years. Around this time, many employers were using automated data systems to sort through job applications. As the automated systems ranked job applicants, they were frequently programmed to instantly reject applicants who had not worked in the previous six months. This obviously exacerbated the problem of re-employment facing the long-term unemployed, and proposals were floated to bar employers from undertaking this kind of application sorting.[7]

But this proposed solution was severely flawed relative to the optimal response. The reason why long-term unemployment was high in the 2010s was because overall unemployment was high. Aggregate demand (spending by households, businesses, and governments) was too low to absorb enough willing workers to meaningfully push down unemployment (either short- or long-term). Policy efforts to boost aggregate demand could have quickly lowered overall unemployment, and long-term unemployment would have quickly followed suit. We know this is true because as unemployment fell steadily (if slowly) into the late 2010s, long-term unemployment fell even more rapidly.[8]

Essentially, a severely damaged macroeconomy was inundating employers with far more applications for each job than they felt capable of processing efficiently, so these employers used a technological advance (automated hiring software) as a shortcut for sorting applications based on long-term unemployment. Barring employers from using this coping strategy for dealing with the excess of applications over job vacancies would not have solved any society-wide problem. Employers still would have faced too many applicants per job and would likely have just moved onto some other application sorting shortcut. A common one was ratcheting up educational credentials required for the job despite the underlying work not really demanding these credentials.[9]

Crucially, while barring employers from using unemployment duration as a criterion in their automated application sorting processes might have resulted in some long-term unemployed worker getting a job, this job would have come directly at the expense of another worker who was also unemployed. Again, the main labor market problem in the mid-2010s was too few jobs per jobseeker. Changing how these too few jobs were allocated would have done little to improve aggregate human welfare over this period. But generating *more jobs* through expansionary macroeconomic policy would have solved this underlying problem and improved aggregate human welfare enormously. Focusing on the technological fad (automated hiring systems) and missing the deeper economic problem (a shortfall of aggregate demand) led to a much less constructive policy debate.

We worry that concerns about AI's potential effects on labor markets will prompt a rush to construct targeted AI-specific policies—as has happened over and over again in U.S. policy debates on technological change. These policies mostly will not materialize at all because policymakers will soon be distracted by the next fad. Even if some policies do get constructed, they would be mostly ineffective in making labor markets better for workers and will divert valuable attention away from other policies that would actually improve labor market functioning.

Is it possible we're wrong and AI will be the technological change that finally drives bad labor market outcomes for the vast majority? It's possible. But there's no evidence of it doing that yet and the nature and history of how technology affects labor markets argues that it is policies bolstering typical workers' bargaining position in labor markets—not the newest development in AI—that should preoccupy policymakers who aim to deliver better labor market outcomes for workers in the years to come.

Key definitions, issues, and questions about technological change and labor markets

The remainder of this report will focus on key concepts, definitions, issues, and questions about technological change and labor markets.

In section 2, we provide a brief overview of two competing models of the labor market. The choice of which model best describes the functioning of real-world labor markets is crucial in assessing how technological changes can affect labor market outcomes, and which influences (technology or institutional change) have driven historical trends in wage growth and inequality.

In section 3, we explain how economists tend to measure technological progress—as movements in productivity growth.

In section 4 we assess broad empirical correlations between faster productivity growth (or an increased pace of technological progress) and overall labor market outcomes.

In section 5 we evaluate the claims of some economists that particular forms of technological change have altered the relative demand for large classes of workers in competitive labor markets, and hence have driven much of the rise in inequality of pay seen in recent decades. We find these claims lacking in key evidence.

In section 6, we note some arguments surrounding the effect of technological change on labor markets that have not received enough attention from mainstream economists: the role of technology as a tool for employers to boost their leverage in pay-setting versus typical workers. However, we note that the root cause of this problem is unbalanced labor market power, not technology qua technology, which could in theory be just as easily used to boost workers' power as degrade it.

Finally, we sum up what this analysis implies for policy and what should preoccupy policymakers looking for real solutions to boost workers' pay and improve their labor market outcomes.

Competing models of the labor market

In recent decades, a key debate in labor economics has been determining which changes in the economy are responsible for the large rise in wage inequality since 1979. Since the late 1970s, only workers at the top of the wage distribution (those earning more than 90% of all other workers) have seen growth in wages that approaches growth in economywide productivity. Wage growth for workers below the 90th percentile has substantially lagged productivity growth.[10]

Two competing explanations for this rise in wage inequality are: first, technological change that has decreased the relative demand for less credentialed labor (sometimes called *skill-biased technological change*, or SBTC) and second, institutional changes (like the decline of unions, the erosion of the federal minimum wage, and a change in macroeconomic policy priorities) that undercut typical workers' leverage and bargaining power in the labor market. [11]

It is often underrecognized that the outcome of the debate over the sources of wage inequality hinges almost entirely on what one assumes is the correct underlying model of the labor market: one where labor markets are competitive and power is roughly balanced between workers and employers, or one where employers have structurally greater power than typical workers.

Those who emphasize technological change as the root of wage inequality are invariably working with a model that assumes labor markets are competitive. In these models, workers and employers are equally powerless, and wages and employment are set by the intersection of demand and supply curves in competitive markets for labor, with very little scope for the economy to diverge from these competitively determined levels without adverse consequences. Crucially, this means that only one employment level is consistent with a given wage level and vice-versa—wages and employment are jointly determined by the same underlying forces, and this means that any influence that affects one of these necessarily affects the other. Given this model of the labor market, it is natural to react to large changes in wages or employment for any group of workers by postulating that something must have shifted either relative labor demand or supply.

"Relative" labor demand or supply means demand or supply of one type of labor relative to other types of labor. So, for example, if employers decided that college-educated workers were growing more productive and valuable over time (say because they had more facility with new forms of technology), relative labor demand would increase for workers with college credentials, while relative labor demand would decrease for those without these credentials. The result would be both wage and employment levels rising for college workers and falling for noncollege workers.

Much of the economic research making strong claims that the rise in wage inequality over recent decades is driven by technological change relies on competitive models of the labor market. It is important to realize the strong role that this *assumption* of a competitive labor market plays in this research. Real-world trends in relative labor supply are easy to observe in data on the size of the workforces with and without college degrees. The relative wage can also be seen in the data—it's the ratio of average wages for workers with a four-year college degree to the average wages of other workers. However, these two observable datapoints are often combined with the *assumption* of competitive labor markets to infer trends in relative demand for different types of labor. Often these inferences of trends in labor demand are incredibly influential in public debate. For example, the claim that the introduction of personal computers drove inequality in the 1980s and 1990s is often a direct statement about the inference that technology shifted the relative demand for workers without a college degree. Yet direct evidence of economic influences that reliably shift demand or the timing of when they might have happened is extremely thin.[12]

Those who emphasize the importance of institutional change for wage inequality are nearly always working with a labor market model that includes substantial employer-side market power. The source of this power may vary. It can include traditional monopsony power stemming from too few employers, dynamic monopsony power stemming from informational and logistic frictions associated with job changing and search, employer choice in how effort is elicited from workers (through costly monitoring or higher wages), or some other source.

Frictions and unbalanced power in labor markets mean that a range of influences besides workers' own productivity affect wage levels and their evolution over time. Manning (2003) has argued that frictions in real-world labor markets make changing jobs costly to workers, and hence effectively grant employers substantial "monopsony power" over their employees. Some of these frictions that make job changes more costly include things like researching and applying for alternative jobs, changing commuting schedules, rejiggering child care arrangements, switching health insurance plans, and breaking social ties with work colleagues.[13] No single one of these frictions imposes costs that are high enough to prevent *any* job switching from happening, but the accumulated drag of some or all of them can substantially blunt the potential of labor market competition to boost workers' wages. Further, even quite small reductions in competition spurred by these frictions can lower wages significantly.

Of course, a literal labor market monopsony would be one in which only one single employer existed, which would obviously keep competitive pressures from working to help workers bargain for higher wages. The Manning (2003) model does not require just one employer or even some arbitrarily small number of employers; it only requires that some employers are able to exploit the real-world fact that the costs of switching jobs for workers is nonzero. If this cost of job switching is a part of the baseline model of labor markets, it grants employers considerable power.

Besides this baseline reality of costly job changing, other forms of employer power stem from realities of the production process within firms in capitalist economies.

For example, Bowles (1985) points out that employers must hire workers to produce output, but must also elicit effort from these workers. Employers' main leverage to elicit effort is the threat to fire workers found to be shirking. Firms have two main instruments to maximize leverage from this threat: they can monitor workers intensely—so that any shirking is highly likely to be detected—or they can pay high wages to intensify the pain of losing a job. Either higher monitoring or higher wages can incentivize workers to expend more effort and shirk less. Both strategies are costly to employers: to implement the monitoring strategy, they must hire managers who do not contribute directly to production, but instead just oversee workers' effort, whereas to implement the high-wage strategy, they must increase the pay of workers directly involved in production. In some cases, if the "outside" wage available to workers is one generated in a labor market characterized by substantial monopsony power, a high-wage strategy adopted by a firm to elicit effort can counterbalance the depressing wage effects of this monopsony power.

Regardless of the source of market power, recent cutting-edge research has demonstrated how far from the competitive ideal most labor markets truly are. The key effect of this employer-side power is to make the range of possible wage-employment level combinations set in the labor market much wider than is possible in competitive models. A given employment level can be consistent with a wide range of wage levels. This "range of indeterminacy" can explain, for example, why large increases in mandated minimum wages are often found in empirical studies to have no significant effect on employment levels.[14] This noneffect of minimum wages on employment, conversely, would be hard to explain with competitive models where a single combination of wage and employment level is determined jointly by the intersection of demand and supply curves. Hence, the potential role for institutional change to significantly drive inequality—even absent any change to underlying demand and supply for labor—is much larger in models of labor markets with employer-side power.

For decades, the assumption that labor markets are best represented by competitive models was widely adopted across the economics profession, and this naturally channeled much research about rising inequality into searches for "demand-shifters" like technology. More recently, models with employer-side power have become much more prominent in labor market debates, and the possible scope for institutional change to drive trends in wages and inequality has been more widely recognized.[15] Now, debates over the drivers of wage inequality in recent decades require a much higher empirical burden of proof than they did in the past, when the assumption of competitive labor markets lead almost inevitably to the conclusion that technology played a key role.

To put our cards on the table, we believe the evidence strongly supports a view of labor markets where employer power is significant, and that direct evidence of technological change having first-order effects in changing relative demand for labor is extremely thin (we highlight some of this evidence and its thinness in a later section). But putting this debate front and center when discussing the potential effects of technological change on wages and employment is a useful practice going forward regardless.

How economists typically measure technological progress: productivity growth

Economists generally measure technological progress as an increase in economywide *productivity*. There are two main ways that productivity and productivity growth are measured. First, *labor productivity* is the amount of income generated in an average hour of work in the U.S. economy. This income includes wages, but also business income (including corporate profits), rents accruing to landlords, and other forms of income. Second, *total factor productivity* (TFP, sometimes also called multifactor productivity) growth measures how much output has grown *after accounting for the growth of all measurable inputs*, such as labor and contributions from capital services (like factories and machines). Economists often focus more on TFP as a measure of pure technological change. However, in the rest of this paper, we will focus more on labor productivity and argue that it maps more directly onto popular conceptions of how technology might influence economic outcomes.

Labor productivity—or the income generated in the average hour of work in the U.S. economy—rises consistently over time. These increases are why the current generation is on average so much richer than their ancestors—the average hour of work in the economy of 2023 generated far more income than the average hour worked in (say) 1960. Labor productivity has grown steadily—if inconsistently—for the last century or more in the United States.

The main drivers of growth in labor productivity are *labor quality*, *capital-deepening*, and *TFP growth*.

Labor quality increases over time reflect the growing average level of educational attainment in the economy—more highly educated workers tend to be more productive workers, and increasing educational attainment is one reason why an average hour of work in 2023 generated more income than an average hour of work did in 1960.

Capital-deepening reflects the fact that workers in 2023 had access to much better tools with which to do their jobs than workers had in the past. An obvious example is digital scanners at retail establishments, which allow faster and more accurate pricing at checkouts. Another example is word processing (particularly editing and redrafting) that can be done much more efficiently with personal computers than with manual typewriters. Both examples—digital checkout scanners and the replacement of typewriters with personal computers—illustrate why the measure of labor productivity is likely more aligned with what people think of when they envision technological progress and its effect on the economy, as neither of these effects would be reflected in looking solely at trends in TFP growth.

Total factor productivity reflects the fact that a given set of inputs (a particular number and type of workers, and a particular bundle and type of capital goods) produced more output in 2023 than it did in years past. Because it accounts for tangible inputs (hours worked and capital used), TFP growth is sometimes referred to as the influence of "ideas," or as the purest form of "technological progress." Many economics papers refer exclusively to TFP when they purport to measure trends in technological progress.

This paper focuses more on trends in labor productivity, because we think most people understand the broader determinants of growth in labor productivity as being reflective of technological change.[16] For example, most people would see the introduction of digital scanners and computers as a key way that technological progress has changed how people perform work in recent decades. In contrast, many people would find it odd or too limiting to hold constant the effect of computers when assessing the influence of technological change on the labor market.

Figure A highlights trends in labor productivity growth and the contribution of its drivers over U.S. business cycles since World War II. The most striking finding from this analysis is that productivity growth over the most recent business cycles has been historically *slow*, not fast. This alone is provides key context for current debates about how the economy can absorb technological progress: any technology-induced job destruction allowing a given hour of work to produce more income—and hence substitute more sharply for labor—has substantially *slowed* in recent decades. Yet breathless reporting on today's technological advances often ignores this, or even outright claims the opposite.

Figure A

Last two decades have seen historically slow productivity growth

Average annualized change in each component's contribution to productivity growth and their total, by business cycle

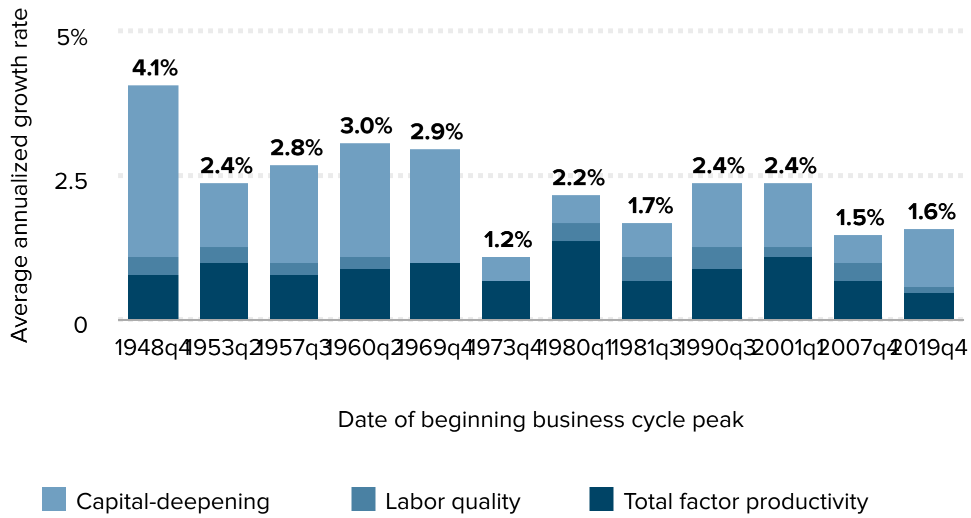

Source: Fernald (2023) data from the Federal Reserve Bank of San Francisco.

Economic Policy Institute

By far the biggest slowdown in the contributors to labor productivity growth has been in the category of TFP growth—or the "pure" form of technological change. The first implication of these trends is obvious: if rapid technological progress is feared to cause labor market problems, were these labor market problems more pronounced in past business cycles, when this technological progress ran faster? The next sections address this question.

Can accelerating technological progress cause mass joblessness?

If we define technological progress as the ability to produce more output in a given hour of work, this often raises an obvious concern: Won't less labor be needed over time, causing mass joblessness?

The answer is a clear "no." While it is true that the level of unemployment at any given point in time is in part a function of the economy's productivity, there is another variable—*aggregate demand*—that policymakers have significantly more control over and which can be adjusted to keep unemployment low, regardless of productivity trends.

Unemployment rises when the economy's *potential output* exceeds *aggregate demand*. Potential output is a measure of how much an economy could produce if nearly all the

Figure B

When GDP falls below potential GDP, unemployment rises

Actual and potential GDP since 1979 ($2017)

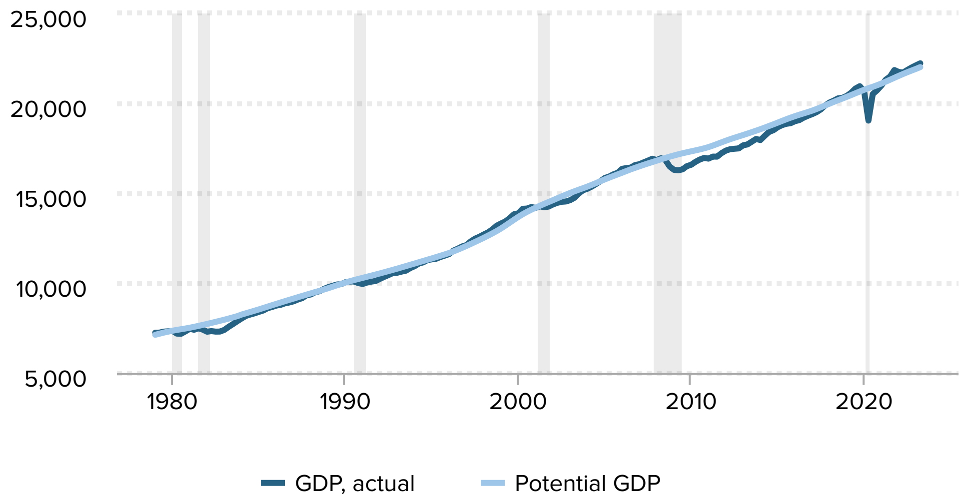

Note: Recessions are shaded in grey.

Source: GDP data from the National Income and Product Accounts (NIPA) of the Bureau of Economic Analysis (BEA). Estimates of potential GDP from the Congressional Budget Office.

Economic Policy Institute

economy's willing workers were fully employed.[17] A key determinant of potential output is productivity—any given employed workforce can produce more if productivity is higher. Aggregate demand is the amount of spending by households, businesses, and governments. When aggregate demand falls beneath the economy's potential output, then unemployment rises. Say that there is a hotel with staff and rooms for 50 parties. If only 45 parties offer to rent these rooms, then five rooms and the workers to staff them will be unneeded. If this deficiency of demand is widespread across most sectors of the entire economy, then unemployment will rise as unneeded workers are laid off and not reemployed in other sectors.

Figure B shows estimates of the economy's potential output, as well as actual measures of gross domestic product (GDP)—the value of the nation's output and income in a given period. When actual GDP falls beneath potential, this means that aggregate demand is running more slowly than growth in the economy's supply side, resulting in rising unemployment. Recessions are indicated by grey shading in the figure and they are defined by actual GDP falling beneath potential output.

This logic might at first glance seem to buttress fears about technological progress generating unemployment: technological progress boosts the economy's potential output, and if this boost pushes it above the economy's aggregate demand, then unemployment can result. But the data show clearly that sharp changes in potential output (which is how technology-driven productivity jumps would show up in this data) is not behind the mismatches in aggregate demand and potential output that lead to recessions.

The determinants of potential output move slowly. The size of the labor force and the nation's capital stock (and its state of technological sophistication) do not whipsaw around year to year. Instead, they tend to grow at a slow and predictable rate over time. Aggregate demand is far more volatile and *can* whipsaw quickly from year to year. For example, when the bubble in home prices began deflating in late 2006 and 2007, households immediately began spending less money and saving more to make up for the lost value of wealth, leading quickly to the severe 2008–2009 recession.[18] Similarly, in early 2020, the labor force available to firms in the face-to-face services sector did not disappear and cause an employment collapse. It was customers who disappeared as fears of COVID-19 spread, and it was this demand shock that led to mass layoffs in the early part of that year.[19]

But just as aggregate demand can fall quickly, policy efforts can boost it quickly to ensure recessions are short-lived and recoveries are rapid. Aggregate demand can be boosted through either monetary or fiscal policy interventions to boost demand, with the Federal Reserve using monetary policy tools (like interest rate cuts), and Congress and the president setting taxes and spending at the levels needed (in practice, fiscal policy turns out to be the more powerful tool). Support for the statement that policy can quickly restore falls in aggregate demand is provided by the U.S. economic recovery from the COVID-19 recession. In December 2020, after the low-hanging jobs created by simply reopening the economy after the first wave of the pandemic had been restored, the unemployment rate was 6.7% and job growth had turned negative. Absent further policy efforts, there was a real possibility of stagnation at this high rate of unemployment. But due to further fiscal recovery packages passed in December 2020 and March 2021, by the end of 2021, the unemployment rate had already fallen below 4% again—essentially matching the immediate pre-pandemic level.[20]

Table 1 highlights this point about which variable—aggregate demand or potential output—moves more quickly (and in the right direction) to cause periods of joblessness. It shows growth rates for both actual and potential GDP in the year before recessions have hit the U.S. economy, and then over the subsequent recession. It then calculates the "swing" in these growth rates—how much they changed as the economy entered recession. Crucially, any sharp divergence of real GDP from potential output is caused by changes in aggregate demand.

In all cases, real GDP growth has swung sharply from positive to negative in the first year of recessions, by an average of 4.6% in the five business cycles before the COVID-19 pandemic (the COVID-19 recession was so extreme that we will set it aside for now). Estimates of potential output slowed as well, but only by an average of 0.4% over these same business cycles. Further, *slowing* potential output growth can *reduce* unemployment if it represents a slowdown in productivity growth, so this slowdown in estimated potential output puts downward—not upward—pressure on joblessness. In short, the wrenching change that causes recessions and rising unemployment is *not* an acceleration of technological progress making labor unnecessary—again, potential output *decelerated* in each of these periods. Instead, the pronounced change is the rapid deceleration and

Table 1

When GDP changes, it's because aggregate demand falls

Changes in actual and potential GDP as recessions hit

	1980q1	1981q3	1990q3	2001q1	2007q4	2019q4
Last year before recession						
Real GDP	1.4%	4.3%	1.7%	2.2%	2.1%	3.2%
Potential GDP	3.2%	2.5%	3.0%	3.7%	2.0%	1.9%
Peak-to-trough change (annualized)						
Real GDP	-4.3%	-2.0%	-2.7%	0.7%	-2.6%	-17.5%
Potential GDP	2.3%	3.0%	2.5%	3.0%	1.7%	1.8%
"Swing"						
Real GDP	-5.7%	-6.3%	-4.5%	-1.5%	-4.7%	-20.6%
Potential GDP	-0.9%	0.5%	-0.5%	-0.8%	-0.2%	-0.1%

Note: Author's analysis of data sources from Figure B.

Economic Policy Institute

outright *fall* of real GDP, which, given trends in potential GDP, must by definition have been caused by a fall in aggregate demand.

Figures C and **D** provide some slightly more systematic looks at the relationship between productivity growth and joblessness. In both figures, average values over an entire business cycle peak—from one peak to the next—are assessed. The dates on the dots in the figure mark the beginning of the business cycle. Figure C shows the average rate of productivity growth and the average rate of unemployment across business cycles since World War II. Contrary to worries about tradeoffs between fast productivity growth and low unemployment, fast productivity growth is associated with *lower* average rates of unemployment across business cycles.

Figure D shows the relationship between average productivity growth and the *change* in unemployment rates between business cycle peaks. That is, it looks to answer the question: On average, fast productivity growth may be associated with lower unemployment, but does fast productivity growth over a business cycle keep unemployment from falling as fast as it could have? Again, there is no systematic relationship between the average pace of productivity growth and the decline of unemployment over an entire business cycle.

Over the last completed business cycle (from 2007–2019), productivity growth averaged roughly 1.5%. The most highly optimistic projections for how much AI can boost the pace of productivity growth are about 1% per year (most other projections are quite a bit lower). This would move productivity growth from 1.5 to 2.5%—a level that the U.S. economy saw for decades following World War II, and which was accompanied by lower average

Figure C

Fast productivity growth and low unemployment do not trade off against each other

Average unemployment and average annualized productivity growth across business cycles

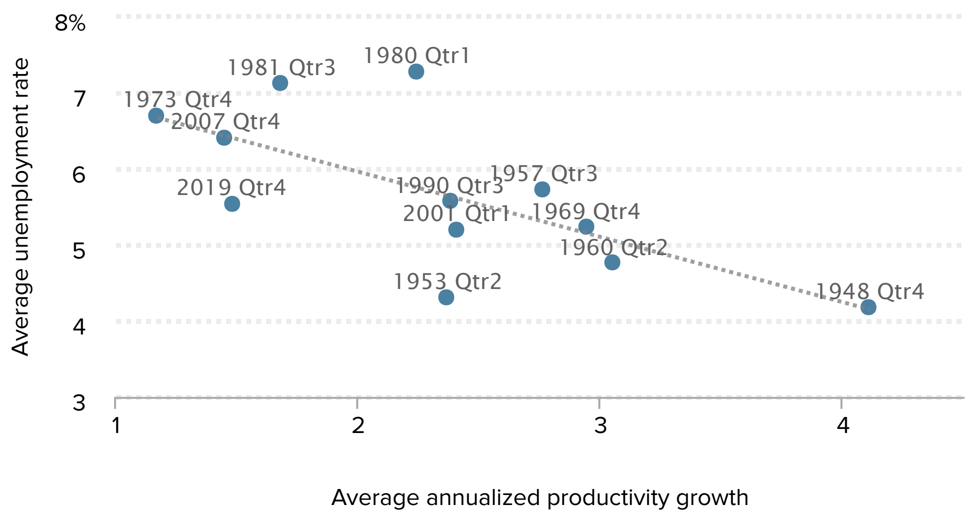

Source: Unemployment data from the Bureau of Labor Statistics (BLS) Current Population Survey, productivity data from the BLS productivity program. The precise productivity measure used is real output per hour worked in the nonfarm business sector.

Economic Policy Institute

unemployment than has persisted in recent decades.[21] In short, there is nothing in either the historical relationship between productivity growth and unemployment or in projections of AI's impact on productivity growth that indicate that this technological change will prevent policymakers from sustaining low rates of unemployment—should they choose to do so.

Does technological change ever displace jobs?

None of this is to say that *specific jobs* are not threatened by technological progress. Rapid technological change concentrated in any specific sector can reduce employment *in those sectors*. The analysis above simply says that the aggregate number of jobs and the overall rate of unemployment is unlikely to be threatened by an acceleration of technological progress, as long as policymakers respond appropriately by boosting aggregate demand.

As productivity rises following an acceleration of technological progress, job losses within sectors experiencing the productivity increase will be counterbalanced (or more than counterbalanced) by expanding employment in other sectors *as long as aggregate demand is maintained*. Autor and Salomons (2018) empirically estimate how employment responds to a sectoral productivity shock. They find that the *own-effect* of a productivity

Figure D

Fast productivity growth and lower unemployment do not trade off against each other

Peak-to-peak unemployment change and average productivity across business cycles

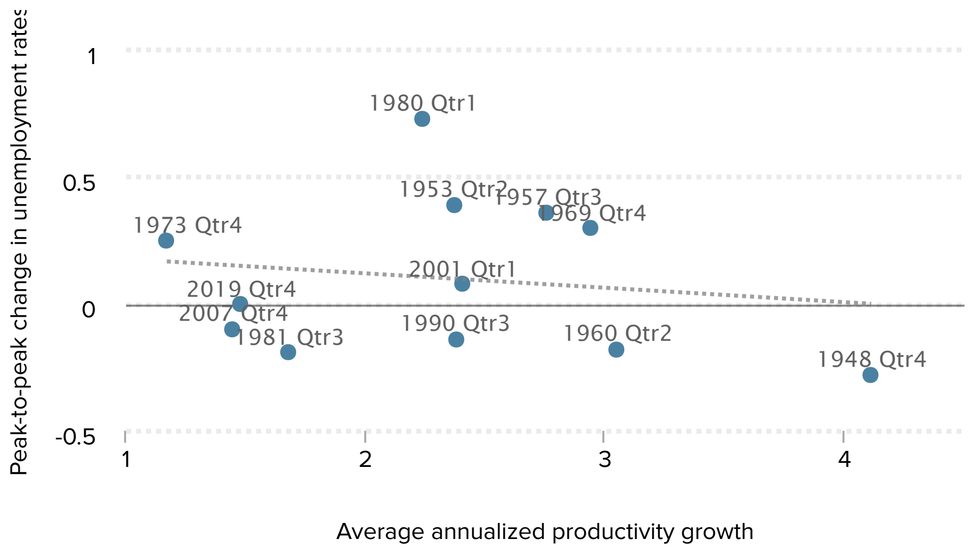

Source: Unemployment data from the Bureau of Labor Statistics (BLS) Current Population Survey, productivity data from the BLS productivity program. The precise productivity measure used is real output per hour worked in the nonfarm business sector.

Economic Policy Institute

shock within a sector is indeed modestly negative, with the reduction in hours of work needed to produce output in the sector not fully offset by the rise in demand for the sector's output, made cheaper by productivity growth. However, the *cross-effect* of productivity growth within a sector—the effect of its own productivity growth on employment *in other sectors*—is strongly positive, and outweighs the negative own-effect in terms of aggregate employment trends.

Take the example of a 1% increase in productivity in a specific economic sector like manufacturing. Autor and Salomons (2018) find that the average first-order effect of a sectoral shock (the own-effect) is to decrease employment in that sector by 0.1%. This is the intuitive effect most people think about when they worry that introducing more automation into production might displace human labor *in that sector*.[22]

But this 1% rise in sectoral productivity means that more income is being generated in each hour of work in that sector, and this extra income boosts employment when it is spent in other sectors. This positive "final demand effect" on jobs alone almost completely counterbalances the direct effect, adding almost 0.1% to employment. Additionally, the combination of productivity growth and competition in product markets lowers the prices of goods from the sector that has seen the positive productivity shock. In the example of manufacturing, this would provide a boost to employment in sectors that use manufactured goods as intermediate inputs (for example, a falling price of computers makes it less expensive to produce accounting services). These “upstream effects” boost employment by almost twice as much as the direct effects reduce it. Overall, the economywide net impact of these effects is an *increase* in overall employment stemming from productivity growth within a given sector.

How to reduce damage from sectoral job displacements in labor markets—whatever its cause

It is certainly true that some individual workers may suffer from sector-specific job displacements, even if aggregate unemployment or employment is unaffected. The labor market is not frictionless, and it may take some painful time before comparable employment in a new sector is obtained. Some workers (particularly older workers) may never find a specific job as good as the one they lost. Yet much of this individual suffering could be ameliorated with broad policies that provide better protective social insurance, more widespread collective bargaining, and sustained high-pressure labor markets—policies that are highly desirable regardless of the pace of technological change.

One reason specific jobs are occasionally highly valued in the U.S. labor market is because they come bundled with nonwage compensation—like health and retirement benefits, which are not universally available. But if these benefits were universally available through more protective social insurance systems, the damage done by the loss of any particular job would be greatly lessened. Another key social insurance system—unemployment insurance (UI)—is too stingy in the U.S., causing large income losses while workers search for alternative employment. Boosting the protectiveness of UI would be a key win for those looking to reduce the pain caused by the loss of particular jobs.

Another reason some specific jobs can be highly valued in the U.S. economy is because they are unionized. This should not be as rare as it currently is, but recent decades have seen a combination of employer hostility and policy indifference lead to a near shutdown of organizing unions in newly created jobs at any large scale. This means that sectors today that remain unionized do so largely because of a historical legacy that saw their unions formed decades ago; the chances of workers leaving this sector finding a unionized job elsewhere are slim indeed. In short, there are only rare pockets of unionized jobs in the U.S. economy and new ones are not being created fast enough. Hence, anything (including technological progress) which leads to the destruction of today’s unionized jobs are likely to leave many of their former holders worse off.

Additionally, the U.S. economy has spent much of the past four decades with excessively high unemployment, which radically increases the cost of losing a job. When unemployment is low and vacancies are high, a worker who has lost their job can quickly find alternative employment, putting employers under constant pressure to keep job quality high enough to retain and attract new employees.

If U.S. policymakers created a more protective social insurance system, restored the effective right to organize unions, and maintained high-pressure labor markets with low unemployment, then a large part of the damage done by technologically induced job displacements would disappear.

Finally, despite all the possible challenges faced by workers who are displaced from specific sectors by technological change (or anything else), it is possible to both overstate how widespread these challenges are, and underestimate the value of new jobs and the higher productivity created by technological change.

How widespread is sector-specific "churn" caused by technology and has it increased?

Were technology responsible for a reallocation of jobs toward certain industrial sectors or occupations, we should expect to see an increased amount of employment flows with workers increasingly separating from jobs, and certain sectors losing and gaining shares of employment in the labor market. The U.S. labor market has always and everywhere been characterized by tremendous rates of job "churn"—workers separating from employers either voluntarily or involuntarily. For example, in the last year before the COVID-19 pandemic, 3.7% of all workers left their jobs *each month*. Similarly, 3.9% of all workers were newly hired each month (for a net change in employment each month of roughly 0.2%). Over a year, this is a huge amount of churn, with more than a third of the entire workforce changing jobs (or changing their employment status) each year. Yet this churn has been a feature of the U.S. labor market for decades, and most data indicates that it has actually slowed, not increased, in the 2000s—despite the proliferation of the Internet and large advances in computer hardware and software.

Figure E, reproduced from Bivens and Mishel (2017), clearly emphasizes this point. It shows the sum of the (absolute) change in occupational employment shares over various decades. To construct this metric, Bivens and Mishel examined the shares of total employment for 250 occupations at the beginning and end years of each decade and computed the changes in these shares. The metric shown in the figure is half of the sum of the absolute value of changes in occupational employment shares (taking only half of the sum avoids double counting gains and losses). This metric measures the share of total employment exchanged between occupations—or the measure of job churn between occupations—for each decade.

The decadal rates of occupational employment shifts, starting in the 1940s, are shown in Figure E. The rate of change was fairly uniform over the 1940–1980 period, and far more rapid than for any period since 1980. The period since 2000 has seen the lowest rate of change—half the rate of change of the 1940–1980 period.[23]

Were technology causing massive displacements or reallocation, the data would have exhibited the opposite pattern. One important reason for the lack of widespread displacements is that technological increases can complement the tasks of workers, rather than permanently substitute away from particular occupations or industries. As a result,

Figure E **Change in occupational employment shares, by decade, 1940–2015**

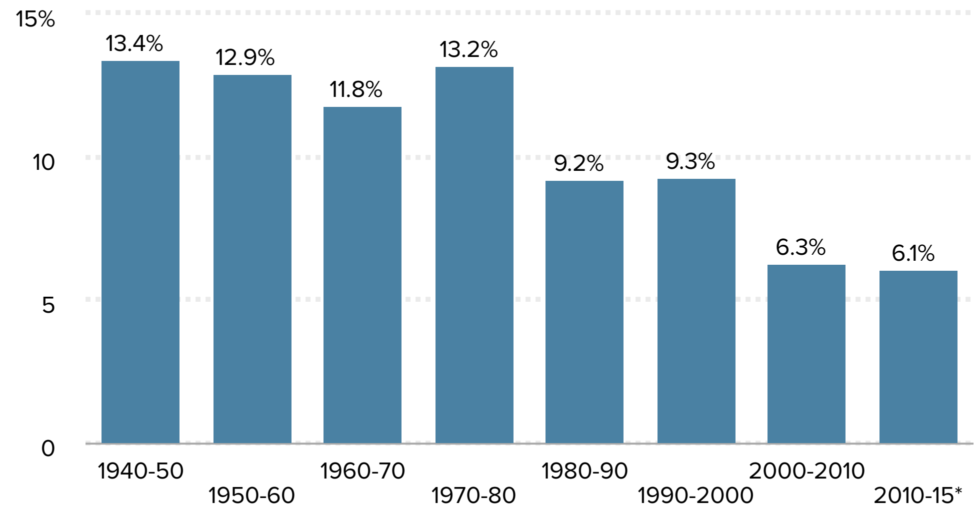

* Converted to decade rate by multiplying by two.

Source: Authors' analysis of data from Atkinson and Wu (2017)

Economic Policy Institute

even though large shares of the labor market may be exposed to new technologies, much smaller shares of jobs would be destroyed entirely by automation. Indeed, some observers have in fact argued that AI provides "an opportunity to complement worker skill and expertise" (Acemoglu, Autor, and Johnson 2023).

Does technology reduce demand for workers without the right credentials or skills?

We argued in the previous section that technological change and increased productivity has not led to aggregate job loss or increased unemployment. Moreover, even in recent years, these forces have not even led to more rapid occupational churn in the labor market. However, many economists have argued that technological change was a major cause of growing wage inequality in the U.S. labor market in the post-1979 period, and that this technology-induced rise in inequality was the result of technological changes that boosted relative demand for workers with higher skills (almost always proxied by a four-year college degree). This technology narrative has been extraordinarily influential among policymakers, even as cutting-edge research increasingly casts doubt on it.

This shift in relative demand toward college workers, sometimes called *skill-biased technological change* (SBTC), has been a major focus of economic research in understanding the growth of U.S. wage inequality. The SBTC-based explanation of inequality relies on a model of competitive labor markets, where wages and employment of workers of different skill levels have their relative wages and employment levels set by the intersection of supply and demand. The SBTC theory claims that technological change has caused an increase in relative employer demand for college workers (presumably because these allegedly more skilled workers have greater facility with using new forms of technology), and this rise in turn led to higher relative wages (or a higher *college wage premium*) over the last several decades.

This stylized story simply does not fit the data. First, basic estimates of the relative demand for college labor suggest that the bulk of the growth in the college wage premium in the 1980s and 1990s is not due to an acceleration in employer demand for college labor, but a slowdown in the supply of college labor (therefore raising the price or wage of college labor). As Autor, Goldin, and Katz (2020) explain, “rapid and disruptive technological change from computerization, robots and artificial intelligence is not to be found” during these periods of massive innovation in computing technology. These authors (and others) often present this set of facts as demonstrating that inequality is the result of a “race between technology and education,” with technology presumed to raise relative demand for college graduates, while education conditions the supply. However, recent decades have clearly seen much more marked changes in the education/supply side of this race—and that leads to a narrative about the driver of inequality that departs significantly from stories that center technological change as the driving force.

Second, compared with earlier time periods, there has been little change in wage inequality between college and noncollege workers since 2000. **Figure F** shows the annual college wage premium over 1979–2023, controlling for demographic differences in the college versus noncollege population within each year. There was a sharp increase in the college wage premium in the 1980s and 1990s, but a much smaller rise since 2000, during the widespread adoption of computing at the workplace. In fact, there has been essentially zero change in college/noncollege wage inequality since 2010, so if anything, these wage patterns suggest a decline in the relative demand for college labor over the last one to two decades.

In a preview of recent concerns over AI, the early recovery from the COVID-19 recession saw many expressing worries that employers would respond to the organizational changes they made in the era of social distancing to replace workers with technology. Casselman (2021), for example, wrote that:

> An increase in automation, especially in service industries, may prove to be an economic legacy of the pandemic. Businesses from factories to fast-food outlets to hotels turned to technology last year to keep operations running amid social distancing requirements and contagion fears… But some economists say the latest wave of automation could eliminate jobs and erode bargaining power, particularly for the lowest-paid workers, in a lasting way.

As support, Casselman (2021) pointed to a 2021 working paper from the International Monetary Fund that argued: “Our results suggest that the concerns about the rise of the

Figure F

The college wage premium has stagnated in recent years

The log wage difference between workers with and without a college degree

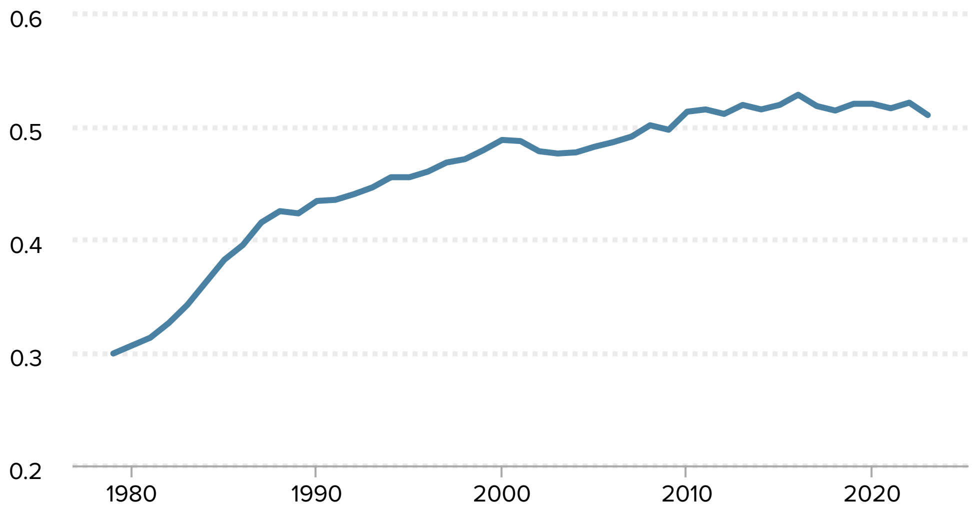

Notes: The college wage premium in Figure F is estimated from year-specific sample-weighted regressions of Version 1.0.47 of the EPI Current Population Survey extracts of the log hourly wage on college degree attainment, a quartic polynomial in age, and gender, race, marital status, and state fixed effects.

Economic Policy Institute

robots amid the COVID-19 pandemic seem justified” (Sedik 2021).

And yet, almost three years on, the post-pandemic labor market has actually been a huge source of strength for low-wage and low-credential workers. Autor, Dube, and McGrew (2023) show that after accounting for changes in the demographic composition of the workforce, the college/high school wage premium fell during the last two years. Instead of technological change widening the gap between those with more or fewer credentials, a tighter labor market during the 2021–2023 period compressed wages. Young, noncollege workers saw significant wage increases because the tighter labor market provided more opportunities to switch to higher paying jobs.

Employment rates for workers without a college degree are still worse than they were decades ago, but in aggregate, they are largely not determined by technological changes. An easy way to see this is comparing the United States to other advanced economy countries who have faced similar technological shocks but who have very different macroeconomic policy and social support systems. **Figure G** shows that in the United States, the share of the population with a high school but no college degree that is employed has dropped dramatically since 2000. In contrast, low-credentialed employment in other G7 countries has not experienced such falls. In some cases, like in Austria, Germany, and Great Britain, employment rates have grown for workers with just a high school education.

Figure G

Employment rates for ages 25-64 with high school but no college degree, by G7 country

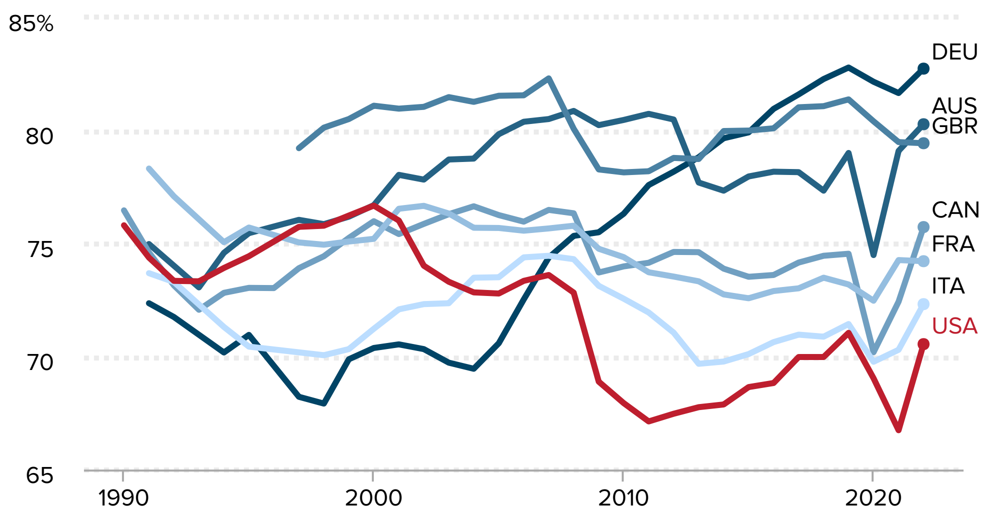

Note: Employment rates are for those with an upper secondary but nontertiary education level.

Source: OECD (2024).

Economic Policy Institute

Technology is a tool—the balance of labor market power determines who it helps

While most economic analysis of technology's potential role in generating greater inequality in recent decades has focused on its effect in shifting demand and supply curves in competitive labor markets, this is often not how most informed consumers of news about the U.S. economy think about the effect of technology.

Instead, many media reports about technology's role in the workplace—and how it might exacerbate inequality—focus on how it might be used as a tool for employer monitoring and speed up of workflow. For example, a well-known example of technology-enhanced monitoring is the "black box" installed in long-haul cargo trucks. Originally, these boxes were installed to validate that long-haul truckers were meeting mandated rest requirements for safety reasons. Now, however, one of the main appeals of the box for employers is to ensure that they only pay truckers for the time they spend actually moving cargo forward. Skott and Guy (2013) note that the producer of one of these black boxes boasts on their website that for trucking company managers, this technology "is like being able to sit next to every one of your drivers every second they drive."

Another example is robots in an Amazon warehouse essentially setting the pace for human workers when processing packages. As Anway (2022) writes: "The clock was always ticking…As soon as she'd filled a rack, she'd press a button and one robot would zip it away while another robot would bring a new one to fill." This high pace of work has been implicated in high rates of injury in these warehouses.

These examples have nothing to do with technology's effect in shifting demand or supply curves for labor in competitive markets, and yet show technology as enabling exploitation and degrading job quality. A much smaller body of economic research highlights labor market models that shed light more directly on these situations.

Earlier, we described models where employers had market power and could choose between monitoring or high wages as strategies to elicit effort from workers. One could imagine technological change that reduces the cost of monitoring. This could induce firms to lean more heavily on the intensive monitoring strategy and less on the high-wage strategy. This in turn would lead to more workers having their wages set directly by the labor market outside the firm—an outside labor market which itself might be riven with employer-side power—and hence lower wages. This heavier reliance on monitoring would *lower* measured productivity, as more work hours in the firm would be spent monitoring other workers rather than producing output for sale. In this case, technological change would not be boosting productivity (it would in fact be lowering it) and would instead only be leading to a zero-sum (or even negative-sum) redistribution away from workers and toward employers and managers.

However, technological advance is not the only—and likely not the primary—determinant of whether firms choose a high-wage or a high-monitoring strategy. A bigger determinant is the *relative bargaining power of workers*. If the workers at a given firm manage to organize a union, for example, the choice is essentially decided: wages will be higher, and workers will value the unionized job more than what they can get in the outside labor market (and hence will expend more effort).

The example of the trucking black box monitoring technology makes this clear. Originally the box was thought necessary to keep truckers from breaking safety rules regarding how much rest they got between spells of driving. In more recent incarnations, it is marketed as a device to ensure trucking companies do not have to pay for any time spent that does not move cargo forward. However, what this example makes clear is that it is the underlying power relationship, not the new technology, that determines wages and job quality. For example, if workers were paid sufficiently for the entire time commitment of hauling cargo (and not just for time actively driving), the worry that they would skimp on rest requirements to earn more money would be blunted. Further, while the black box is often referenced as a tool for employer control, the underlying technology could in theory solve a pressing problem for truckers: proving that large swathes of time they're not actively hauling cargo is in fact necessary "company time," as they are forced to wait to pick up loads at ports.

In the supply chain breakdown of 2021 and 2022, a key bottleneck to moving goods from producers to consumers was a backlog at ports. This backlog led to truckers often having to wait long hours (or even days) idling in a queue waiting to have their cargo loaded. And often, this wait time went uncompensated. One employer objection to paying for this wait time could be verification—the company only “knows” when the trucker is really working for them when the load is transferred to the truck. But that’s obviously not true—the same technology that verifies whether long-haul truckers are spending enough time actively driving to meet their contractual demands could also verify that truckers are indeed in an active queue waiting for cargo to be loaded.

The real reason why truckers have not been compensated for these wait times in recent decades is not the technical impossibility of verifying wait times, but instead simply the power of employers. Trucking was once a highly unionized, high-wage job. The push to deregulate and deunionize beginning in the 1970s substantially eroded the relative wages in this sector. One imagines that if the black box had been invented in the 1950s and employers tried to force its adoption by a more heavily unionized trucking workforce, it would have been successfully rejected by the then considerably more powerful Teamsters union.

In short, it is true that technology exists that might aid employers engage in zero-sum redistributions away from typical workers. But the exact same technology used by employers to wring more effort and profit out of workers could often in theory be used by workers to wring higher wages and productivity out of employers. The same robots that are implicated in a work pace that is injurious to workers at Amazon could be a genuine boon to worker safety if robots handled all heavy loads and *did so at a pace that did not put undue stress on human workers*. This pace is not dictated by technology, it is set by employers, too often in the context of highly unbalanced power. In short, the problem is almost never in the technology itself, and nearly always in the relative power relationships.

Recent flashpoints about hiring discrimination and IRS audits highlight that unbalanced power is the root problem

Discrimination in hiring processes and in federal tax enforcement are two key examples highlighting that it is power—of bosses and policymakers—that determines whether or not technology (including variants of AI) are used to ameliorate or exacerbate existing inequalities in U.S. society.

It is known that automated processes for employer hiring can use embedded discriminatory criteria when sorting applicants.[24] This is obviously a real problem. Yet discriminatory criteria do not appear by magic in automated data processing systems; their logic is explicitly or implicitly programmed somewhere along the way. The best response to this issue has very little to do with the automated process itself: it is making firms legally responsible for the outcomes of their own hiring software’s decisions and providing regulators enough access and information to perform audits and measure the magnitude of bias in the hiring process. There are even reasons to believe that discriminatory criteria embedded in automated hiring systems will be easier to detect and solve than old-fashioned employment discrimination that largely happened inside the heads of hiring managers.[25] Again, the underlying problem here is not technology, it’s the broader social context this technology operates in, which could in fact benefit from the use of technology in combatting some of its problems.

A final example highlights a potential danger of focusing on technology as the problem rather than the more foundational decisions embedded in technology. A recent paper looked at IRS audit rates for Black and non-Black taxpayers. They found Black taxpayers audited at substantially higher rates. They found this disparity (and other key features of IRS audits) could possibly be explained if the IRS was picking taxpayers for audits based on an algorithm that sought to maximize the share of underreported income that was accounted for by refundable tax credits (like the Earned Income Tax Credit).[26]

This is an odd target to maximize if you thought the point of audits should be to simply generate as much appropriate revenue as possible. Potentially, however, it is an understandable thing to maximize if you are an agency that has been swayed by unrelenting Republican attacks on refundable tax credits in the name of minimizing “fraud” perpetrated by low-income taxpayers. The authors also find that the audits fall heavily on returns with zero business income. This is again odd if you want to maximize unclaimed revenue, as business income is rife with underpayment. But this choice might make sense for an agency that has been starved of resources—business income returns require a lot more resources to audit than individual returns.

Both plausible maximization goals (focusing on refundable tax credits and not focusing on returns with business income) cut sharply toward increasing the share of Black households that would be selected for audits. They suggest (implicitly) a much better maximization goal to guide the audit selection algorithm: maximize underreported income, period. This goal would not only raise more revenue (the larger point of audits), but would also erase the race-based disparity in current audits.

Again, the issue is not algorithms or automation per se, it is the human choices behind them. More broadly, is there any question that advances in information processing (AI or otherwise) could be a hugely helpful tool for using IRS audits to maximize revenue if that is what the agency wanted? To put it simply, banning or severely constraining the use of AI or any other information-processing tool in the conduct of tax enforcement would be a huge win for wealthy tax cheats looking to escape taxation. It may be fanciful to worry about such bans, but given that Republicans in Congress have routinely sought to hamstring tax enforcement for decades, if fears and generalized bad feeling about AI becomes widespread across society, it may provide an opening for such destructive proposals.

Conclusion and policy implications—looking through the latest technological fad to see the real threats to workers' well-being

It is no doubt a useful exercise to make sure public policies are tailored to specific forms of significant new technology that arise. So, a recent spate of proposed legislative and regulatory activity around AI has many sensible elements. But it is also extremely easy to focus too much on the latest form of technology and get distracted away from more structural reasons for why U.S. workers struggle to secure a decent living in the labor market.

Addressing these structural issues—the too-thin social insurance systems of the U.S., the impediments to organizing unions, and the failure to sustain low rates of unemployment—would not only boost workers' wage growth across the board. It would also address most of the stress on smaller groups of workers experiencing job displacement due to technological change.

There is real harm to public analysis and policymaking that focuses so much attention on each new mini wave of technological advance as a cause for workers' problems. The most obvious harm is that it's a clear misdiagnosis as source of wage suppression. If one could somehow completely ban progress on AI, this would do nothing to improve workers' lot in the future. If we could go back in time and ban research on robots or autonomous vehicles, wages today for workers would be no higher. Yet AI and robots and autonomous vehicles have sucked up more attention than the structural issues we referenced above from many who should sincerely be concerned with how U.S. workers are faring. The attention of policymakers, researchers, and advocates is a scarce resource, and every minute they are convinced they need to be constructing plans around the newest technological fad is a minute they are not working on issues of deeper concern to workers.

When this is recognized and how technology is used by workplaces becomes a focus of empowered workers, smarter workplace policy can result. Key examples of clear-eyed stances toward AI can be seen in recent negotiations between the AFL-CIO and Microsoft, and the negotiated role of AI in contract agreements between the Writers Guild of America (WGA) and the Alliance of Motion Picture and Television Producers (AMPTP). The AFL-CIO and Microsoft have recently come to an agreement that Microsoft would remain neutral in future organizing campaigns and agreed to future discussions about how AI can be used to improve workplace productivity and workers' pay and working conditions (rather than be used as a cudgel to reduce workers' leverage and bargaining power). The WGA contract with the AMPTP states that writers can use AI as a tool in their own work, but that AI cannot be used to undermine writers' claims to credit for what they produce. There will clearly be unforeseen issues that will arise going forward, but these are encouraging first steps that clearly show that, in a balanced labor market (like when a union is present), issues regarding AI (and any other technological change) have a strong chance of being settled in ways that benefit workers.

Notes

1. See Marcus (1983) for a contemporaneous account of fears concerning "technological unemployment."

2. See Engardio, Bernstein, and Kripalani (2003) for an example of fears being raised over the prospect of "white-collar offshoring."

3. See Gilbert (2013) for a piece detailing the alleged threats robots pose to human employment.

4. See Bivens and Mishel (2015) for a review of the historical interplay between wage growth for typical workers and productivity, and for a decomposition of where the wedge between these workers' pay and productivity growth went.

5. See Bivens and Mishel (2021) for a review of the research supporting the case that it is this policy-induced degradation of typical workers' labor market power that drove the sharp slowdown in wage growth for these workers and the resulting increase in wage inequality.

6. See Krueger and Solow (2002) for a deep examination of the "Roaring Nineties."

7. See Ghayad (2013) for documentation of this employer sorting by duration of unemployment.

8. See Bivens and Shierholz (2014) for real-time evidence on how deficient demand, not worker skills or employer behavior, was the real cause of elevated long-term unemployment.

9. See Modestino, Shoag, and Balance (2020) for evidence of this type of employer "upskilling" during periods of too-slack labor markets.

10. See Gould (2019) for an overview of U.S. wage trends since 1979.

11. For a time in the 1990s, globalization—not institutional change—was generally seen as the main competitor to SBTC as the dominant driver of wage inequality in the U.S. Since then, however, economists have increasingly settled on (what we consider to be) the correct view that globalization has had significant effects on wage inequality but remains insufficient to explain most of the rise in inequality. See Bivens (2017a) for a review of some of this debate on globalization.

12. Card and DiNardo (2002) and Schmitt, Shierholz, and Mishel (2013) provide extremely detailed examinations of the direct evidence supporting the SBTC view and find it lacking.

13. See Bassier, Dube, and Naidu (2022) on how individual firms do have discretion over wage levels and cannot simply hire as many workers as desired at exogenously set "market wages."

14. Lester (1952) first coined the term "range of indeterminacy" to describe the situation where a single wage might be consistent with many different employment levels. Schmitt (2013) has reviewed the research on the many margins of adjustment available to accommodate increases in mandated minimum wages.

15. For example, Ashenfelter, Card, Farber, and Ransom (2022) edited a symposium in the *Journal of Human Resources* on "Monopsony in the Labor Market," and Mishel (2022) edited a symposium in

the *Journal of Law and Political Economy* on “Not So Free to Contract: The Law, Philosophy, and Economics of Unequal Workplace Power.”

16. There is also the important issue that labor productivity is much more straightforward to measure and interpret than is total factor productivity. Because it measures increases in output growth after accounting for all observable inputs, in many measures, total factor productivity is simply a quantity representing what we cannot truly explain—it has been labelled a “measure of our ignorance.”

17. We say “nearly all” instead of “all” because potential output is not actually the maximum feasible output an economy could produce (say under conditions of wartime and price controls). Instead, it’s how much an economy can produce without spurring accelerating inflation. As aggregate demand gets extremely high relative to potential output, unemployment can be driven so low that workers’ wage demands exceed productivity growth, spurring inflation. This level of unemployment that maps onto maximum output that can be produced without accelerating inflation is sometimes called the “natural rate” of unemployment.

18. See Baker (2010) for the best macroeconomic narrative of how the housing bubble’s burst reduced aggregate demand and caused the Great Recession.

19. This dichotomy between determinants of potential output growth and aggregate demand growth is not quite as strict as this section indicates. Long periods of time when aggregate demand is depressed, for example, can actually reduce productivity growth and labor force growth as businesses invest less in labor-saving technologies and potential workers stay on the sidelines if wage growth is sluggish (see Bivens (2017b) for evidence on some of these links). Yet in well-functioning economies with responsible policymakers, policy decisions can effectively make the determinants of potential output and aggregate demand mostly separate.

20. See Bivens (2022) for an overview of the U.S. economic situation before the American Rescue Plan passed and the law’s subsequent effect on labor markets.

21. A very useful discussion and possible scenarios for AI’s effect on productivity growth over the next decade is provided by Briggs and Kodnani (2023). They estimate a 1.5 percentage point potential annual productivity growth rate boost due to the adoption of AI in the U.S., which would be likely significantly dampened by AI’s substitution away from other technologies and the possibility of a slower adoption period (say 20 years, rather than 10 years).

22. The numbers referenced here for direct, final demand, and upstream effects are very rough estimates taken from Figure 1B in Autor and Salomons (2018).

23. There is some suggestive evidence that some measures of churn—like job-to-job moves—have been increasing since roughly 2015, and that churn jumped enormously in response to the COVID-19 shock. However, as Figure E shows, even a pronounced uptick in churn relative to the recent past will likely not approach past historical peaks.

24. See, for example, the discussion of AI and employment discrimination in Kim and Bodie (2021).

25. See Mullainathan (2019) for this argument that it may be easier to correct discrimination occurring by algorithm relative to discrimination occurring by personal decision-making.

26. See Hadi et al. (2023). Note that they obviously could not assess the true IRS algorithm as this was kept confidential. Instead, they constructed their own algorithms and assessed them for how closely their predicted outcomes matched actual audit patterns.

References

Acemoglu, Daron, David Autor, and Simon Johnson. 2023. *Can We Have Pro-Worker AI? Choosing a Path of Machines in Service of Minds.* MIT Shaping the Future of Work Initiative, September 2023.

Anway, Nicholas. 2022. "Amazon's Approach to Robotics Is Seriously Injuring Warehouse Workers." *OnLabor* (blog post), May 5, 2022.

Ashenfelter, Orley, David Card, Henry Farber, and Michael Ransom, eds. 2022. "Special Issue: Monopsony in the Labor Market." *Journal of Human Resources* 57, supplement 2022: S1–S10.

Atkinson, Robert D., and Wu, J. John. 2017. *False Alarmism: Technological Disruption and the U.S. Labor Market, 1850–2015*. Information Technology & Innovation Foundation ITIF, May 2017.

Autor, David, Arindrajit Dube, and Annie McGrew. 2023. "The Unexpected Compression: Competition at Work in the Low Wage Labor Market". National Bureau of Economic Research Working Paper # 31010, November 2023.

Autor, David, Claudia Goldin, and Lawrence Katz. 2020. "Extending the Race Between Education and Technology." *American Economic Review*, AEA Papers and Proceedings 110 (May): 347–51.

Autor, David, and Anna Salomons. 2018. "Is Automation Labor-Displacing? Productivity Growth, Employment, and the Labor Share." *Brooking Papers on Economic Activity: BPEA Conference Drafts*, March 2018.

Baker, Dean. 2010. *False Profits: Recovering from the Bubble Economy.* Sausalito, CA: PoliPoint Press.

Bassier, Ihsaan, Arindrajit Dube, and Suresh Naidu. 2022. "Monopsony in Movers: The Elasticity of Labor Supply to the Firm Wage Policies." *Journal of Human Resources* 57 (S): s50–s86.

Bivens, Josh. 2017a. *Adding Insult to Injury: How Bad Policy Decisions Have Amplified Globalization's Costs for American Workers.* Economic Policy Institute, July 2017.

Bivens, Josh. 2017b. *A "High-Pressure Economy" Can Help Boost Productivity and Provide Even More "Room to Run" for the Recovery*. Economic Policy Institute, March 2017.

Bivens, Josh. 2022. "Will Friday's Wage Growth Numbers Stop the Fed from Snatching Defeat out of the Jaws of Victory?" *Working Economics Blog (*Economic Policy Institute), July 2022.

Bivens, Josh, and Lawrence Mishel. 2015. *Understanding the Historic Divergence Between Productivity and a Typical Worker's Pay: Why It Matters and Why It's Real*. Economic Policy Institute, September 2015.

Bivens, Josh, and Lawrence Mishel. 2017. *The Zombie Robot Argument Lurches On: There Is No Evidence that Automation Leads to Inequality or Joblessness.* Economic Policy Institute, May 2017.

Bivens, Josh, and Lawrence Mishel. 2021. *Identifying the Policy Levers Generating Wage Suppression and Wage Inequality*. Economic Policy Institute, May 2021.

Bivens, Josh, and Heidi Shierholz. 2014. *Lagging Demand, Not Unemployability, Is Why Long-Term Unemployment Remains So High*. Economic Policy Institute, August 2014.

Bowles, Samuel. 1985. "The Production Process in a Competitive Economy: Walrasian, Neo-Hobbesian, and Marxian Models." *American Economic Review* 75, no.1 (March): 16–36.

Briggs, Joseph, and Devesh Kodnani. 2023. “Upgrading Our Long-Run Growth Forecasts to Reflect the Impact of Generative AI.” *Global Economics Analyst* (Goldman Sachs), October 29, 2023.

Bureau of Economic Analysis (BEA). 2023. “Real Gross Domestic Product, Chained Dollars” [Table 1.1.6.], National Income and Product Accounts (NIPA). Accessed December 2023.

Bureau of Labor Statistics, Current Population Survey (BLS-CPS). 2023. “Labor Force Statistics from the Current Population Survey” (web page). Accessed December 2023.

Bureau of Labor Statistics, Office of Productivity and Technology (BLS-OPT). 2023. “Productivity,” (web page). Accessed December 2023.

Card, David, and John DiNardo. 2002. “Skill Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles.” *Journal of Labor Economics* 20, no. 4: 733–783.

Casselman, Ben. 2021. “Pandemic Wave of Automation May Be Bad News for Workers.” *New York Times*, July 3, 2021.

Congressional Budget Office (CBO). 2023. *Potential GDP and Underlying Inputs*, August 2019.

Engardio, Peter, Aaron Bernstein, and Manjeet Kripalani. 2003. “The New Global Job Shift: Is Your Job Next?,” *Businessweek Economics* (Bloomberg), February 3, 2003.

Fernald, John. 2023. Utilization-Adjusted Quarterly-TFP Series for the U.S. Business Sector, Federal Reserve Bank of San Francisco. Accessed December 2023.

Ghayad, Rand. 2013. “The Jobless Trap.” Working Paper.

Gilbert, James. 2013. “Robot Job Growth Is an Accelerating Problem, Researchers Tell 60 Minutes.” *Huffington Post*, January 14, 2013.

Gould, Elise. 2019. “Decades of Rising Economic Inequality in the U.S.” Testimony before the U.S. House of Representatives Ways and Means Committee, March 27, 2019.

Hadi, Elzayn, Evelyn Smith, Thomas Hertz, Arun Ramesh, Robin Fisher, Daniel Ho, and Jacob Goldin. 2023. “Measuring and Mitigating Racial Disparities in Tax Audits.” Stanford Institute of Economic Policy Research Working Paper, January 2023.

Kim, Pauline, and Matthew T. Bodie. 2021. “Artificial Intelligence and the Challenges of Workplace Discrimination and Privacy.” *Journal of Labor and Employment Law* 35, no. 2: 289–315.

Krueger, Alan, and Robert Solow, eds. 2002. *The Roaring Nineties: Can Full Employment Be Sustained?* New York, N.Y.: Russell Sage Foundation.

Lester, Richard. 1952. “A Range Theory of Wage Differentials.” *ILR Review 5*, no. 4 (March): 483–500.

Manning, Alan. 2003. *Monopsony in Motion: Imperfect Competition in Labor Markets*. Princeton, NJ: Princeton University Press.

Marcus, Steven. 1983. “Technology: As Computers Eliminate Jobs.” *New York Times*, June 2, 1983.

Mishel, Lawrence, ed. 2022. “Special Issue: Not So Free to Contract: The Law, Philosophy, and Economics of Unequal Workplace Power.” *Journal of Law and Political Economy* 3, no. 1.

Modestino, Alicia Sasser, Daniel Shoag, and Joshua Balance. 2020. “Upskilling: Do Employers Demand Greater Skill When Workers Are Plentiful?” *Review of Economics and Statistics* 102, no. 4:

793–805.

Mullainathan, Sendhil. 2019. “Biased Algorithms Are Easier to Fix Than Biased People: Racial Discrimination by Algorithms or by People Is Harmful – But That’s Where the Similarities End.” *New York Times*, December 6, 2019.

Organization for Economic Cooperation and Development (OECD). 2024. “Employment by Education Level” (indicator), https://doi.org/10.1787/26f676c7-en (accessed on February 4, 2024).

Schmitt, John. 2013. *Why Does the Minimum Wage Have No Discernible Effect on Employment?* Center for Economic Policy Research, February 2013.

Schmitt, John, Heidi Shierholz, and Lawrence Mishel. 2013. “Don’t Blame the Robots: Assessing the Job Polarization Explanation of Growing Wage Inequality.” Economic Policy Institute-Center for Economic Policy Research Working Paper, November 2013.

Sedik, Tahsin Saadi. 2021. “Pandemics and Automation: Will the Lost Jobs Come Back?” International Monetary Fund Working Paper No. 2021/011, January 2021.

Skott, Peter and Fred Guy. 2013. “Power, Luck, and Ideology: Technological and Institutional Parameters of the Agency Problems for CEOs.” *Review of Radical Political Economics* 45, no. 3.

Stansbury, Anna M., and Lawrence H. Summers. 2020. “Declining Worker Power and American Economic Performance.” *Brooking Papers on Economic Activity: BPEA Conference Drafts*, 2020.