THE AMERICAN ECONOMY IN A NEW ERA

ASPEN ECONOMIC STRATEGY GROUP
aspen institute

CHAPTER

Skills Alignment for the AI Economy: A Framework for US Labor Market Policy

by Robert Seamans

Disclaimer: This paper was written for the Aspen Economic Strategy Group. The views expressed herein are those of the authors and do not necessarily reflect those of the Aspen Institute or the Aspen Economic Strategy Group members.

Suggested Citation: Seamans, Robert. 2026. "Skills Alignment for the AI Economy: A Framework for US Labor Market Policy" In *The American Economy in a New Era*, edited by Melissa S. Kearney and Luke Pardue. Washington, DC: Aspen Institute.

SEPTEMBER 2026

Skills Alignment for the AI Economy: A Framework for US Labor Market Policy

SEPTEMBER 2026

AUTHOR

Robert Seamans*

ABSTRACT

Generative AI is diffusing rapidly across the US economy, and speculation abounds as to its likely effects on the labor market. This paper presents evidence from the most recent data available, documenting strong task-level productivity gains but limited evidence of broad, aggregate labor-market disruption to date. This combination of large effects at the task level but limited aggregate disruption so far calls for policies that improve labor market adjustment, rather than policies that assume mass job losses. In light of this evidence, I argue that the appropriate policy response should be designed to address information asymmetries about future skill demand and barriers to worker reskilling. Importantly, any policy response should be designed to benefit both firms and workers. Rather than creating AI-specific assistance programs or slowing adoption through targeted taxes, policymakers should strengthen the institutions that connect workers, employers, and training providers. This approach includes piloting AI-enabled skill training, supporting employer-based training, modernizing transition supports, and investing in real-time labor-market data infrastructure so that information about changing skill demand flows more effectively through the labor market.

* Professor of Management and Organizations, New York University, Stern School of Business

Acknowledgements: Thank you to Harry Holzer, Chris Hughes, Ben Hyman, Melissa Kearney, Rachel Lipson, Sam Manning, Luke Pardue, and John Soroushian for comments, suggestions, and edits.

Introduction

Generative artificial intelligence (AI) is poised to dramatically alter the nature of work and the demand for skills, but in ways that remain highly uncertain. What is clear, however, is that there is a need for our country to be "policy-ready" for whatever changes AI brings to workers and firms. This policy readiness should aim to buffer workers from job losses in areas where demand for certain types of skills is shrinking and redirect workers to areas in which demand for other types of skills is increasing. Policies that accomplish these goals will not only help workers weather any disruption from AI but will also help firms to more quickly adjust to changing market conditions and new opportunities, thereby helping to drive economic growth. Rather than viewing labor policy primarily as helping workers adapt, this article argues that policy should improve the process through which employers communicate changing skill needs and workers acquire the skills needed to meet them.

Generative AI is widely considered to be a general-purpose technology (GPT), much like the steam engine, electricity, and personal computers. Though AI presents its own specific opportunities and challenges, the history of GPTs yields three broad lessons that ought to guide our policy response to AI's diffusion. First, diffusion takes time: Firms must redesign processes, train workers, and make complementary investments before new technologies translate into broad productivity gains. Second, employment effects may lag task-level change: AI may reshape what workers do before it changes how many workers firms hire. Third, the effects will be uneven: Some occupations, firms, regions, and groups of workers will adjust more easily than others. The implication for policy is that rather than bet on a single forecast of AI-driven job loss or job creation, policymakers should build institutions that help workers and firms adjust as the evidence comes in.

Those general lessons notwithstanding, there are also reasons to expect that the case of generative AI may be different—and consequently have different implications for the labor market. First, occupations that involve cognitive tasks are likely more directly exposed to generative AI than previous technologies, which, in the case of steam and electricity, primarily affected physical tasks (Felten et al. 2024). Second, generative AI's user interface allows workers to experiment with it before firms have formally reorganized around it. This informal experimentation creates a gap between worker use and firm adoption, between task-level productivity and firm-level performance, and between technical exposure and realized labor-market effects. Because this informal adoption outpaces formal adoption, policy may already be behind by the time aggregate statistics clearly reveal the direction of change.

In addition, generative AI is being commercialized against an economic backdrop different from those of prior GPTs. Labor market participation has been declining over the past 70 years for men and the past 20 years for women, while business concentration has increased (Furman and Seamans 2019). People are moving from one US geography to another at slower rates than in the past (Jones et al. 2025). Occupational churn has slowed over time. The years spanning 1990 to 2017 were less disruptive than any prior period measured, going back to 1880, as Deming, Ong, and Summers (2024) show in an earlier AESG paper. Gig work, temporary work, and other forms of independent work have grown substantially over the past decade (Sundararajan 2025). Political polarization and US federal debt are rising (Klein and Obstfeld 2023).

These are all important trends to monitor and consider when evaluating policy proposals. Because of the rise in gig work, there are now more workers not covered by traditional safety-net programs such as unemployment insurance (UI). Workers appear to face a number of frictions in the labor market, as evidenced by declining churn and labor market participation. Less frequent moves suggest that training and retraining programs need to be considered in light of market demand at the local level, as opposed to the national level. The implications of increasing business concentration are less clear, but one potential worry is that it leads to increased monopsony, whereby there are fewer firms competing for workers, leaving workers in a worse bargaining position. Political polarization suggests that policies traditionally favored by constituents of one political party but not the other will be harder to pass. Rising debt makes it harder to pay for new policies. With this backdrop in mind, policies will need to not only benefit workers but also benefit firms looking to hire workers with the skills needed to meet changing demand.

This chapter proceeds in four steps. Sections 1 and 2 present data about the extent of AI adoption by US firms and workers, as well as survey evidence regarding productivity, exposure, job displacement, and job creation. Section 3 develops the Skills Alignment Framework, which emphasizes reducing information asymmetries between workers, employers, and training providers. Section 4 proposes a set of policies within this framework, highlighting the need to strengthen matching between skill supply and employer demand by expanding training and skill-discovery systems, and to use AI where possible to help do so; encouraging employer investment in training; modernizing unemployment insurance; and building better real-time labor market data. I conclude the paper by emphasizing a few simple principles: Robust labor-market policies benefit workers and also firms, and a robust governmental data infrastructure, which can include partnering with private companies, should be a top priority.

1. AI Adoption

In this section, I present data on how AI has been adopted by firms and workers to date.

1.1 Firm-Level Adoption

The best official data suggest that business use of AI has risen quickly but remains far from universal. The US Census Bureau's Business Trends and Outlook Survey (BTOS) has become the most important high-frequency source for firm-level AI adoption statistics. Since September 2023, BTOS has included questions about the current use of AI and expected use of AI in six months. BTOS updated its questions in November 2025, so figures showing the trend over time exhibit a discrete jump up around this time period.

A recent study using the BTOS data reports that, during the November 2025 to January 2026 reference period, roughly 18 percent of firms used AI in at least one business function, equating to approximately 32 percent on an employment-weighted basis (Bonney et al. 2026). The study also reports substantial heterogeneity in adoption by industry, with Information and Professional, Scientific, and Technical Services being the leading industries.[1] The most recently available BTOS data as of August 2026 suggests that approximately 22 percent of firms use AI, and 26 percent expect to use it within six months.[2]

The BTOS estimates are higher than the pre-generative-AI baseline. Using the US Census Bureau's 2018 Annual Business Survey, one study documented relatively low adoption of AI technologies prior to the "ChatGPT moment" in late November 2022 (McElheran et al. 2024). They report 8 percent adoption, and approximately 18 percent on an employment-weighted basis. That earlier evidence is useful because it shows that the availability of AI technology does not imply immediate organizational uptake. Generative AI has lowered the cost of experimentation, but the movement from experimentation to production remains uneven.

1.2 Worker-Level Adoption and "Shadow AI"

Worker-level adoption appears to exceed formal firm-level adoption. Many workers use ChatGPT, Claude, Gemini, Copilot, or other tools without a firm-wide deployment decision. The Council of Economic Advisers (CEA) combines several data sources on the use of generative AI at work and shows a large post-ChatGPT increase (CEA 2026). This "shadow AI" matters for two reasons. First, workers may be learning to use AI faster than firms are learning how to reorganize around it. Second, worker-level use can create productivity gains and new skill demands before human resources systems or occupational data record any change.

One notable survey of individual generative AI use is the individual-level Real-Time Population Survey (RPS), which has asked respondents about generative AI adoption on a quarterly basis since August 2024 (Bick et al. 2026). They report that, as of May 2026, the share of US adults using generative AI at work is approximately 45 percent (and approximately 62 percent use generative AI at home or work). A recent Federal Reserve report suggests that the differences between BTOS and RPS are likely due to a number of factors, including survey sample, question framing, survey respondent characteristics, and other factors (Allen 2026). Directionally, these

[1] The five sectors exhibiting the highest AI adoption are Information (38 percent); Professional, Scientific, and Technical Services (34 percent); Educational Services (31 percent); Finance and Insurance (30 percent); and Real Estate (24 percent). According to the data, among adopters, scope remains limited: 57 percent use AI in three or fewer functions, most often Sales and Marketing (52 percent), Strategy (45 percent), and IT (41 percent).
[2] See US Census Bureau, "Business Trends and Outlook Survey," Census.gov, updated August 27, 2026, https://www.census.gov/hfp/btos/data.

surveys provide similar results: AI adoption is increasing over time but is still lower than press headlines might suggest. The differences suggest that definitional boundaries about what one means by “AI” are unstable, and they highlight that careful definition and measurement need to become a policy priority.

The policy implication is that firm-level adoption rates understate the degree to which AI is already entering production. However, worker-level use also should not be mistaken for full economic transformation. A worker using AI to draft emails or summarize documents is not the same as a firm redesigning workflow, changing staffing ratios, hiring employees with new sets of skills, etc. For labor market forecasting, an important question is whether informal “shadow” use becomes embedded in production systems.

1.3 Uneven Individual Adoption Across Geographies and Demographics

A recent Microsoft report, based on telemetry data that captures individual users’ interactions with generative AI tools such as ChatGPT, Google Gemini, Anthropic Claude, and Microsoft Copilot, estimates that approximately 31 percent of US adults use generative AI (Misra et al. 2026). This figure is lower than the figure reported by RPS (62 percent), a difference that may be due to sample construction and/or definitions of use. For example, Microsoft restricts their sample to users with at least 90 minutes of usage time in a given month, a constraint that may be more restrictive than any imposed on the RPS sample. Also, Microsoft bases their data on laptop and personal computer usage, whereas it is possible that more respondents rely on generative AI via smartphones—a use pattern that the RPS data may be picking up.

The Microsoft report highlights substantial heterogeneity across geographies. There is a stark urban-rural divide that is reminiscent of the broadband digital divide. Counties with large metropolitan areas have approximately 33 percent adoption, whereas rural counties have approximately 16 percent adoption. Counties with college towns, which have a higher proportion of young adults, tend to have higher adoption rates, including multiple counties with greater than 60 percent adoption.

Other industry reports include those released by Anthropic and OpenAI, which use data on Claude and ChatGPT, respectively, to document patterns in adoption and use (Appel et al. 2025; Chatterji et al. 2025). These reports also highlight uneven adoption across countries and US states, as well as uneven adoption across demographics including age group and gender.

The reports from these companies are all notable as they provide unique insights into adoption and usage patterns of generative AI, and how these patterns evolve over time. For a variety of reasons, including the protection of individual user privacy, the use of the underlying data is restricted to employees at these companies, and in some cases to selected external economists. As will be discussed in a later section, mechanisms to allow greater access to close-to-real-time data usage, at a granular level, by a wider range of government and external economists will allow for more widespread and timely research on the effects of AI on the economy.

Table 1 below summarizes the results from selected surveys to highlight the differences in reported adoption rates, as well as the heterogeneity in adoption across firms, individuals, and geographies.

2. AI and Labor Market Effects

AI can affect jobs through multiple channels. It can directly substitute for work in some cases, leading to job losses. AI can also increase demand for other work or lead to the creation of new work, leading to job gains. More broadly, AI can lead to lots of changes in terms of how we all do work; that is, the job remains but the nature of the job changes.

The labor market evidence can be summarized as follows. First, generative AI can raise productivity substantially in controlled settings and in some real workplaces. Second, these gains are heterogeneous, often benefiting less experienced workers most. Third, aggregating from studies in controlled settings to the firm level or economy level is not straightforward, and much more research is needed. Fourth, aggregate labor-market data have not yet shown broad AI-driven displacement. Fifth, specific groups and occupations may nevertheless be experiencing early effects, especially entry-level workers and software-related occupations, though it is difficult to assess how much observed changes are due to AI or other macroeconomic factors. Finally, some evidence suggests that new, AI-specific jobs are being created and experiencing growth.

2.1 Layoff Announcements Versus Realized Labor-Market Changes

Layoff announcements have created a public impression that AI is already driving widespread job loss. Reports by Challenger, Gray, and Christmas, which aggregate public announcements of layoffs, show an increase in announced layoffs attributed to AI or automation (Challenger et al. 2026). However, announced layoffs are not the same as realized separations, and firms might have incentives to frame restructuring as AI-related, rather than as being due to macroeconomic uncertainty, lower-than-expected consumer demand, or poor strategic planning.

Table 1: Generative AI adoption rates across different units of analysis

Survey	Unit	Date	Adoption rate & notes
US Census BTOS	Firms	August 2026	22% (26% expected in six months)
Microsoft	Counties	Q1 2026	31% (33% metro vs. 16% rural counties)
RPS	Individuals	May 2026	62% (45% at work)

ASPEN ECONOMIC STRATEGY GROUP aspen institute

Sources: US Census BTOS: Bonney et al. 2026; Microsoft: Misra et al. 2026; RPS: Bick et al. 2026..

Official labor-market data have not yet shown any evidence of AI-driven displacement. Every month the US Bureau of Labor Statistics produces estimates of job openings, hires, and separations as part of its Job Openings and Labor Turnover Survey (JOLTS) program. JOLTS data is more representative of the economy than the Challenger data is. For example, Challenger, Gray, and Christmas data records 60,000 job cut announcements in March 2026, whereas the JOLTS data suggests there were approximately 1.8 million job cuts in the same time period (Darling 2026).

Another way to track AI-related displacement is through notices filed under the federal Worker Adjustment and Retraining Notification (WARN) Act, which requires large employers to provide 60 days' advance notice of plant closings or mass layoffs. Some states impose additional WARN requirements. Since 2025, New York has required employers to report whether layoffs are attributable to AI. Yet, according to an analysis published in *Wired*, no company filing a layoff notice in New York State had attributed those layoffs to AI as of January 2026 (Dave 2026).

2.2 Exposure Measures

Since we are in the early innings of AI adoption and do not yet know where and how AI will be adopted, economists often rely on "AI exposure" to identify which sectors will be most affected.[3] A consistent finding from these studies is that generative AI exposure is concentrated in cognitive, professional, and white-collar work (see, for instance, Eloundou et al. 2024; Felten et al. 2024; Webb 2019).

For example, as depicted in figure 1, occupations with higher median salaries, greater educational requirements, and a higher emphasis on creative abilities are those most exposed to generative AI (Felten et al. 2024).

[3] AI exposure measures are attempts to quantify how much an occupation, task, industry, or geography is potentially affected by AI capabilities. They do not directly measure adoption, productivity effects, job loss, or whether AI substitutes for or complements workers. Examples include Felten et al. 2021, which measures AI exposure by linking progress in AI application areas to O*NET abilities used in occupations; Eloundou et al. 2024, which measures LLM exposure by using humans and LLMs to assess whether LLMs could affect occupational tasks; and Webb 2019, which measures AI exposure by comparing the text of AI patents to the text of job task descriptions.

The Yale Budget Lab's synthesis of exposure measures emphasizes three points: Exposure metrics broadly agree on which occupations are not exposed, disagree more on the magnitude of exposure among highly exposed occupations, and do not by themselves imply automation or job loss (Gimbel et al. 2026). The distinction between exposure and displacement is important. An occupation can be highly exposed because AI can accelerate tasks performed by workers in that occupation. That exposure could increase demand if AI raises productivity, lowers prices, improves quality, or expands output. It could reduce demand if the same output can be produced with fewer workers. It could change the composition of tasks without changing headcount.

Figure 1: Occupational characteristics of exposure to generative AI

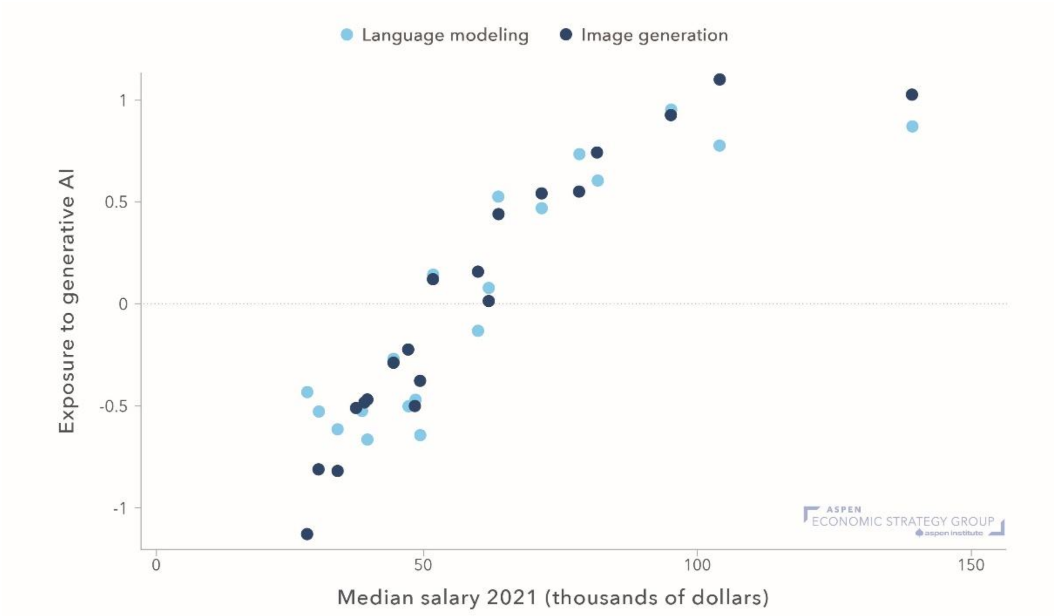

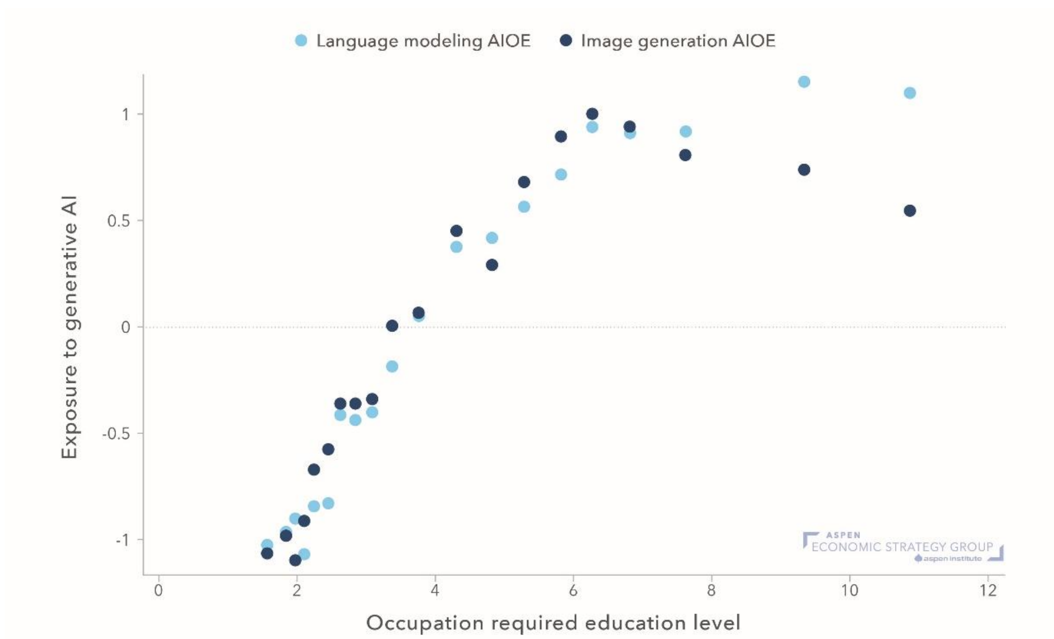

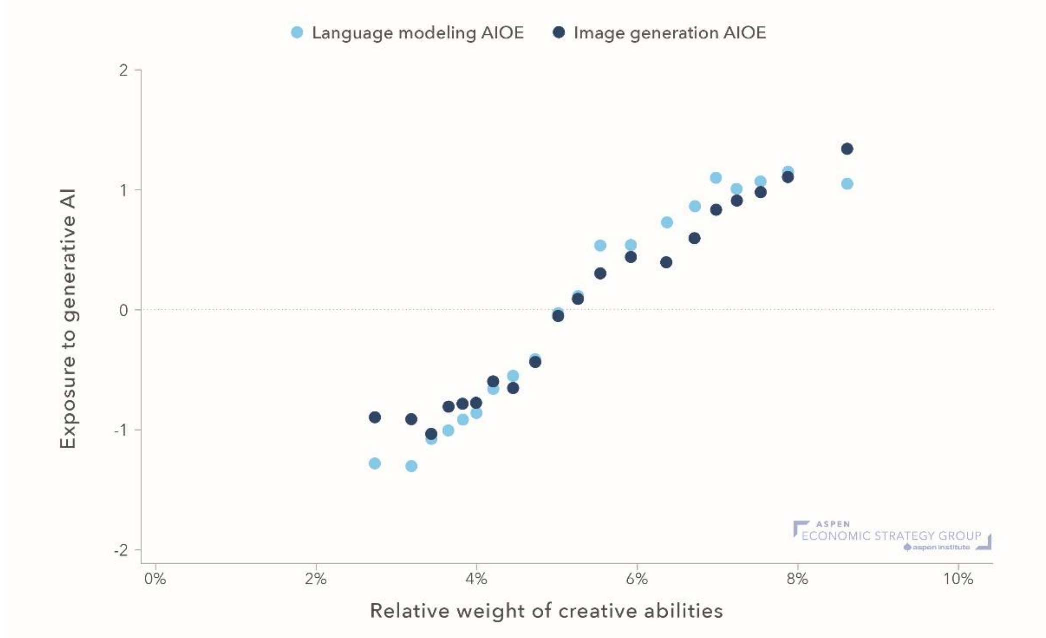

Source: Felten et al. 2024.
*Notes: This figure plots the relationship between occupational exposure to generative AI (y-axis), which includes language modeling or image generation as indicated in the legend, and occupational characteristics (x-axis). Occupations are grouped into 20 equal-sized bins based on their x-axis value. Data on median salary comes from 2021 Bureau of Labor Statistics estimates. Data for required level of education is constructed as the weighted average of required education reported across an occupation by O*NET. Data for the presence of creative abilities is constructed as the relative weight of the two creativity-oriented abilities (originality and category flexibility) among the portfolio of the 52 abilities O*NET uses to characterize occupations.*

Outcomes are shaped more by adoption patterns, costs, organizational readiness, regulatory constraints, and complementary investments than by technical feasibility alone. The exposure literature is therefore best used as a map of where in the economy to monitor for job gains and losses, and more generally for where work will likely change the most.

2.3 Entry-Level Work and the Job Ladder

An important current debate concerns whether AI is weakening the first rung of the job ladder. This concern is plausible because entry-level roles often consist disproportionately of tasks that are codifiable, reviewable, and useful for training. Examples include drafting, summarizing, coding, checking, preparing slides, conducting preliminary research, and responding to routine customer issues. If AI can adequately perform these tasks, firms may reduce junior hiring in favor of deploying AI.

Several recent studies provide suggestive evidence by comparing how employment for different types of workers changed following the release of ChatGPT in late November 2022. The idea is not that ChatGPT is directly responsible for these changes, but rather that the intense press in the weeks and months that followed, coupled with the subsequent release of similar large language models (LLMs) by Google and others, caused businesses to become much more aware of the power of this technology, and potentially to experiment with using it for labor-replacing purposes.

A 2025 study entitled "Canaries in the Coal Mine?" documents a slowdown in hiring of younger workers in AI-exposed industries (Brynjolfsson, Chandar, et al. 2025). Using private administrative data from ADP, a payroll processing firm, the study shows that, following the arrival of ChatGPT, employment lagged for entry-level workers in these exposed industries, while workers at other career stages did not experience similar effects (Brynjolfsson, Chandar, et al. 2025). A study of coders, who are particularly exposed to the type of work that generative AI can perform, uses publicly available data from the Current Population Survey (CPS) and finds that coder employment has continued to grow, but at a much slower rate than before the arrival of ChatGPT (Crane and Soto 2026).

However, other research finds little evidence of these effects. Using CPS data, one study finds that unemployment growth since 2022 has been concentrated in the least AI-exposed occupations (Eckhardt and Goldschlag 2025). Another study finds that job postings began declining more in AI-exposed occupations in early 2022, well before the ChatGPT shock, and argues that these patterns are more likely explained by interest-rate-sensitive sectors and the timing of monetary tightening than by the release of ChatGPT itself (Iscenko and Millet 2026). Real-time CPS tracking finds no statistically significant relationship between measures of AI exposure, automation, or augmentation and changes in employment or unemployment at the aggregate level (Yale Budget Lab 2026).

Recent evidence suggests that the slowdown in hiring for junior-level, AI-exposed jobs such as software engineering may be less a story about AI replacing entry-level workers than about the organizational consequences of remote work. Occupations highly exposed to generative AI are also likely to be highly exposed to work from home. One recent study finds that once both exposures are included in the same empirical design, the apparent AI effect on junior hiring largely disappears while the work-from-home effect remains (Lambert and Schindler 2026). Consistent with this interpretation, recent work by the New York Fed argues that remote work makes it harder for managers to train and mentor young workers, and that the timing of rising unemployment among young college graduates points more to remote work than to generative AI (Emanuel et al. 2026).

One recent research paper links firm-level AI spending data from Ramp with workforce outcomes data from Revelio across roughly 21,600 US companies (Kharazian et al. 2026). The analysis finds that faster post-adoption employment growth is concentrated almost entirely among high-intensity AI adopters, who see employment gains of about 10 percent, while lowerintensity adopters show no measurable effect. The disparity is even sharper for entry-level roles, where high-intensity AI adopters post a 12 percent increase in headcount. Taken together, the results described in this subsection suggest, at face value, that AI may already be affecting particular entry-level pathways, especially in software and some professional services, while not yet producing any broad labor-market displacement signal. As will be discussed in a later section, policy should therefore focus not only on things like unemployment insurance after job loss but on the architecture of transitions, which might include internships, apprenticeships, early-career training, on-the-job learning, and credible skill signals.

2.4 AI Increases Individual Productivity

The strongest evidence for labor market effects comes from carefully designed productivity studies, not employment studies. A random assignment study of software coders using GitHub Copilot, a generative AI software coding tool that suggests blocks of code in real time to coders, much in the way that predictive typing suggests sentence completion to writers, found that coders with access to the tool were able to complete a specified task approximately 50 percent faster, with much of the improvement accruing to coders with less experience (Peng et al. 2023). A random-assignment study of college-educated professionals using ChatGPT on writing tasks found large reductions in time and improvements in output quality (Noy and Zhang 2023). The gains were not uniform: Lower-performing workers benefited more, compressing the productivity distribution. A study of the staggered introduction of a generative AI assistant among customer-support agents found an average productivity increase of about 15 percent, with especially large gains for novice and lower-skilled workers (Brynjolfsson, Li, et al. 2025). Figure 2 summarizes the findings from these studies. A variety of other papers mostly confirm this pattern: Generative AI appears to boost productivity, especially for workers with less experience.

It is less clear how this pattern might inform the labor market, because a number of difficulties arise when generalizing from a carefully controlled experiment to work in an organization. Moreover, aggregating from some areas in an organization to the organization as a whole is also difficult. For example, in the carefully constructed research settings above, the workers in question were either placed in a “treatment” or a “control” group (the former received use of the technology and the latter did not). Outside an experimental setting, a worker would need to be incentivized to adopt and use a technology. In addition, while the tasks in the experiments above were chosen in part because they are easy to measure (e.g., lines of code or calls handled), actual work might be less easy to measure, or the outcome might be measurable only years in the future (this is especially the case with R&D-type tasks).

Figure 2: Generative AI gains are larger for less experienced individuals

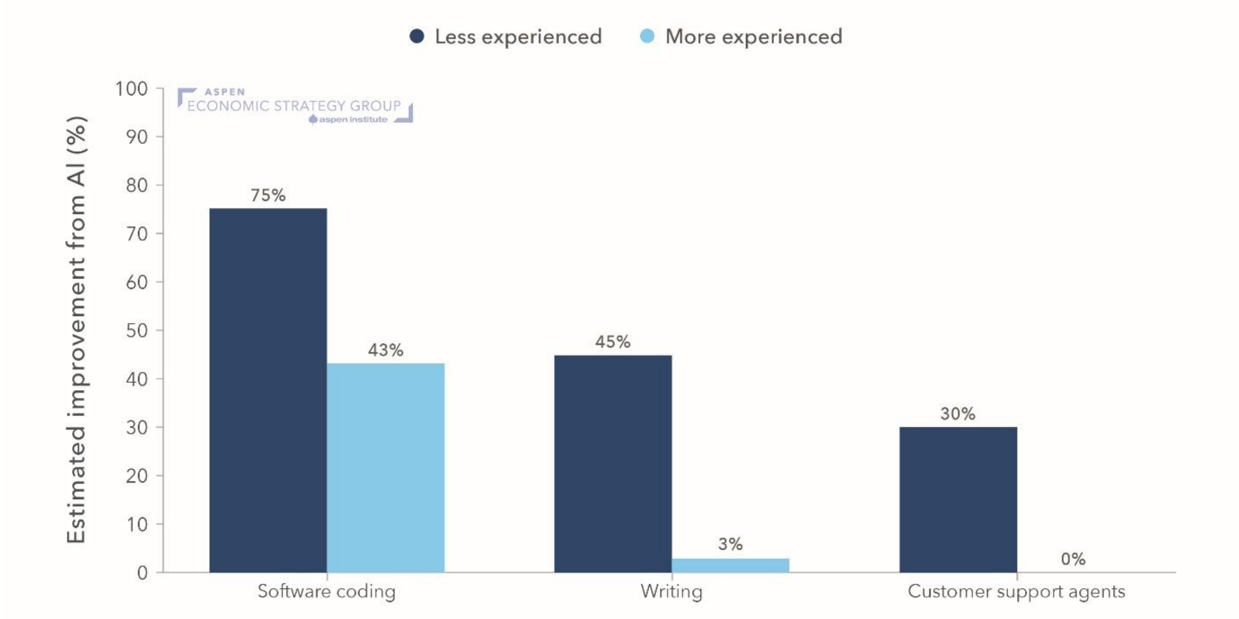

Sources: Software coding estimates are inferred percent reductions in completion time from Peng et al. (2023) for roughly the 15th vs. 18th percentile of programming experience. Writing estimates are grade-quality gains from Noy and Zheng (2023) for lower vs. higher performers. Customer support estimates are Brynjolfsson, Li, et al.(2025) resolutions-per-hour gains for least vs. most experienced agents.

Some studies reinforce how these organization-level effects are difficult to discern. AI can improve performance inside what has been called its "jagged frontier" while reducing performance when users apply it to tasks where it is unreliable (Dell'Acqua et al. 2026). Using data on individual Danish users mapped to firms, one study finds widespread adoption in exposed occupations but negligible effects on earnings and hours (Humlum and Vestergaard 2025). The implication is that task-level productivity gains do not automatically or even necessarily aggregate into employment, wage, or firm-level effects.

It is possible that these findings are specific to the short run, and that effects would look different over a longer time horizon. As generative AI improves, the "jagged frontier" may become less pronounced. Similarly, the Humlum and Vestergaard findings of negligible short-run effects are perhaps not surprising, since prior research on general-purpose technologies suggests that productivity gains often require time to emerge after adoption, as firms make complementary investments and reorganize production (David 1990; Brynjolfsson et al. 2021).

In addition, as more firms experiment with new uses for generative AI, there may be spillovers between firms. An interesting example of such a spillover was documented in one study whereby a set of AI startups received information about how other firms had reorganized their production around AI, which prompted the startups to search for use cases across a broader set of functions than they had before (Kim et al. 2026). The authors find that the changes resulted in meaningful performance gains, including completing more tasks, acquiring more customers, and generating higher revenue.

2.5 Job Creation, New Work, and Firm Formation

Historically, technological change has eliminated some tasks and occupations while creating others. By the 2010s, most people worked in occupations that did not exist in 1940, and these new occupations emerged in response to technological innovations that complement occupational outputs as well as to demand shocks that raise occupational demand (Autor et al. 2024). We are already seeing evidence of new jobs in the AI context. Emerging AI-linked roles—including data annotators, AI forensic analysts, heads of AI, AI engineers, and forward-deployed AI engineers—are among the largest or fastest-growing AI-related job categories, alongside increased demand for data center construction and maintenance roles (LinkedIn 2026).

The AI landscape is also likely to lead to the creation of many new firms. Generative AI can reduce the cost of software development, content production, customer support, and market testing, thereby potentially reducing barriers to entry for new firms. New business applications by solo founders tracked by the US Census Bureau have been increasing since late 2024, and recent analysis by Stripe suggests that sectors of the economy with the most AI adoption are seeing the largest increase in these solo ventures (Tedeschi et al. 2026). Recent research using Chinese firm registration data shows a large increase in startup entry following the launch of ChatGPT, but primarily in areas that have a large preexisting stock of AI-related human capital (Cai et al. 2025).

Much more research is needed in this area, which raises a variety of follow-up questions: What are the characteristics of the entering firms? How well do they survive and grow? What types of infrastructure and support do these firms need to survive and grow? More real-time data from US statistical agencies would be quite valuable in approaching these questions—a point I return to below.

Taken together, the evidence points less toward a single, immediate labor-market shock than toward a potential transition problem: AI is likely changing tasks faster than institutions can observe, classify, and respond to those changes. The relevant policy challenge is therefore not simply to compensate workers after displacement but to help workers, firms, and training providers discover where skill demand is moving and to adjust before mismatches become costly.

3. The Skills Alignment Framework

The “Skills Alignment Framework” starts from this distinction between technical exposure and realized labor-market change. If AI’s effects depend on adoption, complementary investment, local labor-market conditions, and organizational redesign, then policy should focus on the institutions that help workers and firms continuously update their expectations about which skills are valuable.

Consider the following thought experiment, which will help motivate the Skills Alignment Framework. Imagine a company that produces a wide range of products and services for a range of different types of customers, using a variety of workers with different skillsets. The company regularly surveys its customers to understand what demand for its products and services will be 12, 24, and 36 months in the future. It then translates this demand into a set of skills that it needs in order to meet this future demand. The company also regularly surveys its workforce to understand what set of skills the workforce has and matches the existing set of skills with the forecasted demand for skills. The company identifies where there are gaps and develops training modules to address these gaps. The company then incentivizes its workforce to undertake these training modules to ensure it has the workforce with the skills necessary to meet its customers' anticipated demand. The incentives could be via higher wages for workers who develop proficiency in certain skills, or via a mandate, or via leaders or other high-performing workers in each unit who demonstrate the importance of the new skills. Imagine also that the surveys the company uses to assess what their customers need and what skills its workforce has are accurate, timely, and can adjust to changing needs of customers and changing skillsets of the workforce.

The purpose of the thought experiment is to motivate how we should think about policy responses to AI or any other new technology. The thought experiment highlights a few important points that together comprise the Skills Alignment Framework.

First, in the example above, if the company cannot obtain accurate and timely data on what its customers want and on the skills that its workforce has, it cannot then identify gaps in skills. For our economy, this point highlights the importance of a robust statistical infrastructure, without which we will be flying blind. Policymakers need data to distinguish cyclical weakness from technology-driven displacement. Without better information, subsidies may fund low-value training and workers may spend time and money on low-value credentials.

Second, in the example above, in order to produce the goods and services required by its customers, the company wants to know the gaps between skills demanded by its customers and skills its employees have. For our economy, this point highlights that robust training and retraining programs are not just to the benefit of workers but also to the benefit of firms hiring those workers. In other words, labor policies should benefit both workers and employers, and thus both should have skin in the game.

Third, the hypothetical company needs to identify what skills its workforce already has and incentivize its workers to undertake training for new skills. In the example above, a mechanism exists to align these incentives. For our economy, this point highlights the presence of an information asymmetry about demand for and supply of skills. In the US economy, training is often done by a third party, not the employers themselves (with the exception of employersponsored workforce development programs, which will be discussed in more detail below). Workers and training providers cannot easily observe which skills employers will value in the future. Employers themselves may not know which AI use cases will scale. Standard occupational categories provided by statistical agencies may be updated too infrequently and may be too coarse to capture task reallocation within jobs. This information problem has multiple sides. Workers need information about which transitions are feasible given their current skills. Training providers need information about demand. Employers need credible signals about the skills that workers have.

Recent place-based industrial policies point in a similar direction. The workforce provisions of the CHIPS and Science Act explicitly recognize that successful workforce development requires coordination among employers, educational institutions, and government, with training designed around the skill needs of regional industries rather than around generic credential attainment (Ross and Muro 2024). This point illustrates a broader principle: Effective labor policy should improve the flow of information from employers to workers and training providers so that skill investments reflect evolving labor demand rather than historical patterns.

Finally, in the hypothetical example, there were no barriers for the workers to undertake retraining, whereas in the real world, acquiring new skills may be costly for workers. If the courses are taken via an educational institution, tuition or other out-of-pocket costs may be associated with the training. The training programs themselves take time, and the opportunity cost of the time the worker spends in the training program is the foregone wages from employment.

In summary, the Skills Alignment Framework considers whether policies help identify emerging gaps between the skills workers have and the skills employers need; reduce information asymmetries among workers, firms, and training providers; create incentives for both workers and employers to invest in relevant training; and build the data infrastructure needed to monitor where interventions are most needed and whether they are working. The framework will help us assess policies in light of how they address these issues. Viewed this way, the purpose of labor market policy is not simply to increase human capital but to improve the matching process through which workers acquire the skills employers need and employers communicate the skills they expect to need.

4. Policy Recommendations

A mix of policies will be needed to address the imminent needs of both workers and companies in the era of generative AI. In what follows, I lay out several specific policy pillars. The policies below should be read as complementary pieces of a transition architecture. Some policies, such as AI-enabled skill training and workforce development programs, directly address skill discovery and acquisition. Others, such as unemployment insurance and wage insurance, reduce the risks workers face while moving between jobs. Still others, especially data and measurement infrastructure, make the entire system more adaptive by helping policymakers observe which workers, occupations, regions, and firms are experiencing change. The common thread is that policy should improve matching and adjustment rather than attempt to identify a single class of "AI-displaced" workers after the fact.

4.1 Promote AI-Enabled Skill Training

One of the most notable findings from the research on generative AI and productivity is that, for the most part, generative AI appears to provide the largest productivity benefits to those with the least amount of experience. The implication of these findings is potentially profound when it comes to labor policy around training and retraining. These studies suggest that generative AI could be used to retrain workers in new skills and for new jobs in a relatively quick and efficient way. In addition, machine learning (ML) is well known for its predictive capabilities. Together, these points suggest that ML could be used to more accurately predict how a worker's skills and interests fit with a potential job opening. All these observations highlight that AI itself could be a useful tool that is part of labor policy responses to the potential disruptions from AI. This policy directly applies the Skills Alignment Framework: AI-enabled training can help identify the gap between a worker's current skills and the skills demanded in nearby occupations, while also lowering the time and information costs of closing that gap.

The experience of private companies using AI as such a tool, together with nascent academic research on the topic, provides confidence that AI-enabled skills training can be an effective labor policy. IBM's "Your Learning System," for example, uses AI to evaluate employees' existing skills, recommend individualized pathways for skill development, and help design and update training programs (Qin and Kochan 2020).

Academic research also shows how workers' existing skills can be used to identify higher-paying occupational opportunities. One study focuses on workers without bachelor's degrees (BAs) and combines O*NET (Occupational Information Network) measures of occupational skill requirements with Current Population Survey labor market data to identify higher-wage jobs that are relatively close to a worker's current occupation in terms of "skill distance." The authors find that many non-BA workers may already have skills compatible with higher-paying jobs, including roughly 30 million "Rising STARs" with plausible pathways into middle- or high-wage work (Blair et al. 2020). A more recent study applies a similar approach to AI upskilling under Workforce Innovation and Opportunity Act (WIOA) training programs, estimating which occupations are most likely to yield successful transitions based on workers' backgrounds and the historical success rates of similar workers moving into new roles (Hyman et al. 2025). The goal is not simply to teach workers how to use AI but to use AI to make the training system itself more responsive to changing labor demand.

While federal support will be necessary for any large-scale AI-enabled training program, much of the actual implementation should be done at the local level. This approach would allow the training programs to be tailored in a regional way (e.g., county, metro area, or state) so as to better address skill demand from local and regionally based firms. Moreover, the local matches between workers and firms would mean that workers do not need to move to new geographic locations, which is responsive to the fact that fewer workers move to new geographies than did a generation ago, as pointed out earlier.

Community colleges are particularly well positioned to serve as intermediaries in this process. They operate at the intersection of education and local labor markets, serving both new entrants to the workforce as well as displaced workers seeking new skills (Goolsbee et al. 2019). Community colleges' relationships with regional employers can help them to identify emerging skill needs and translate those needs into curriculum updates, short-term credentials, and career pathways (Lipson 2026).

Pilot studies that implement these ideas would be useful so that we can see what works and what does not and how such programs need to be tailored differently across geographies and industries. In addition, businesses already implementing such programs could be incentivized to share best practices with others in the industry and with government providers. Community colleges could provide one natural institutional home for these pilots.

If they work, these AI-enabled skill training programs would address many of the points laid out in the Skills Alignment Framework. These training programs also link back to the evidence on "shadow AI": Workers are already experimenting with these tools, but policy can help convert informal experimentation into portable, measurable, and labor-market-relevant skills. These programs would benefit businesses by providing a steady supply of well-trained workers. They also benefit workers by informing them about possible career pathways and the outcomes associated with each, and by training them in the skills businesses need in those pathways.

4.2 Scale Proven Workforce-Training Models

Workforce development programs help workers acquire skills that are in demand among employers, often through apprenticeships, sectoral partnerships, on-the-job instruction, customized training, or employer-designed programs supported in part by public funds. Employer-linked workforce development programs are a natural complement to AI-enabled training because they embed demand signals from firms directly into the design of training. Community colleges can play an important role here as well. By working with regional employers, they can translate employer demand into curriculum updates and credentials while providing workers with skills that are portable across firms (Lipson 2026). This arrangement can also help address a classic training externality: Individual firms may underinvest in general skills because workers can take those skills to competing employers, whereas publicly supported community colleges can provide general and transferable training that complements firm-specific training. These workforce development models operate in several different ways.

Apprenticeship programs pair classroom or technical instruction with paid, work-based learning at an employer, giving participants a chance to apply newly acquired skills in a real job setting while employers evaluate and train potential hires. If the program is a Registered Apprenticeship, which are vetted and approved by the US Department of Labor, the worker receives a portable, nationally recognized credential. Apprenticeships remain much less common in the US than in several peer countries, suggesting substantial room for scale-up if supported by stronger branding, occupational standards, public funding for related instruction, and better program infrastructure (Holzer 2026). Expanding apprenticeship in the US could improve school-to-career transitions, raise youth employment and wages, better align workers' skills with employer demand, and generate positive returns for both workers and firms (Lerman 2019).

Public-private training grants operate somewhat differently. In these programs, firms apply for public funds to train incumbent or newly hired workers. Such grants are designed to address a classic market failure in training. Workers may not know which skills are valuable, may not be able to finance training, or may be reluctant to invest in skills whose returns are uncertain. Firms, meanwhile, may know what skills they need but underinvest in training because trained workers can be hired away by competitors. Public funding can therefore help align training with actual labor demand while reducing the private cost to firms and workers.

In principle, these programs can help firms adapt to technological change, expand production, and fill skill gaps, while helping workers acquire more valuable and potentially portable skills that may generate spillover benefits to the local economy over time. In terms of the Skills Alignment Framework, these programs are valuable because they give firms skin in the game while giving workers more credible information about which skills employers are actually willing to reward.

While academic research on workforce development programs is limited, the existing evidence is encouraging. State workforce development training grants are more common in competitive labor markets where firms face greater risk that trained workers will be poached, consistent with the idea that public funding helps overcome underinvestment in training (Dillon et al. 2025). After receiving grants, firms grow for an extended period and reduce skill requirements in job postings, suggesting that training helps firms expand and creates more opportunities for less skilled workers (Dillon et al. 2025).

Evidence from Tennessee's customized job-training program points in a similar direction. The program raised worker earnings by about 3 percent per quarter over five years, with especially large gains in small firms and in programs emphasizing industry-specific skills (Millar 2025). Much of the training content appears to be transferable, potentially generating spillover benefits for other firms. Increased earnings also raise tax revenues, which likely more than offsets program costs (Millar 2025).

While more research is needed, the available evidence suggests that well-designed workforce training programs can benefit both workers and firms. They can help address information asymmetries, reduce underinvestment in training, and create pathways for workers to acquire skills that are valuable in a changing economy. A key policy challenge is to ensure that public funds support genuinely additional training rather than subsidizing training firms would have provided anyway. Although there is little evidence of this concern in the Tennessee customized job-training program, the risk is especially relevant for programs that provide direct payments or tax incentives to firms.

Looking ahead, these programs could be especially valuable if paired with AI-enabled skill training and better information about career pathways. Such an approach could help firms adapt to new technologies while giving workers clearer guidance about which skills are in demand, what training options are available, and what outcomes are associated with different pathways.

4.3 Modernize WIOA

The Workforce Innovation and Opportunity Act (WIOA) is a federal workforce-policy framework that helps adults, dislocated workers, and other job seekers access employment services, career counseling, and occupational training. It is intended to connect eligible workers to good jobs through job search assistance, individualized career services, and training opportunities. Its programs are typically short-term, job-focused, tied to occupations that states or local workforce boards identify as in demand, and designed to result in portable credentials, such as registered apprenticeships, commercial driver's licenses, or information technology (IT) credentials. WIOA trainees often earn higher post-training wages, but the returns are uneven and depend on destination industries and occupations (Hyman et al. 2025).

Despite these goals, WIOA's training system has well-known shortcomings. The system is highly fragmented: More than 7,000 eligible providers and roughly 75,000 eligible programs compete for a relatively small pool of training dollars, with the average program serving only a few WIOA-funded learners per year (Deming et al. 2023). This makes it difficult for participants to distinguish strong programs from weak ones and limits incentives for providers to build high-quality offerings around WIOA funding. Performance data are often missing or unreliable, and state websites, often the first point of contact for a job seeker, are frequently unhelpful (Deming et al. 2023). WIOA already has many of the institutional pieces needed for a skills alignment system, but it lacks the information infrastructure and accountability mechanisms needed for a fast-moving AI transition.

While WIOA training systems would appear to be the type of policy that benefits both workers and firms, there is need for better measurement and accountability of the third parties that provide programs, better tracking of participant outcomes, and clearer and more real-time information for workers considering programs. Modernizing WIOA should therefore be understood not only as a funding issue but as a market design problem: Workers need clearer information, providers need stronger incentives to offer valuable programs, and employers need more confidence that credentials reflect usable skills. Moving forward, these aims should determine immediate next steps to modernize and update WIOA. Longer term, if the AI-enabled skill-training pilot studies show promise, one could imagine replacing much of the WIOA training delivery infrastructure with AI-enabled skill training. If AI-enabled skill-training pilots prove successful, WIOA could become the delivery channel through which personalized skill assessment, training recommendations, and local labor-market information are scaled.

4.4 Modernize Unemployment Insurance

A top policy priority should be modernizing unemployment insurance. Modernized unemployment insurance should be viewed as both a safety net policy and a data infrastructure policy. While the data reviewed in sections 1 and 2 of this chapter do not support the idea of a rapid AI-led labor displacement, if such an event were to occur, displaced workers would rely heavily on unemployment insurance. Unemployment insurance is attractive because it is automatic relative to many targeted programs. Yet, we now know from the Covid-19 pandemic that the unemployment insurance system is outdated. Eligibility varies by state, benefit duration differs, and IT administrative systems remain outdated. During the pandemic, Congress created temporary programs such as Pandemic Unemployment Assistance to fill gaps and cover contractors, gig workers, and many part-time workers who are normally excluded from regular unemployment insurance.

Those temporary programs helped many workers but also proved vulnerable to fraud and administrative failure. The Government Accountability Office (GAO) documented large improper payments in pandemic unemployment insurance programs and has suggested ways for the Department of Labor to address these issues (GAO 2023). One important lesson from the pandemic experience is that the Department of Labor and state agencies that administer unemployment insurance are in desperate need of a technical and human capital overhaul. As of 2020, less than half of states had modernized their unemployment insurance IT systems. In some cases, workers still need to be physically mailed a password to log in to their unemployment insurance accounts (Rouse 2024).

A modernization agenda should consider both expansion of eligibility and upgrades to unemployment insurance systems. On the expansion of eligibility, eligibility rules could be expanded to better cover part-time workers, workers with variable hours, and workers with mixed W-2 and 1099 income. As indicated earlier, an increasing portion of the workforce relies on such earnings (Sundararajan 2025). On the second point, states should upgrade IT systems so claims can be processed quickly, securely, and flexibly, which will benefit workers adjusting to a changing labor market landscape. These updated unemployment insurance systems could, in principle, help displaced workers connect more effectively to reemployment services, training information, and job matching programs, all of which are better for the worker and for the firms searching for workers with specific skill sets.

In principle, better IT systems for unemployment insurance would allow policymakers to monitor claims by occupation, industry, geography, age, and other characteristics in near-real time, ensuring that policymakers are equipped with the data they need to make labor policy adjustments. For example, California Governor Newsom's recent executive order on AI and the workforce explicitly calls on the California Employment Department to use data from the state's unemployment insurance program to study the effect of AI on labor (Newsom 2026). An example of such work is shown in recent research by the California Policy Lab that links California unemployment insurance claims with AI exposure measures (Hyman et al. 2026). Updated systems would also help combat fraud. All these points connect directly to the Skills Alignment Framework: Better unemployment insurance systems can reduce the cost of transition for workers while generating the information needed to target training and reemployment services more effectively.

4.5 Expand Wage Insurance

Wage insurance is a form of insurance for displaced workers who find a new job that pays less than the job they lost. Rather than replacing income only while a worker is unemployed, as unemployment insurance does, wage insurance continues to support the worker after reemployment by paying a temporary subsidy tied to the size of the wage loss. Conceptually, the policy is meant to reduce the "cliff" workers face when accepting a lower-paying job, making faster reemployment more attractive while still preserving incentives to work. Research shows that wage insurance can smooth income losses, shorten unemployment spells, help workers avoid the scarring effects of long-term joblessness, and provide a bridge while workers rebuild earnings in a new job or sector (Boeri and Cahuc 2023). Policymakers including President Obama have highlighted the potential benefits of wage insurance (Obama 2016). Wage insurance complements but does not substitute for skill alignment policies.

The academic literature on wage insurance remains limited because wage insurance programs have been rare, small, and often bundled with other labor market policies. Hence, it is difficult to generalize from the limited evidence on small, targeted wage insurance programs to the US economy as a whole. A recent US study of wage insurance under Trade Adjustment Assistance (TAA) provides the strongest evidence to date. Using US Census data that link employers and employees, one study finds that wage insurance eligibility increased short-run employment and raised cumulative earnings over the following four years, mainly by shortening nonemployment spells (Hyman et al. 2024).

Wage insurance is a policy that should be beneficial for both workers and firms. A worker may be more willing to accept a job at a firm that cannot match their old wage, while the firm does not have to bear the full cost of compensating the worker's earnings loss. Moreover, the firm may benefit from an expanded pool of potential workers to hire from. Wage insurance reduces the financial penalty workers face when moving into a new occupation, firm, or sector, thereby making adjustment faster and less risky.

However, while wage insurance gives workers more financial room to accept a lower-paid job in a new firm or field, it does not necessarily help reduce information asymmetries about which skills are in demand. At a minimum, it would be useful to experiment more with wage insurance programs to gather additional data on how they may or may not be helpful during periods of rapid job transformation. However, because wage insurance does not by itself tell workers which jobs or skills are likely to pay off, it should be paired with AI-enabled skill assessment and local labor-market information.

4.6 Rely More Heavily on Capital Taxation

Tax policy addresses a different margin than training policy: It affects firms' incentives to substitute capital for labor, rather than directly helping workers acquire new skills. The US tax system currently favors investment in capital over labor, which may distort a firm's choice between workers and machines. Labor income is generally subject to individual income taxes and payroll taxes, while capital income often benefits from lower rates, deferral, accelerated depreciation, expensing, and other preferences. This wedge can make investment in capital, such as AI and robots, privately attractive to firms even when the social-productivity gain from replacing workers may be small. Effective US labor taxes are estimated to exceed 25 percent, while effective capital taxes on software and equipment were about 10 percent in the 2010s and fell to about 5 percent after the 2017 tax reforms (Acemoglu et al. 2020). According to some accounts, reducing the capital bias in the tax code could increase employment (Acemoglu et al. 2020).

There are several ways to pursue these changes, each with tradeoffs. One set of approaches is to reduce taxes on labor, especially payroll taxes or employer-side employment taxes, which would lower the tax penalty on hiring and could increase after-tax wages. The challenge is that the tax revenue that is used to finance Social Security, Medicare, and other spending would drop. A second set of strategies is to raise the effective tax burden on capital income via a variety of strategies, such as limiting accelerated depreciation and expensing, narrowing preferential treatment for capital gains and dividends, strengthening taxation of business and pass-through income, and preserving or expanding taxes such as the net investment income tax. One reason to focus on expensing is that allowing firms to immediately deduct the full cost of capital investments can effectively exempt the normal return to those investments from tax, thereby further privileging capital over labor in a system where wages remain subject to individual income and payroll taxes (Batchelder 2026). These effective tax increases on capital would raise revenue and reduce incentives to use capital in place of labor but could discourage some investment and face political and design constraints. A deeper discussion of the fiscal consequences of changes to taxes on capital are provided in this volume by Hubbard and Pardue (2026) and Clausing and Sarin (2026).

Adjusting tax policy in the ways described above would likely benefit labor and could be desirable even if it imposes some costs on firms. Changes in tax policy would likely slow adoption of some new technologies, which, while not a policy goal of this chapter, could provide some breathing room for policymakers to put in place appropriate training and retraining programs. The slower adoption of some new technologies might also mean fewer benefits from the technologies in the short term, and these tradeoffs would need to be weighed against the benefits of a less capital-biased tax system. There are substantial political constraints around changing tax policy, but raising taxes on capital while lowering taxes on labor could potentially have large benefits for workers. If policymakers reduce the tax bias toward automation, they may slow some inefficient labor-replacing investment, but the broader challenge remains to help workers and firms identify and build the skills that will be valuable as AI adoption proceeds.

4.7 Avoid Narrowly Targeted Labor-Market Policies

The evidence reviewed earlier cautions against policies that require policymakers to determine whether a particular job loss was "caused by AI." AI exposure is measurable only imperfectly, adoption is often informal, and displacement may occur through task reorganization rather than through easily observed layoffs. A variety of "targeted" policies have been suggested by policymakers and scholars, with the idea that government intervention is necessary to either provide assistance to certain populations or to slow down adoption of certain technologies. In general, while these policies may be well-intentioned, their implementation—in particular the definition of the "target"—is politically fraught, and hence they should be avoided.

For instance, policymakers should *not* pursue an AI-specific assistance program. Such an assistance policy would step in when a worker loses a job to AI, offering support as that worker looks for a new job. AI-specific assistance programs are similar in spirit to the US Trade Adjustment Assistance (TAA) program, which existed for more than 60 years in various forms. The appeal of a TAA-like program for AI is clear: If workers are displaced by a national economic transformation that generates aggregate gains at the expense of specific workers, policy should help them adjust. The problem with a TAA-like program for AI is eligibility and take-up. TAA has long struggled to identify which workers qualify, and AI would make these problems harder. Firms may have incentives to overstate or understate AI's role. Workers may not know whether AI caused their displacement. Administrative agencies would be asked to determine causality in cases where measurement is likely difficult. These challenges are why the Skills Alignment Framework favors broad adjustment capacity over AI-specific eligibility categories. The policy goal should be to help workers and firms respond to changing skill demand, regardless of whether the proximate cause is AI, remote work, trade, interest rates, or some combination of forces.

Another idea to be avoided is an "AI tax" or a "robot tax," which would place a levy on firms that substitute machines or software-based automation for human labor. Its proponents view automation as a source of employment displacement and fiscal erosion: When AI or robots replace workers, payroll tax revenues fall, displaced workers require support, and firms may internalize too little of the social cost of automation. In its simplest form, the tax would raise the relative price of adopting AI or robots, thereby slowing labor-replacing investment, preserving some employment, and generating revenue that could finance retraining, wage insurance, or other transition assistance for affected workers. More broadly, the proposal reflects a concern that the tax system may already favor capital over labor, encouraging inefficiently high automation relative to the social optimum, a point discussed above.

The case against a targeted AI or robot tax is that it targets the wrong margin with a poorly defined instrument. The empirical evidence suggests that robot-adopting firms increase rather than decrease employment and often improve productivity as well (Seamans 2021). A robot tax would therefore risk penalizing productive investment, slowing diffusion of efficiency-enhancing technologies, and reducing growth, with potentially adverse effects on workers as well as consumers. As highlighted above, the empirical evidence on AI and labor is still too nascent to arrive at any conclusions about effects on employment and productivity. While in some cases AI may substitute for human labor, in other cases AI may augment human labor. An AI tax would penalize both, and this is a problem if it turns out there are more cases of AI augmentation.

4.8 Improve Data, Research, and Measurement Infrastructure

Data infrastructure is the backbone of the entire policy framework. An overarching concern with all the policies considered is that we do not know if we are spending money wisely if we cannot measure outcomes. Hence, a top priority should be to ensure that the US statistical agencies are well-funded and well-supported. The American Economic Association (AEA) Committee on Economic Statistics (AEA 2025) and others have made this call repeatedly. As described above, the Census Bureau's Annual Business Survey and Business Trends and Outlook Survey are currently the benchmark of what we know about AI adoption across the US economy, but budget restrictions or reductions could put these fundamentally important data collections at risk. Continued federal support is not just a matter of keeping existing surveys alive; it also means ensuring that AI-related questions are asked consistently over time, so that researchers can track adoption curves and link them to changes in employment, wages, and productivity. Without timely, granular, and credible information, policymakers cannot know where AI adoption is occurring, which workers are affected, which training programs work, or whether public dollars are improving labor market matching.

In addition, more efforts should be made to incorporate private companies into the national statistical infrastructure. Private companies hold vast quantities of data that could substantially accelerate research on AI's labor market effects. Finding ways to unlock that data without compromising proprietary interests is one of the most tractable near-term opportunities. One idea is to encourage and/or incentivize externally supported academic researchers to work directly with company counterparts to build partially processed datasets that subsequent research teams could use as starting points. This strategy should help speed up research on the economic effects of AI. An example of such an approach can be found in the work of Howell et al. (2026), who use licensed vendor data (Lightcast) to construct a county-year panel of skill specialization, relatedness, and complexity from 433.6 million job postings (2010–2024). Their approach illustrates how licensed vendor microdata can be transformed into an openly available research and policy resource.

Another idea is to establish a secure research center, potentially organized by a body like the American Economic Association, National Association of Business Economists, or National Bureau of Economic Research, that could house firm-supplied datasets under well-defined data security provisions, allowing multiple research teams to study the same data while giving firms the assurances they need about confidentiality and intellectual-property protection. Such an approach would be similar in spirit to the restricted data that the Census makes available in its Research Data Centers for approved projects and has antecedents in the university sector's data-sharing enclaves such as UMETRICS. This subsection therefore closes the loop with the chapter's opening argument: Because AI's labor market effects remain uncertain and uneven, the first requirement of policy readiness is the capacity to observe change as it happens.

Conclusion

This chapter reviews the limited evidence on the labor market impacts of AI and proposes that potential labor-market policies focus their attention on (1) addressing information asymmetry about which skills will be valuable as AI reshapes tasks within and across occupations and (2) ensuring that policy serves the aims of both workers and employers simultaneously.

The evidence paints a picture of a technology that is diffusing rapidly but unevenly, and whose labor market effects are still difficult to measure clearly. Firm-level adoption of generative AI has risen quickly, reaching roughly 18 percent of US firms by early 2026, or about 32 percent on an employment-weighted basis, but for now remains concentrated in larger, more digitized firms and in sectors like Information. Productivity studies consistently show large gains from AI use, especially for less experienced workers, but these task-level gains have not yet translated into measurable aggregate effects, either in terms of economic growth or labor market disruption. On balance, the labor market picture is one of a technology in early but accelerating diffusion, whose full economic effects will depend heavily on organizational adoption decisions, complementary investments, and the policy environment. The evidence to date suggests that the United States does not yet face broad AI-driven labor-market disruption, but it does face a growing need for better skill discovery, more adaptive training, stronger transition supports, and more timely data.

Looking ahead, the most promising near-term policy combines several reinforcing elements. AI itself can be a tool for workforce policy: Given that generative AI disproportionately raises productivity for less experienced workers, AI-enabled skill training programs, which use machine learning to match workers to feasible career pathways and deliver personalized instruction, could prove both efficient and scalable. In the near term, AI-enabled skill training programs should be piloted at regional levels before being scaled up in the longer term if pilots prove successful. AI training can be coupled with workforce development programs that align training with actual employer skill demand, ideally pairing public funding with accountability mechanisms that give workers reliable information about program quality and expected outcomes. Other programs, such as those funded by the Workforce Innovation and Opportunity Act, should be modernized in the near term and potentially folded into AI training programs in the longer term.

The central argument of this chapter is that AI does not require an entirely new labor policy architecture. Instead, it requires improving the process through which information about changing skill demand flows from employers to workers and training providers. Better information, stronger employer engagement, and more responsive training systems can reduce mismatches before they become costly, allowing workers and firms to adapt together and mutually benefit as AI continues to evolve.

References

Acemoglu, Daron, Andrea Manera, and Pascual Restrepo. 2020, April. "Does the US Tax Code Favor Automation?" Working paper no. 27052. National Bureau of Economic Research. https://www.nber.org/papers/w27052.

Allen, Jeffrey S. 2026, April 3. "Monitoring AI Adoption in the U.S. Economy." *FEDS Notes*, Board of Governors of the Federal Reserve System. https://www.federalreserve.gov/econres/notes/feds-notes/monitoring-ai-adoption-in-the-u-s-economy-20260403.html.

American Economic Association (AEA) Committee on Economic Statistics. 2025, February 21. "Measuring the Economic Effects of Artificial Intelligence." AEA. https://www.aeaweb.org/content/file?id=22228.

Appel, Ruth, Peter McCrory, Alex Tamkin, Miles McCain, Tyler Neylon, and Michael Stern. 2025, September 15. *The Anthropic Economic Index Report: Uneven Geographic and Enterprise AI Adoption*. Anthropic. https://www-cdn.anthropic.com/7b76335c444876a93fa22a63aabb4aeb820aff25.pdf.

Autor, David, Caroline Chin, Anna Salomons, and Bryan Seegmiller. 2024. "New Frontiers: The Origins and Content of New Work, 1940–2018." *Quarterly Journal of Economics* 139, no. 3: 1399–465.

Batchelder, Lily L. 2026, Summer. "Expensing—Not Worth the Expense." *Democracy Journal* 81. https://democracyjournal.org/magazine/81/expensing-not-worth-the-expense/.

Bick, Alexander, Adam Blandin, and David J. Deming. 2026, January 20. "The Rapid Adoption of Generative AI." *Management Science*. https://doi.org/10.1287/mnsc.2025.02523.

Blair, Peter Q., Tomas G. Castagnino, Erica L. Groshen, Papia Debroy, Byron Auguste, Shad Ahmed, et al. 2020, March. "Searching for STARs: Work Experience as a Job Market Signal for Workers Without Bachelor's Degrees." Working paper no. 26844. National Bureau of Economic Research. https://www.nber.org/papers/w26844.

Boeri, Tito, and Pierre Cahuc. 2023. "Labor Market Insurance Policies in the Twenty-First Century." *Annual Review of Economics* 15: 1–22.

Bonney, Kathryn, Cory L. Breaux, Emin Dinlersoz, Lucia S. Foster, John C. Haltiwanger, and Aditya A. Pande. 2026, April. "The Microstructure of AI Diffusion: Evidence from Firms, Business Functions, and Worker Tasks." Working paper no. 35141. National Bureau of Economic Research. https://www.nber.org/papers/w35141.

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen. 2025, November 13. "Canaries in the Coal Mine? Six Facts About the Recent Employment Effects of Artificial Intelligence." Working paper. Stanford University. https://digitaleconomy.stanford.edu/app/uploads/2025/12/CanariesintheCoalMine_Nov25.pdf.

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond. 2025. "Generative AI at Work." *Quarterly Journal of Economics* 140, no. 2: 889–942.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2021. "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies." *American Economic Journal: Macroeconomics* 13, no. 1: 333–72.

Cai, Junhui Jeff, Xian Gu, Liugang Sheng, Mengjia Xia, Linda Zhao, and Wu Zhu. 2025. "AI as 'Co-Founder': GenAI for Entrepreneurship." *arXiv* preprint. https://doi.org/10.48550/arXiv.2512.06506.

Challenger, Gray, and Christmas. 2026, April 2. *Job Cut Announcement Report*. Challenger, Gray, and Christmas. https://www.challengergray.com/wp-content/uploads/2026/04/Challenger-Report-March-2026-1.pdf.

Chatterji, Aaron, Thomas Cunningham, David J. Deming, Zoe Hitzig, Christopher Ong, Carl Yan Shan, and Kevin Wadman. 2025, September. "How People Use Chatgpt." Working paper no. 34255. National Bureau of Economic Research. https://www.nber.org/papers/w34255.

Clausing, Kimberly, and Natasha Sarin. 2026. "$15 Trillion on the Table in 2029: Ordinary Taxes to Meet an Extraordinary Fiscal Moment." *The American Economy in a New Era*, edited by Melissa S. Kearney and Luke Pardue. Aspen Institute Press.

Council of Economic Advisers (CEA). 2026, January 21. *Artificial Intelligence and the Great Divergence*. White House. https://www.whitehouse.gov/research/2026/01/artificial-intelligence-and-the-great-divergence/.

Crane, Leland Dod, and Paul E. Soto. 2026. "AI and Coder Employment: Compiling the Evidence." Working paper. Board of Governors of the Federal Reserve System. https://www.federalreserve.gov/econres/feds/files/2026018pap.pdf.

Darling, Matt. 2026, March 11. "What the Challenger Layoff Tracker Really Measures." Substack. https://besttrousers.substack.com/p/what-the-challenger-layoff-tracker.

Dave, Paresh. 2026, February 9. "No Company Has Admitted to Replacing Workers with AI in New York." *Wired*. https://www.wired.com/story/no-company-has-admitted-to-replacing-workers-with-ai-in-new-york/.

David, Paul A. 1990. “The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox.” *American Economic Review* 80, no. 2: 355–61.

Dell’Acqua, Fabrizio, Edward McFowland III, Ethan Mollick, Hila Lifshitz, Katherine C. Kellogg, Saran Rajendran, et al. 2026. “Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality.” *Organization Science* 37, no. 2: 403–23.

Deming, David, Alexis Gable, Rachel Lipson, and Arkadijs Zvaigzne. 2023, March. *Navigating Public Job Training*. Harvard Project on Workforce. https://pw.hks.harvard.edu/post/publicjobtraining.

Deming, David, Christopher Ong, and Lawrence H. Summers. 2024. “Technological Disruption in the US Labor Market.” In *Strengthening America’s Economic Dynamism*, edited by Melissa S. Kearney and Luke Pardue., 169–196. Aspen Institute. https://doi.org/10.5281/zenodo.13973975.

Dillon, Eleanor W., Lisa B. Kahn, Joanna Venator, and Michael Dalton. 2025, July. “Do Workforce Development Programs Bridge the Skills Gap?” Working paper no. 34012. National Bureau of Economic Research. https://www.nber.org/papers/w34012.

Eckhardt, Sarah, and Nathan Goldschlag. 2025, August 10. *AI and Jobs: The Final Word (Until the Next One)*. Economic Innovation Group. https://eig.org/ai-and-jobs-the-final-word/.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2024. “GPTs Are GPTs: Labor Market Impact Potential of LLMs.” *Science* 384, no. 6702: 1306–8.

Emanuel, Natalia, Emma Harrington, and Amanda Pallais. 2026, June 1. “Remote Work Leaves Younger Workers Sidelined.” Liberty Street Economics. https://doi.org/10.59576/lse.20260601.

Felten, Edward, Manav Raj, and Robert Seamans. 2021. “Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses.” *Strategic Management Journal* 42, no. 12: 2195–217.

Felten, Ed, Manav Raj, and Rob Seamans. 2024. “Generative AI Requires Broad Labor Policy Considerations.” *Communications of the ACM* 67, no. 8: 29–32.

Furman, Jason, and Robert Seamans. 2019. “AI and the Economy.” Innovation Policy and the Economy 19, no. 1: 161–91.

Gimbel, Martha, Joshua Kendall, and Ryan Kulsakdinun. 2026, February 19. “Labor Market AI Exposure: What Do We Know.” Budget Lab, Yale University. https://budgetlab.yale.edu/research/labor-market-ai-exposure-what-do-we-know.

Goolsbee, Austan, R. Glenn Hubbard, Melissa S. Kearney, and Amy Ganz. 2019. “A Policy Agenda to Develop Human Capital for the Modern Economy.” In *Expanding Economic Opportunity for More Americans*, edited by Melissa S. Kearney and Amy Ganz, 16–39. Aspen Institute Press.

Government Accountability Office (GAO). 2023, September 12. *Unemployment Insurance: Estimated Amount of Fraud During Pandemic Likely Between $100 Billion and $135 Billion*. GAO-23-106696. Government Accountability Office. https://www.gao.gov/assets/gao-23-106696.pdf.

Holzer, Harry J. 2026. “Workforce Development in the United States: Recent Trends and Evidence.” *Journal of Economic Literature* 64, no. 2: 602–36.

Howell, Anthony, Maryann Feldman, Lauren Lanahan, Nikhil Kalathil, and Evan Johnson. 2026, June 6. “An Economic Geography Dataset of U.S. Skill Specialization, Relatedness, and Complexity.” *arXiv*. https://arxiv.org/abs/2606.09918.

Hubbard, Glenn, and Luke Pardue. 2026. “Advancing Economic Participation in the AI Era: A Private- and Public-Sector Agenda.” *The American Economy in a New Era*, edited by Melissa S. Kearney and Luke Pardue. Aspen Institute Press.

Humlum, Anders, and Emilie Vestergaard. 2025, May. “Large Language Models, Small Labor Market Effects.” Working paper no. 33777. National Bureau of Economic Research. https://www.nber.org/system/files/working_papers/w33777/revisions/w33777.rev0.pdf.

Hyman, Benjamin G., Brian K. Kovak, and Adam Leive. 2024, May. “Wage Insurance for Displaced Workers.” Working paper no. 32464. National Bureau of Economic Research. https://www.nber.org/papers/w32464.

Hyman, Benjamin G., Benjamin Lahey, Karen Ni, and Laura Pilossoph. 2025, August. “How Retrainable Are AI-Exposed Workers?” Working paper no. 34174. National Bureau of Economic Research. https://www.nber.org/papers/w34174.

Hyman, Ben, Till Von Wachter, Diana Herrera, Roozbeh Moghadam, Jacob Morris, Swapnil Motghare, et al. 2026, June. *Tracking AI-Related Job Loss Using Unemployment Insurance Claims Data in California*. California Policy Lab. https://capolicylab.org/wp-content/uploads/2026/06/Tracking-AI-Related-Job-Loss-Using-Unemployment-Insurance-Claims-Data-in-California.pdf.

Iscenko, Zanna, and Fabien Curto Millet. 2026, January. “Looking for the Ladder: Is AI Impacting Entry-Level Jobs?” Economic Innovation Group. https://eig.org/wp-content/uploads/2026/01/TAWP-Iscenko-Millet.pdf.

Jones, John, Yue Li, and Urvi Neelakantan. 2025. “Why Is Geographic Mobility Declining?” *Economic Brief* 25, no. 25-19.

Kharazian, Ara, Lisa Simon, and Ryan Stevens. 2026, June. *A New Look at AI’s Impact on Jobs: Firm-Level AI Spending and Workforce Adjustment*. Ramp Economics Lab. https://ramp.com/data/ai-jobs-impact.

Kim, Hyunjin, Dahyeon Kim, and Rembrand Koning. 2026. “Mapping AI into Production: A Field Experiment on Firm Performance.” Working paper no. 2026/20/STR. INSEAD. https://ssrn.com/abstract=6513481.

Klein, Michael, and Maurice Obstfeld. 2023, June 12. “Why Is the U.S. Debt Expected to Keep Growing?” *EconoFact*. https://econofact.org/why-is-the-u-s-debt-expected-to-keep-growing.

Lambert, Peter John, and Yannick Schindler. 2026, May 20. “The Broken Ladder: AI, Remote Work, and Early-Career Hiring.” Social Science Research Network. https://ssrn.com/abstract=6787638.

Lerman, R. I. 2019. “Scaling Apprenticeship to Increase Human Capital.” In *Expanding Economic Opportunity for More Americans*, edited by Melissa S. Kearney and Amy Ganz, 56–75. Aspen Institute.

LinkedIn. 2026. *Labor Market Report: Building a Future of Work That Works*. Economic Graph Research Institute. https://economicgraph.linkedin.com/content/dam/me/economicgraph/en-us/PDF/linkedIn-labor-market-report-building-a-future-of-work-that-works-jan-2026.pdf.

Lipson, Rachel. 2026, July. *Getting in the Front Door in the Age of AI: Policies to Help Early Career Workers*. IE University Center for the Governance of Change. https://static.ie.edu/CGC/CGC_FrontDoorAI_Lipson.pdf.

McElheran, Kristina, J. Frank Li, Erik Brynjolfsson, Zachary Kroff, Emin Dinlersoz, Lucia Foster, and Nikolas Zolas. 2024. “AI Adoption in America: Who, What, and Where.” *Journal of Economics and Management Strategy* 33, no. 2: 375–415.

Millar, N. 2025. “The Labor Market Returns to Customized Job Training.” Working paper. Stanford University. https://drive.google.com/file/d/18RjxrsrIniaNda_fmCSw5j60immjdgoR/view

Misra, Amit, Jane Wang, Will Fein, Scott, McCullers, Kevin White, and Juan Lavista Ferres. 2026, May 27. *Measuring AI Diffusion Across U.S. Geographies: County, State, and Metro Estimates*. Microsoft AI for Good Lab. https://www.microsoft.com/en-us/research/wp-content/uploads/2026/05/AI-Diffusion-US-Technical-Report.pdf.

Newsom, Gavin. 2026. Executive Order N-6-26. https://www.gov.ca.gov/wp-content/uploads/2026/05/5.21.26-AI-Workforce-EO-FINAL-SIGNED.pdf.

Noy, Shakked, and Whitney Zhang. 2023. "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence." *Science* 381, no. 6654: 187–92.

Obama, Barack. 2016, January 12. "Remarks by the President in the State of the Union Address, January 12, 2016." Obama Library. https://www.obamalibrary.gov/sites/default/files/uploads/documents/State%20of%20the%20Union%20Address%202016%20(TRANSCRIPT).pdf.

Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer. 2023. "The Impact of AI on Developer Productivity: Evidence from Github Copilot." *arXiv* preprint. https://doi.org/10.48550/arXiv.2302.06590.

Qin, Fei, and Thomas Kochan. 2020, December 3. *The Learning System at IBM: A Case Study*. MIT. https://mitsloan.mit.edu/sites/default/files/2022-06/Qin%20and%20Kochan%20The%20Learning%20System%20at%20IBM%2012%20 2020.pdf.

Ross, Martha, and Mark Muro. 2024, January 4. "How Federal, State, and Local Leaders Can Leverage the CHIPS and Science Act as a Landmark Workforce Opportunity." Brookings Institution. https://www.brookings.edu/articles/how-federal-state-and-local-leaders-can-leverage-the-chips-and-science-act-as-a-landmark-workforce-opportunity/.

Rouse, Cecilia E. 2024. "Lessons for Economists from the Pandemic." *NBER Reporter* 3: 1–9.

Seamans, Robert. 2021, August 25. "Tax Not the Robots." Brookings. https://www.brookings.edu/articles/tax-not-the-robots/.

Sundararajan, Arun. 2025, April 22. *Workforce Capacity Development and Occupational Transitions with Dignity*. Brookings Institution. https://www.brookings.edu/articles/workforce-capacity-development-and-occupational-transitions-with-dignity/.

Tedeschi, Ernie, Marisa Rama, and Chris Cruickshank. 2026, June 22. "The Age of the Solopreneur." Stripe Economics. https://www.stripeeconomics.com/p/the-age-of-the-solopreneur.

Webb, Michael. 2019. “The Impact of Artificial Intelligence on the Labor Market.” Social Science Research Network. https://dx.doi.org/10.2139/ssrn.3482150.

Yale Budget Lab. 2026, updated August 19. “Tracking the Impact of AI on the Labor Market.” Budget Lab, Yale University. https://budgetlab.yale.edu/research/tracking-impact-ai-labor-market.