

AI ADOPTION, UPSKILLING, AND WORKFORCE TRANSFORMATION IN THE GLOBAL ECONOMY

COMPANION PAPER: **THE AI LABOR STACK**

May 18th, 2026

Dustin Ferrone^{1,2}, Lindsey Lacey¹, Christophe Combemale^{1,2}

Carnegie
Mellon
University



¹ Valdos Consulting, LLC

² Carnegie Mellon University

CONTENTS

Executive Summary	3
Introduction	5
Literature Review	6
I. Training and education	6
II. Adoption of AI	8
III. Labor market impacts	9
IV. International Perspective	11
Pathways from Productivity Gains to Labor Impact	14
Labor Conditions For AI Adoption	15
Categories of AI Workers	15
Policy & Open Research Questions	18
Adoption	18
Workforce Development	20
New Industries and New Tasks	21
Labor Strategies for Global Trade	22
Conclusion	23
Workshop Questions	24
References	25

EXECUTIVE SUMMARY

The diffusion of artificial intelligence (AI) is unfolding as one of the fastest technological transformations in history: measured on utilization by consumers and the workforce, the pace of adoption may be 2.8 times faster than the personal computer (PC) and twice as fast as the internet (Bick, Blandin, and Deming, 2024). AI differs from previous technological advancements in its ability to automate a wide range of high-skill, less-routine cognitive tasks, with a high tolerance for input variability (compared to past information technologies, which primarily automated routine, middle-skill tasks).

Despite ongoing improvements in AI capabilities, there is lingering uncertainty about the labor market impacts of AI, how it will diffuse across the global economy, and how nations can position themselves to participate in global productivity gains.

AI is already showing strong evidence of being a General Purpose Technology (GPT): historical GPTs (e.g., the steam engine, electricity, and the microprocessor) are foundational innovations that spur complementary advances across industries, reshaping productivity, business models, and the workforce. In these previous cases, diffusion depended on highly skilled *innovators* who pushed the technology's capability frontier forward; *facilitators* who deployed, integrated, and otherwise enabled its use (often through complementary technologies); and *end-users* whose work was altered by the technology. AI diffusion may follow a similar pattern, in which each level of this labor pyramid must be satisfied before its impacts can be fully realized.

Current literature on the workforce impact of AI often fails to account for firms' economic incentives to adopt technology. In other words, what is technically feasible may not be economical. Without plausible users, there is

no opportunity for productivity gains. Without facilitators, the technology will face insurmountable logistical challenges. Without innovators, a nation faces a make-or-buy dilemma in which developing domestic expertise is a slow and uncertain process, and failure to do so raises concerns about technological dependence. This parallels the Cloud Stack in that end users primarily interact with commercial-off-the-shelf applications built by more-savvy developers who leverage available infrastructure and tooling developed by highly skilled innovators (AI Adoption Initiative [AIAI], 2025).




Figure 1: The GPT Labor Pyramid

Programs such as Amazon's Future Ready initiative seek to bolster a global workforce corresponding closely to the "facilitator" layer of the labor pyramid. Associated education programs reflect several roles essential to making AI adoption technically feasible at the regional level (e.g., Data Center Operations and Fiber-Optic Technicians, and IT workers), as well as more specialized roles that will support emerging AI applications (e.g., industrial

automation). Once a critical mass of workers gain the skills to support deployment of a general technology, their region can become a magnet for investments far beyond the initial industrial scope.

For most of human history, access to technical expertise was constrained by geography. This constraint changed significantly in the 1980s as advances in telecommunications led to a more globalized service economy. These changes have enabled nations to adopt new technologies despite a lack of domestic expertise, and to build new industries on top of incumbent industries that may not exist domestically. Countries with low labor costs may continue to adopt AI, further improving the competitiveness of their service sectors. However, with new and more capable AI models being released

regularly, the international exchange of services faces an uncertain future: the service tasks that have been offshored globally in recent decades (especially from higher- to lower-income countries) were often those that were easier to separate from other business processes. These are the same operational characteristics that facilitate AI automation. The result is that economies built on exporting services may need to compete on skills beyond the current capabilities of AI, while simultaneously adapting to changing customer business models that integrate AI. Firms and industries that have historically relied on low-cost labor to perform simple tasks for customers may be directly replaceable with AI, if the delivery of services has already been intermediated digitally. At lower volumes of work, the low-labor-cost company remains competitive. However, at scale, a large customer support sector might quickly be dominated by a few AI-powered firms that can deliver the same services at lower cost.

At the forefront of the policy discussion on the workforce impact of AI is whether the technology will directly replace workers (through automation) or be used as a tool, augmenting how jobs are performed and increasing individual worker productivity. Critically, whether the employment of augmented workers increases or decreases depends on whether the demand for goods and services increases with productivity (i.e., due to reduced prices) – if not, then not only will absolute labor demand for an augmented occupation fall, but wages may decline as disrupted workers compete for fewer positions. This relationship between demand and the employment outcomes from augmentation suggests that policy interventions in economies with high augmentation potential should emphasize growth to maintain employment. Trading growth for apparent labor market protection through regulation may lose both in the long run in a competitive global environment.

Much of the global discussion around AI labor impacts is focused on the notion of “exposure” of workers whose tasks can be performed by (or using) AI. The workers exposed to AI may be most at risk for disruption *and* most important to capturing the productivity gains from AI: hence, workforce policy to accelerate AI adoption and diffusion may be inextricable from policy for mitigating employment effects. Disruptions to one occupation may affect comparable work for the same reasons, leaving

vulnerable workers stranded with skills that are no longer in demand. Policies that encourage rapid reskilling and align firms’ incentives to recognize and support transitions for disrupted workers can help reconnect workers more quickly to the AI-user layer and other job growth areas that are less exposed. Because exposure is correlated within labor markets (the jobs that exposed workers would typically switch to are themselves exposed), governments may need to provide pathways for workers to find (or build toward) “next-best” options that are decoupled from the same technological risks affected incumbent occupations.

We identify three strategic questions for policymakers to answer about their own country and four sets of open research questions to help anticipate the global impact of AI diffusion, drawing on the labor pyramid framework and prior literature.

Questions for Policy Makers:

1. Does the national economy have access to the workers necessary to facilitate the use of AI?
2. What share of the national economy is truly “AI exposed”?
3. How do a nation’s labor costs compare to those of the global economy?

The answers to these questions lead to vastly different national strategies for AI adoption and workforce development. The first question establishes the feasible scope for an AI adoption strategy, and the second question informs the potential gains to be had from AI. A country whose national economy is highly exposed to AI and whose exposed elements are traded globally has the most urgent case for AI adoption and augmentation to continue competing in global exports. A country with low exposure and low access to facilitators may be insulated from direct disruptions and might position itself to trade goods and services from less exposed industries to countries that have successfully adopted AI. The final question concerns the extent to which the gains from AI may be driven by exports (of augmented, low-cost workers’ services) versus improvements in domestic productivity.

Open Research Questions:

1. Where and how can firms and governments upskill and reskill workers, and which workers should be targeted, to mitigate labor market disruptions and accelerate productivity gains from AI?
2. How does the rate of AI adoption vary across tiers of the labor pyramid, and what individual- and firm-level factors predict early adoption?
3. What are the net implications of AI for jobs, and where are the gains and losses concentrated?
4. How much does AI mirror or differ from other economic substitution choices, such as offshoring or computerization, to understand how AI will reshape global comparative advantage?

The answers to these questions may vary depending on national context, and they matter both for setting expectations about the gains (and disruptions) from AI adoption, and for informing the policy implications of

our three questions for policymakers. If AI follows the same economic and operational channels for diffusion as global trade in goods and services, it may drive significant competitive reorientation away from low labor cost service export economies and increase the competitiveness of a domestic AI base.

Re-skilling is critical, both for addressing worker disruption and for building the pipeline of facilitators and skilled AI end-users that the economy now demands. Past work has shown that occupational exposure to AI occurs in neighborhoods within the labor market, creating both the risk of “islands” of disrupted workers and the potential for productivity gains to be bottlenecked by shortages in a critical skillset. Yet the exact pathway is difficult to discern, because the pace and trajectory of the technology itself is unknown – we argue that policymakers operating under deep technical uncertainty may improve worker outcomes by decreasing barriers to retraining and adopting strategies that empower industries and businesses to do more with more, not the same with less.

INTRODUCTION

AI has historically proven costly to deploy in isolation. Prior to the broad commercial availability of generative models, natural language processing algorithms were almost exclusively trained and deployed by experts to address specific tasks (e.g., sentiment analysis, translation, spelling, and grammar correction). This strategy is costly to firms, as these workers demanded some of the highest wages in the U.S. labor market. Despite this, LLMs have seen rapid international adoption. This phenomenon was made possible by innovators leveraging cloud-computing infrastructure to make their general-purpose models accessible at scale. Once standard access protocols were established, facilitators developed user-friendly applications, increasing the popularity of the Software-as-a-Service business model. As tooling became more accessible (including the growth of low-code or no-code environments), this process accelerated.

Not all countries and regions participated equally in this explosion in the use and diversity of AI applications. Evidence from “superstar” cities, such as San Francisco and San Jose, suggests that local talent pools may play an important role in the diffusion of AI. Both cities have above-average employment in computer and mathematical occupations, which make up a large portion of the facilitators and innovators layers of the AI labor pyramid (Figure 1). This highlights an area where governments may have leverage (with the ability to subsidize training, regulate, or act as market makers), but also face great uncertainty. Unlocking productivity gains from AI will depend not only on adoption at the facilitator and innovator level, though; fostering widespread user-base adoption while also insulating against worker displacement is a critical challenge for policymakers.

Experts hold divergent views on AI and its potential impact on the labor market. Generative AI has the potential to improve workforce productivity by automating tasks, accelerating them, or reducing the skills required to perform them. As is true with any transformative technological change, there will be winners and losers. For example, automated telephone switching displaced switchboard operators but created new technical jobs in telecommunications maintenance and network management.

The rate at which AI diffuses through an economy and the effects it has on a labor market depend on factors such as industry composition, incumbent workers' skills, and the degree to which a nation's infrastructure supports the many technical requirements of AI deployment. These factors influence how a general-purpose technology matures into usable, customer-facing products as well

as the market size that such end products compete for. The United States has long maintained a dominant share of the world's computer scientists, leading to many early AI breakthroughs, tooling, and applications oriented towards use cases within the U.S. labor market. Despite the technological frontier being pushed forward almost exclusively in the U.S. and China, Israel and Singapore have seen the greatest level of AI utilization per capita (Anthropic economic index: Global usage, 2026). This suggests that AI innovators are not necessarily located in labor markets most ready to deploy and use generative solutions. This report outlines a framework for describing occupations in the context of AI, discusses labor conditions for large-scale AI adoption, proposes mitigation strategies that governments might implement to protect disaffected workers, and poses a series of open research questions that may be valuable to industry leaders and policymakers.

LITERATURE REVIEW

AI comprises a wide variety of general-purpose machine learning algorithms that enable digital and cyberphysical systems to perform tasks associated with human intelligence or cognition, including, but not limited to, Large Language Models (LLMs), which dominate much of the media and academic discussion. The expansion of AI capabilities and the wider adoption of AI have the potential to rapidly change how human capital and the global economy are organized. However, the pace, content, and distribution of this change remain the subject of ongoing debate. An emerging body of literature raises four interconnected questions that are central to the debate on AI and its consequences: (1) What is the role of governments in supporting the workforce in the AI transition, and what problems will market forces alone not solve? (2) Who is adopting AI, how quickly, and what individual- and firm-level factors predict early adoption? (3) AI replaces some human tasks, but for which workers is this replacement leading to higher productivity and wages, and who is being displaced? (4) How will AI's

expansion change the organization of firms and of the global economy, and what factors influence a country's optimal policy strategies to manage AI effectively? To help set the stage for a discussion of how workforce and skills policies can accelerate the diffusion of economically useful AI tools, we review the evidence on each question and explore critical gaps in the literature.

I. Training and education

As evidence mounts that AI's impacts are heterogeneous across task types and skill levels, the question of how best to upskill, reskill, or redeploy affected workers remains a pressing policy challenge. In what cases will market forces be sufficient to facilitate transitions, and whether and when governments should step in, are key questions. Amazon's Future Ready 2030 initiative pledged \$12.5 billion to strategically upskill 50 million people, underscoring the importance of this challenge for both the public and private sectors. This emphasis is echoed at the policy level, with the G7 Leaders' Statement on AI for Prosperity highlighting

AI skills development as a central priority for governments, reflecting a growing international consensus on the need to invest in workforce adaptation. Despite this, to date, there remains little empirical evidence on how firms and governments should redesign training and educational programs in response to AI's rapid adoption.

Early research on AI shows evidence of rising demand for AI-related technical skills (e.g., MLops, CSET analysis, and model deployment and integration) and associated wage premiums. Alekseeva et al. (2021) document significant increases in employer demand for AI skills across job postings, while Georgieff and Milanez (2021) find similar patterns globally, with workers who successfully upskill into AI-adjacent roles commanding meaningful wage premiums relative to their peers. Similarly, Acemoglu et al. (2022) use online job posting data and find that AI-exposed firms reduce hiring in non-AI positions and simultaneously increase demand for AI technical skills as AI adoption deepens.

Consistent with earlier studies, new evidence points to continued demand, but with generative AI skills incurring large premiums on top of traditional AI skills. Soares Martins Neto et al. (2026) find that, in digital occupations, AI skills command wage premiums of 11-14%, while generative AI development skills provide an additional 7-9%. In occupations that utilize digital technologies (but are not primarily involved in developing digital solutions), such as finance and marketing, genAI literacy skills generate wage premiums of 25-36%. While meaningful patterns in hiring and skill demand have been documented, research on which existing or new training programs may help meet this demand in the evolving labor market remains scarce. Hyman et al. (2025), using earning records and job training participation spells from the U.S. Workforce Investment and Opportunity Act (WIOA), assemble some of the first evidence on training programs that are effective for reskilling and upskilling. They find an average return to training of \$1,470 per quarter among AI-exposed workers, and, while positive returns exist to upskilling in AI tools, those returns are higher when trainees target non-AI-intensive occupations for career transitions.

Hyman et al. (2025) demonstrate that returns to investing in more general (non-AI) skills may be high. Skills such as resilience, ethics, and analytical thinking complement AI technical skills (and even command a wage premium) and may therefore be important for education and

training programs to emphasize as sources of human competitiveness (Mäkelä & Stephany, 2024). At the other end of the skills spectrum, occupations requiring physical dexterity and non-routine manual tasks, including skilled trades such as electricians, plumbers, and welders, consistently score near zero on AI exposure measures (Eloundou et al., 2024; Massenkoff & McCrory, 2026), and the BLS projects employment in many of these occupations to grow faster than average through 2034, suggesting that vocational training programs may represent a near-term hedge against AI-driven displacement. Vocational occupations are further insulated by their diversity and the low digital capture of their work contexts, a feature that increases the technical difficulty of automating any kind of work. Vocational programs represent a wide range of skilled trades and professions, and a wide distribution of wages. Wage-sustaining or improving transitions options into vocational occupations for displaced workers will vary for workers with different skills and wages. Displaced workers with high wages or specialized skills will require intensive re-skilling, while lower-wage workers may be able to make easier transitions. For example, if an insurance underwriter's job is automated, vocational programs for nursing or electrical work may pay similar or higher wages, but will require intensive training. On the other hand, if a clerical worker is displaced, programs for vocations such as dental hygiene and commercial truck driving, which require less intensive reskilling, represent a wage-improving or sustaining transition. Importantly, the total demand for vocational occupations may not be high enough to absorb all displaced workers, and alternative strategies are critical.

Another potential avenue for reskilling displaced workers is sector-focused training programs that target job seekers in occupational clusters with strong local labor demand (typically manufacturing, IT, and healthcare). Past research from prior periods of economic or technological transformation demonstrates that these programs can create substantial (12-34%) and persistent earnings gains (Katz et al., 2021), particularly when paired with wrap-around services and soft-skills training. These earnings gains are generated by moving employees into higher-wage jobs and industries, and not just by raising employment. However, there is a notable gap in evidence on how such programs interact with AI-driven labor market changes, or how best practices can be applied for AI upskilling. Sector-focused training programs in the U.S. have mostly targeted regional demand, but the disintermediating nature of AI

poses a question of whether localized strategies will incur the same gains. Programs such as Laboratoria (in Peru) and ALX Africa offer reskilling that targets remote work, uncoupling workers from having to rely on local demand. These programs have shown promise in connecting workers to remote jobs in technology, customer service, and basic IT. Whether sectoral employment programs can be leveraged to help workers acquire AI-relevant skills and successfully transition into non-exposed or AI-augmented occupations remains an unanswered question.

The literature on training, education, and upskilling in the age of AI is sparse. Most research to date focuses on task-level exposure estimates to understand shifts in the composition of skill demand. While these estimates provide a useful starting point, there is limited causal evidence on which types of training and education programs may successfully upskill, reskill, or redeploy workers displaced by AI. For instance, quasi-experimental and experimental studies examining the effectiveness of employer-led apprenticeships, online credentialing platforms, and other retraining programs specifically in the context of AI-driven labor market disruption are largely absent from the literature.

II. Adoption of AI

The rapid diffusion and advancement in capabilities of AI tools over the last three years have taken on a global scope. In a survey of approximately 2,000 respondents across 105 countries (Singla et al., 2025), 79% reported that their organization used GenAI in at least one business function, up from 33% in 2023. This rate of adoption is largely attributed to GenAI's characteristics as a general-purpose technology with pervasive utility across a wide range of cognitive tasks (Eloundou et al., 2023). Despite this marked increase, the AI adoption varies across worker and firm types and depends on factors such as the regulatory environment, task complexity and separability, and the availability of skilled workers. Where and how policymakers can deploy training and upskilling solutions depends upon adoption strategies at the firm, process, and individual levels.

The sources of measurement on AI adoption are largely reliant on firm-level surveys: McElheran et al. (2023) utilize data from the Annual Business Survey (ABS) to study five AI-related technologies (automated-guided vehicles, machine learning, machine vision, natural

language processing, and voice recognition) and find that AI usage is highly clustered in emerging technologies such as robotics. Additionally, the authors find significant geographic variation in early adoption, with large “superstar” cities such as San Francisco accounting for a disproportionate share of AI usage. However, limitations in ABS data prevent the authors from studying publicly owned firms and require them to examine a subsample of very young firms (<5 years) to obtain granular information on AI adoption.

Firm-level evidence suggests that AI adoption is concentrated among larger (Zolas et al., 2020) and younger firms (McElheran et al., 2024): firms in the top percentile of industry employment are 9 percentage points more likely to adopt AI than below-median firms, and younger firms are 1–2 percentage points more likely to do so. Adoption rates vary significantly across industries (Acemoglu et al., 2025). For example, the information technology sector has high adoption rates, whereas the agriculture sector has very low adoption rates. These disparities reflect a combination of high fixed implementation costs (e.g., the up-front time, effort, and money required to adopt a solution), organizational barriers to integration, and differences in the degree to which AI tools can meaningfully substitute for or complement sector-specific tasks.

Beyond firm and industry variation, there is significant variability in worker-level adoption patterns. Synthesizing data from a nationally representative survey (the Real Time Population Survey, RPS) on generative AI usage, Bick, Blandin, and Deming (2025) examine characteristics of GenAI users in the labor force. The RPS mimics the structure of the Census Current Population Survey (CPS). The RPS (and worker responses generally) suffer from sample selection and self-reporting biases, but reveal several notable patterns. The authors find that, although genAI adoption speeds significantly outpace those of PCs and the internet, demographic patterns among workers are similar. Namely, AI use is more common among higher-educated, younger, and higher-wage workers: those with a bachelor's or graduate degree are twice as likely to use AI, and adoption rises with income until about the 80th percentile before tapering off. While evidence points to higher adoption rates for the performance of tasks typically associated with computer and mathematical careers (who are likely to be innovators and facilitators),

these signals are from the application layer (e.g., the company hosting the LLM). A student “vibe-coding” or using AI to cheat on their math homework might, in this data, appear similarly to a software developer or a mathematician. As a result, the quantification of adoption rates by professionals across layers of the AI pyramid remains an open question.

The evidence on AI adoption remains nascent. Most studies to date rely on administrative survey-based datasets (such as the ABS) that ask only granular questions about AI adoption to a subset of firms and lack information on the intensiveness of AI adoption. While some measures characterize the magnitude of AI use through firm- or user-level API use, these measures cannot be interpreted as quantitative measures of productivity gains to the firm or the percentage of business processes to which the technology has been deployed. These measures are also insufficient to determine whether a firm is using AI to become more efficient (doing the same scope of work at lower cost) or to expand its portfolio (offering entirely new services). Administrative surveys offer limited questions on generative AI specifically and are subject to long lags, which inhibit tracking the rapidly evolving technical and economic landscape. Anthropic (followed by OpenAI) has made an important contribution to the measurement of AI adoption by providing data on the characteristics of user sessions, with lower latency and higher granularity, but with the critical limitation that the data do not differentiate consumer or “household” utilization from utilization in a business context (or behavior from individual early-adopters acting autonomously within a firm, versus adoption at the firm-level). Thus, the evidence on adoption patterns and task exposure (discussed in the following subsection) is likely to be confounded by differences in the incentives for AI usage.

Consumers, employees, and employers will all adopt for different reasons: what is an acceptable level of performance (e.g., precision, or error) for a household user may not suffice for a business. What saves an individual worker time might not drive firm-level adoption or widespread productivity gains. In turn, task-level characteristics, such as task frequency or co-occurrence, remain largely unmeasured, but they are likely to differ significantly, and with them, the economic returns to automation.

The labor impact of AI will be driven by the economics of

adoption and the corresponding redesign of jobs: Ales et al (2023) identify four relevant dimensions that govern how technologies may substitute or complement labor: (1) fragmentation costs incurred from dividing tasks into steps, (2) sensitivity of the performer (whether human or machine) to the task complexity, (3) sensitivity of the performer to rate, and (4) the cost of relocating a performer once a step is complete. They argue that the technical feasibility of automating a task is not enough to predict adoption patterns: rather, historical automation has enjoyed an advantage when process steps are short (containing few tasks, hence of low complexity) and when the rate of work is high. When volumes are low, humans have competed for the simplest work because it can be done fast (relative to middle- or high-complexity work) and humans are more flexible to reallocate than machines; when complexity is high, humans have competed on their ability to handle uncertainty and cover a wider range of solutions than machines. The result has been a “cone of automation” (Figure 2) in which low-volume work remains manual, and as volume increases, middle-length, middle-wage work is automated first, before widening to automate simpler high-volume work. AI tools have changed the conventional advantages of machine performers, being both less costly to redeploy and better able to handle a wide variety of inputs and contexts, effectively widening the cone of automation (Ales et al, 2024).




Figure 2: Widening Cone of Automation by GenAI
(Ales, Combemale, and Krishnan 2024)

III. Labor market impacts

The uneven adoption patterns of AI raise important distributional questions about who stands to gain and to lose as the labor market changes in response to AI adoption – and how governments should respond through workforce and skills policies. Artificial intelligence

diverges from historical patterns of technological change (especially automation, such as assembly lines at the turn of the 20th century and computers at the turn of the 21st century) because it can perform unstructured, high-variance cognitive tasks. This divergence has led to two competing theoretical perspectives: on the one hand, AI's unique capabilities may substitute for human labor, creating downward pressure on wages and employment. On the other hand, AI may complement human labor by augmenting worker productivity and expanding the range of tasks a worker can perform.

These findings set up an important tension between firm-level and industry-level labor demand. While some firms leverage AI to increase their scale and hiring, growth for early adopters ("first-movers") may be a sign of competitive shifting (from non-adopters to adopters), and not aggregate job creation. Individual firms may still see gains even if industry-level employment shrinks in the aggregate.

A growing body of research suggests that whether AI acts as a complement or substitute for human labor varies significantly across workers, firms, and industries, depending on factors such as skill level, occupational task composition, and the extent to which organizations have the flexibility to redesign workflows around the technology. Noy and Zhang (2023) study the effects of ChatGPT access on college-educated professionals' ability to complete a mid-level writing task and find large productivity gains (a 40% reduction in task completion time alongside an 18% increase in quality). Similarly, Brynjolfsson, Li, and Raymond (2023) study the rollout of a generative AI tool among customer support agents and find a 14% increase in average productivity. Both studies document important heterogeneities: low-skilled workers experience larger productivity gains, suggesting that AI compresses skill-based productivity differentials rather than augmenting productivity uniformly across the labor force.

While augmentation may increase productivity in the short-term, especially for low-skilled workers, Acemoglu and Restrepo (2022), drawing on 29 years of data from the BEA Integrated Industry-Level Production Accounts, document that task displacement driven by the prior era of automation has been associated with increased wage inequality, as capital substitutes for labor in ways that disproportionately harm workers at the lower end

of the skill distribution. Extending upon this, Acemoglu (2024) argues that aggregate productivity gains can be estimated as a function of the share of automatable tasks and their associated cost savings. Conversely, Autor (2024) contends that AI has the potential to restore middle-skill labor demand by enabling a larger set of workers to perform tasks currently performed by experts (e.g., software coding by computer engineers).

A key input to understanding these dynamics is the mapping of AI exposure at the task level. Autor, Levy, and Murnane (2003) pioneered the first framework for thinking about technology substitution at the task-level, dividing tasks into four categories by degree of routineness and level of cognitive input required. They document that the introduction of computing technology caused a sharp decline in demand for routine cognitive and non-cognitive tasks.

Several papers have sought to estimate task-based exposure to new technology, beginning with attempts to quantify potential job loss to automation, prior to the current generation of AI technologies. AI-exposure papers such as Eloundou et al. (2024) construct occupation-level exposure scores by systematically assessing the extent to which large language models (LLMs) can perform or assist with each task that comprises a given occupation. Using this framework and detailed task data from O*NET, the authors estimate that 80% of the US workforce could have at least 10% of their tasks affected by LLMs (meaning LLMs could meaningfully assist with or influence those tasks), while approximately 19% may see 50% of their tasks exposed (meaning LLMs could substantially perform or potentially replace a large share of those tasks).

While Eloundou et al. (2024) measure what tasks *theoretically* can be completed by AI, Anthropic extends upon their measure to create a measure of observed exposure (Massenkoff & McCrory, 2026), utilizing usage data from Claude. They find that AI is far from reaching its theoretical capabilities and that there are no impacts on unemployment rates for workers in the most exposed occupations; however, workforce administrative data (Brynjolfsson et al., 2025) suggest that early-career workers in "AI-exposed" occupations have experienced declines in employment since the release of ChatGPT. Evidence linking firm-level adoption to employment outcomes (both within firms and across industries due

to potential firm out-competition) will be necessary to resolve this empirical disagreement.

Taken together, the existing literature establishes a theoretical and empirical foundation for understanding AI's labor market impacts, but several important gaps remain. Experimental evidence is largely drawn from small, controlled settings that may not generalize to economy-wide labor market dynamics. Exposure-based approaches often rely on datasets subject to latency issues (such as O*NET) that limit their utility in tracking a technology diffusing as fast as AI; while the Anthropic "conversation data" makes important strides to address these limitations, they are limited to a specific form of AI technology and lack pairing with user occupations. While we have evidence of task-level utilization, we do not know whether these tasks are being performed as part of a core occupational activity or incidentally.

Additionally, much of the existing theory focuses on short-term effects in existing industries. The potential for AI to create entirely new industries and sets of tasks is understudied, but often cited as an argument that net job creation from AI will be positive. Indeed, most workers in the U.S. today famously work in occupations that did not exist before the 1940s (Autor, 2015). However, evidence from past labor shocks also suggests that net job creation can mask significant individual-level disruption, with workers who lose out from a transition often unable to regain their prior wages (Autor, Dorn, and Hanson 2021).

Leveraging the task-exposure view suggests a similar potential nuance for the impact of AI on labor markets: Figure 3 illustrates how the exposure score of occupations (using the U.S. Standard Occupational Classification System, which can be crosswalked to the International Standard Classification of Occupations) correlates with the exposure score of the occupations that workers have historically moved into when leaving their original occupation. That is, if exposure (particularly very high exposure) does lead to automation, it may imply the simultaneous automation of neighborhoods of occupations, leaving workers stranded and unable to transition to comparable roles, much as in prior shocks – this resiliency threat may create an intervention point for government to invest in workforce development interventions that decouple worker exposure risk from their current occupational neighborhood.




Figure 3: Correlation of Exposure Score Between Occupations and their "Exit Options" (Ales, Combemale, and Krishnan 2024)

IV. International Perspective

Training and upskilling opportunities and challenges differ internationally in important ways, but most of the literature to date draws on evidence from the U.S. and high-income countries. Whether, and how much, countries should invest in upskilling, reskilling, or other policies that make labor forces resilient varies significantly with the composition of their labor force.

Pizzinelli et al. (2023) provide some of the earliest empirical evidence on the differential impacts of AI across high- and low-income economies (and across economies with higher innovator bases vs. higher user bases). Using worker-level microdata across six countries, they document that high-income economies have higher exposure to AI due to higher shares of professional and managerial occupations, putting them at risk of greater labor substitution but also of productivity gains from AI. For instance, the least exposed 40% of workers account for 70% of employment in India, whereas the same exposure accounts for only 25% of employment in the UK (Pizzinelli et al., 2023). At the same time, research on lower-income countries posits that they face "double vulnerabilities": a concentration of employment in low-skill occupations (which are more easily automatable) and lower preparedness for AI diffusion (due to less technical infrastructure and human capital) (Ganuthula & Balaraman, 2025). Thus, distinctive upskilling and reskilling strategies that account for the international competitive positions of individual countries are critical for resilience in the global economy.

A host of different policies to create an AI-ready workforce are being tested globally. Finland's "Elements of AI" and India's "FutureSkills Prime" programs are two examples of upskilling the workforce in AI tools. Both platforms offer AI courses for firms and individuals, covering topics from basic AI literacy to developing AI tools. FutureSkills Prime offers financial incentives for completing technical courses and focuses not only on AI, but also on cybersecurity, robotics, and other technological advancements. The UK's Skills Bootcamp similarly offers scalable, short bootcamps for entry-level roles in the tech sector, and a job interview post-completion of the program. FutureSkills Prime differs from its counterparts in that it focuses on upskilling into the facilitators layer of the AI labor pyramid. India has a comparative advantage in outsourced, lower-cost technical labor (as evidenced by their IT sector). Thus, the expansion of facilitators presents an opportunity to replicate this comparative advantage in the global AI market.

Beyond incentivizing AI upskilling at the firm-and-worker level, some countries have implemented policies that articulate a tradeoff between oversight and the rate of adoption. A 2025 executive order in the United States created the AI Litigation Task Force, aimed at reducing state-level oversight laws and increasing AI growth. Similarly, the EU's Digital Omnibus Regulation Proposal amends several cornerstone EU regulations, expanding opportunities for real-world testing, streamlining assessment frameworks for high-risk AI systems, and easing use constraints on pseudonymized data – these policies address the facilitator layer of the AI adoption pyramid as much as the innovators.

Training and education to upskill and reskill workers takes time, making safety nets an important part of resiliency strategies for displaced workers. Denmark's Flexicurity model provides an example, as it allows employers to hire and fire at will and allows employees to invest in a safety-net system that pays for two years upon unemployment. The Danish government pairs this program with government-funded training and education programs.

While there is an emerging focus on creating a resilient labor force, upskilling and reskilling training is still limited in availability and scope, and open questions remain about optimal policy choices for different labor force compositions. To date, many AI courses focus on advanced AI skills, such as model development and

deployment, which may be less effective for countries whose labor force is concentrated at the bottom tier of the AI pyramid. Oumiette, Tether, and Allison (2024) point out that "making the most of AI depends not just on advances in frontier AI models, but crucially also on accelerating the workmanlike adoption of AI to solve practical problems in many different industries." For countries with limited access to innovators and facilitators and that face dual vulnerabilities, focusing on basic AI literacy and practical applications (and thus expanding the user tier of the pyramid) offers a more resilient strategy.

For countries with strong bases in innovators and facilitators, there are gains from upskilling in advanced AI skills. First movers in AI can gain advantages such as talent attraction, access to proprietary data, and the ability to set industry standards before competitors. Additionally, first movers can create new markets, as demonstrated by the launch of ChatGPT in 2023. However, these advantages can fade over time as competitors bridge technical gaps, making it important to invest in broad-based AI literacy across user bases as well.

Safety nets that can encourage rapid reemployment of disrupted workers, paired with upskilling and reskilling, are critical to building an AI-ready workforce. Denmark offers a potential solution: allowing employees to invest in safety nets that protect them in the event of displacement, and reskilling them for less exposed occupations or upskilling them in AI tools to make them more competitive. At the same time, growing evidence from the U.S. context suggests that wage insurance may be a more effective safety net to support reemployment of disrupted workers than conventional unemployment insurance – by compensating workers for the difference between their old wage and their new wage, the policy subsidizes employment (rather than making unemployment more appealing) and encourages workers to transition and rebuild their wages over time (Hyman, Kovak and Leive 2024). This lever may be most effective for occupations that experience AI-related disruption systematically (meaning that workers may face at least temporary wage losses), while traditional unemployment insurance may be most effective for supporting workers in gaining the skills to make wage-improving transitions, where these are possible.

While labor market dynamism will not address all of the potential systemic disruptions from AI, policymakers may

be able to accelerate worker transitions into AI user and facilitator roles by aligning worker support with employer incentives, for example, subsidizing employer-led training programs and facilitating direct interviews for graduates of those programs. International evidence suggests that workers treat unemployment support from government and employers as substitutes, meaning there may also be opportunities for incentive-compatible designs that reward employers in vulnerable industries for recognizing and addressing their own disruption risks before they might become apparent to the government (Ellul, Pagano, and Schivardi 2018).

Differences in training adoption strategies (such as whether to invest more in broad user-base AI literacy or to upskill for facilitator roles) depend not only on a country's labor force composition but also on its global labor market position. The relationship between technological advances and the geographic concentration of labor is not new. In the 1980s, the commercialization of telecommunications infrastructure enabled firms to route calls to call centers in India and the Philippines, where workers could be paid lower wages. This technological advancement led to a mass offshoring of labor to lower-income countries.

The parallel with the current AI moment is instructive, but there are stark contrasts, such as the ability to automate cognitive tasks completely away from humans. This factor may alter its global impacts: a decisive open question for the global impact of AI is whether it will tend to substitute for occupations that have historically been offshored (and thus cannibalize the service exports of many lower-income economies for which these types of lower-tier service jobs are a key opportunity for young and growing populations). This presents a challenge for programs like FutureSkills Prime, if jobs in the user and facilitator tiers are likely to be automated away.

Freelance and gig platforms offer a compelling lens for studying the global economic impacts of artificial intelligence. This sector is characterized by a high concentration of non-US workers and a complex dynamic in which AI simultaneously serves as both a substitute for human labor and a tool to augment it. Additionally, AI

reduces translation costs, potentially enabling offshoring, as in the telecommunications industry.

Hui et al. (2024) studied the impacts of generative AI on Upwork, one of the largest online labor markets in the world. Their results provide unique insights into the global impacts of AI, as much of the work on Upwork is characterized by short-term contracts and concentrated in lower-income countries where gig work remains a pathway to middle-class income. The authors find that the release of ChatGPT led to a 2% drop in freelancing jobs and a 5.2% drop in monthly earnings, and that even high-skilled freelancers were not shielded from these impacts. Studying the release on a similar platform, Yui et al. (2025) document increases in entry into these markets, but incumbent freelancers bid on fewer jobs and exhibit greater differentiation in their bids. Additionally, AI helped lower barriers to entry in the freelancing market (such as language), particularly for international participants.

The global diffusion of AI poses a paradoxical challenge for international training and upskilling strategies. Lower-income countries with low labor costs should capitalize on that competitive advantage by investing in facilitators and user bases. Simultaneously, they need to examine areas where competitive advantage may erode (as AI automates tasks) and instead invest in jobs within facilitator layers that will be augmented, but not automated, by AI. Whether facilitators are context-agnostic (e.g., workers who monitor AI system performance for technical failures, such as flagging return code errors) or context-dependent (e.g., localized safety and bias reviewers for AI models) also plays a role in what AI services are exportable.

Conversely, high-income nations with robust technological infrastructure have a unique opportunity to secure first-mover advantages by investing in the 'innovator' layer. However, over-concentrating resources on frontier development while neglecting broad-based AI literacy creates a critical vulnerability; early technological leads will inevitably narrow, making a highly adaptable user base essential for sustained economic resilience.

PATHWAYS FROM PRODUCTIVITY GAINS TO LABOR IMPACT

As outlined by Ales, Combemale, Fuchs, and Whitefoot (2023), tasks may be finely fragmented (subdivided into shorter, less complex tasks), but the cumulative cost of fragmentation increases rapidly. Fragmentation cost is largely driven by the costs of direct interaction (passing parts or information back and forth between performers) and idling (waiting for a predecessor to provide the necessary inputs to perform a task). Fragmenting tasks into less complex steps (in the aggregate) lowers the minimum skill requirements of performers, which is associated with lower wages. Cost-minimizing employers may then optimize the trade-off among fragmentation costs, labor costs, and the rate of completion. The emergence of multi-agent systems reflects these dynamics: System designers use several specialized models to accomplish divisible tasks that would otherwise require a far more powerful model to achieve the same functionality independently. This incentive pattern extends to the other extreme of task divisibility, where it is technically feasible to break a process into increasingly granular tasks. If tasks are divided too much, a system designer could pass the point of diminishing marginal return and achieve negative marginal return. An extreme example of this might be assigning dozens or hundreds of agents to process meeting minutes; many assigned specific chunks of the conversation to summarize (one minute or one statement at a time), and a final set of agents to aggregate important elements and follow-up activities. The number of agents in this case may be so large that the solution is nonperformant, both due to the up-front cost of designing and implementing the system, the fragmentation costs associated with information distortion between agents, and inflated token consumption.

While AI may perform entirely novel tasks, many of the first-order gains to worker productivity will derive from AI (and AI-augmented machines) performing tasks that were once reserved for humans. Whether that task substitution leads to the replacement of labor or the augmentation of labor productivity depends on the

share of tasks being substituted and the compensating demand for their remaining unaffected tasks. AI-driven task substitution may create new supervisory tasks, but the process change only makes financial sense if the total cost per step performed decreases (something most likely to result from a decrease in working hours per task completion).

The downside of most task-replacing technologies is their lack of generalizability, meaning that it is difficult or costly to reconfigure the tool to perform a different task. AI (particularly LLMs) is comparatively more general-purpose, requiring negligible cost to pivot to a different task once an implementation strategy has been established. Furthermore, unlike physical assets, SaaS solutions often follow a pay-as-you-go model, meaning that the end user bears no cost for machine idle time. As a result, the fragmentation cost of tasks performed by AI may be lower than that of the same tasks performed by humans, shifting the optimal degree of fragmentation towards shorter, less complex tasks. We may refer to the fragmentation of a task into more granular elements, performed by a mix of humans and AI, as “task augmentation.” This can resemble past factor-augmenting technologies, in which the combination of human and AI capabilities enables greater output.

When optimal task assignment shifts to a lower complexity, the process demands a lower skill distribution of its workers. This creates downward pressure on wages for this task. Occupations typically involve performing many tasks. Thus, the automation or augmentation of any one of those tasks may not necessarily correlate to a decrease in demand (or potential earnings) for the affected worker. If automating a task allows workers to spend more time delivering higher-value tasks with highly elastic demand, there may be an increase in demand for these workers. Otherwise, efficiency gains from adoption may lead to lower overall labor demand at scale.

Within a firm or economy, the economic opportunity to achieve productivity gains from AI adoption is

proportional to the presence of (what is often referred to as) AI-exposed workers. It is the tasks of these workers that might be made more efficient with AI. This adoption of AI tools necessarily drives demand (in the economy, if not directly within the firm) for workers who facilitate its use or develop new or improved applications (solution integrators, data center workers, prompt engineers). This, in turn, expands the market for tooling suites that make application development and deployment more accessible. In this way, adoption reduces the demand for certain skills held by incumbent workers and increases demand for complementary skills sourced from facilitating occupations.

Wages can only increase for affected workers if their economy experiences compensating demand in response to new productivity gains. Therefore, policymakers may improve worker outcomes by adopting strategies that increase demand. Examples include: promoting economic growth in businesses or industries with highly scalable demand for substituted tasks (e.g., industries that will do more with more, not the same with less), and

improving access to retraining for workers in contracting fields (possibly by subsidizing program attendance).

A major unknown for the impact of AI on economic productivity broadly and outcomes for workers is the development of new tasks (often associated with new goods and services that improve economic welfare). The contours of new tasks might be anticipated, however, in terms of the same adoption principles that drive task-level substitution: many new tasks generated by the use of AI are already being performed by machines themselves (various layers of multi-agent coordination, for example), in part because the medium in which AI products are delivered is itself principally digital, occurring at high volume and under conditions that are highly legible to AI systems. Thus, it is unlikely that workers disrupted by AI will be able to find new employment doing similar work, since those new tasks are likely to be automatable for the same technical and economic reasons – disruption may not preclude reemployment, but it is likely to represent a substantial transformation, perhaps demanding very significant changes in the capabilities of workers.

LABOR CONDITIONS FOR AI ADOPTION

AI is already showing strong evidence of being a General Purpose Technology (GPT). Historical GPTs (e.g., steam engine, electricity, and the microprocessor) are foundational innovations that facilitate the development of new and diverse products and processes across industries. By changing what is considered technically feasible, GPTs reshape productivity, business models, and the workforce. AI diffusion may follow a similar labor market pattern, with highly skilled innovators who push forward the technology's capability frontier; technical workers who deploy, integrate, and otherwise facilitate the technology's use; and end users whose work is altered by the technology.

Categories of AI Workers

We may think of AI workers (those who develop, enable, or use the technology) in terms of their technical expertise

(specific to AI) and the relative share of jobs in the economy that demand that level and type of expertise. The largest share of workers will fall on a continuum in their day-to-day use of AI-powered tools (e.g., lawyers may use more tools that embed AI than construction workers, or work with tools that use AI more explicitly), but generally will not have responsibility for managing, integrating, or deploying those tools. Globally, we would expect fewer workers to be responsible for integrating, managing, and otherwise supporting the deployment and use of AI-powered tools. These enablers of AI technology must have greater AI expertise than the users. However, the innovators who develop those tools will require a greater level of AI expertise than the enablers. We would expect these innovators to be in demand in even smaller numbers than facilitators, given the global workforce. Where innovators number in the thousands, facilitators

that use new AI tooling and maintain critical infrastructure must be employed in the millions, empowering potential billions of application users (AIAI, 2025). We can visualize this relationship between expertise and labor quantity demand as a pyramid – users at the base, innovators at the peak, but with the economic returns to AI technology ultimately generated by end-use at the base (or via automation of tasks traditionally performed by them).

Tier 1 - Innovators

Among these three tiers, innovators will constitute the smallest labor pool and may be expected to demand a wage premium relative to other roles in this labor pyramid. These experts and product developers actively advance the frontier of what is technically feasible with AI technology and require the highest degree of AI expertise. Workers at this tier might be expected to build tools, or tooling platforms, without a specific customer or application in mind (e.g., new LLMs, agentic frameworks, deployment platforms, protein folding algorithms). While the advancements they achieve may be impactful, their absence from a specific economy does not necessarily preclude AI adoption. For example, while the majority of top AI firms are in the US and China, Israel and Singapore have seen the highest levels of AI utilization per capita (*Anthropic Economic Index: Global Usage, 2026*). This is a phenomenon with historical precedent; Jeffrey Ding's book *Technology and the Rise of Great Powers* outlines numerous examples of GPTs developed by innovators in one country that achieved greater diffusion in another. As discussed later in this report, reliance on highly skilled foreign labor can present new risks.

Tier 2 - Facilitators

These are the technicians and engineers who enable the deployment and sustained use of AI-powered tools. Workers at this tier might be thought of on a spectrum from most to least likely to be directly employed by the same firm as the end user. We might expect data management, validation, quality analysis, and other functions that require deep contextual knowledge to be performed "in-house" more often than the equally necessary work performed by IT workers or systems integrators. Meanwhile, almost no end-use firms will directly hire workers responsible for maintaining data centers, telecommunications infrastructure, or performing other roles divorced from the end user's context.

Existing Jobs	New Jobs
■ Systems Architect	■ Model Context Curator
■ Systems Integrator	■ Alignment Engineer
■ Network Engineer	■ AI Site Reliability Engineer
■ Systems Administrator	■ GPU Cluster Manager
■ Cloud Engineer	

The common thread tying all of these roles together is that they must be available in the economy for AI tools to be deployed and maintained, but they are neither advancing the technological frontier nor using the technology to generate a final good or service. The services rendered by this tier of workers may be accessed through the international trade of services. However, this solution introduces other inefficiencies (e.g., system reliability, system latency, language barriers, geopolitical-driven access restrictions, and regulatory hurdles). This form of outsourcing may be most technically feasible for roles that do not require intimate knowledge of the end user's context.

In this tier, we should expect to see some roles persist (albeit with modest education required to keep up with technological advancements) and the emergence and growing prominence of new roles. The following table provides some examples.

Tier 3 - End-Users

Workers in this tier primarily consist of those employed in AI-Exposed Occupations (Eloundou et al., 2023). Exposed workers are those who perform a greater share of tasks that might be changed or entirely automated by AI. The opportunities to achieve AI-driven productivity gains in an economy are proportional to the number of AI-Exposed workers within it and to those workers' relative wages – whether the gains from AI adoption accrue to the countries where these workers are currently employed depends on whether the workers can be augmented rather than automated.

The outcomes of these workers depend heavily on which of their tasks are exposed (e.g., the technically feasible to automate or accelerate with AI powered tools), the share of a worker's time spent performing exposed tasks (factoring into economic feasibility), and the market demand for their remaining, unexposed, tasks (e.g., can they perform more valuable work once an exposed task

takes up less of their work day). How an LLM may reshape a task depends on the model's capabilities. As the technology matures, we may see a greater shift towards task automation and a corresponding increase in market penetration. As we argue in the following section, the conditions driving compensating demand for labor are a salient target for national policy.

When a task requires less human time, the affected workers gain greater availability to perform other tasks. The resulting employment outcomes might then be explained through a discussion of compensating demand for the unaltered tasks. At one extreme, the demand for the worker's remaining tasks does not increase meaningfully as the cost to perform them decreases. In this case, the total worker time demanded decreases, resulting in a lower aggregate demand for that worker's occupation type. At the opposite extreme, the demand for the work's remaining tasks increases substantially as the cost of performing them decreases. This would increase demand for that worker's occupation type. In both cases, the most likely tasks to be automated are those with

the highest desired completion rate and those whose complexity is low enough for AI tools to perform reliably. For past technologies, there was also a minimum task complexity that was economically viable to automate. For example, although the technology has existed for over a decade, meatpacking is still largely performed manually by minimum-wage workers.

Due to the comparatively lower cost of deploying and maintaining AI models, automating some simpler tasks has become economically viable (where it previously was not). For example, call center workers are frequently employed by U.S. firms, but with those call centers located in countries with comparatively lower labor costs. The availability of low-cost labor, compared with the high technical barriers to automation, has historically insulated these positions from some forms of automation. However, the commercial availability of generative language and speech models, combined with the tooling layers for deploying them, has cleared the path for several firms to develop AI call center workers as commercial-off-the-shelf products.




Figure 4: Demand for tasks unaltered by computerization

POLICY & OPEN RESEARCH QUESTIONS

As AI adoption accelerates, policymakers face consequential decisions about workforce development, labor market regulation, and international competitiveness, yet critical questions about AI's economic impacts remain unresolved.

Adoption

Emerging data suggest that AI adoption is uneven across firm types, workers, and geographies. Evidence from “superstar” cities, such as San Francisco, suggests that local talent pools may play an important role in the diffusion of AI. Figure 5 illustrates how labor distributions differ from national averages in two super city Metropolitan Statistical Areas (MSAs): San Francisco and San Jose. Both cities deviate substantially from national averages in computer and mathematical occupations, which make up a large portion of the facilitators and innovators layers of the labor pyramid. A similar, though more muted story, emerges in architecture and engineering occupations; in both cases, San Jose, “the city of the future”, stands out as a geographic outlier, employing more workers from these occupations and below average office and administrative workers and transportation and material moving workers. If some areas have access to the talent needed to adopt and deploy AI quickly, while others do not, will this lead to widening geographic disparities in productivity gains? Local talent pools, beyond influencing the availability of technical skills, could also affect the quality of AI implementation, leading to or exacerbating productivity growth, wage differentials, and job opportunities. Traditionally, place-based policies such as Enterprise Zones (geographic areas characterized by high levels of poverty that receive special tax credits and regulatory exemptions) have been implemented to reduce geographic disparities. Evidence on Enterprise Zones is mixed, but pairing tax credits with other economic supports, such as technical assistance and special staffing, improves the efficacy of these programs (Neumark & Simpson, 2015). Traditional place-based policies that target subsidies to build infrastructure and attract business development in economically distressed

regions may not be effective for AI expansion. The example of broadband internet provides a compelling parallel and critical contrast. The USDA invests heavily in expanding broadband internet access to rural areas by funding physical infrastructure (such as laying fiber-optic cable) and encouraging individual-level adoption (through subsidizing Wi-Fi bills). While the supply-side policies for broadband focused on physical construction, the local bottleneck for AI (above and beyond the preconditions for internet access in general) is largely human capital. However, the demand-side approach of broadband policy offers a potential blueprint: just as subsidizing internet bills drove connectivity, providing local businesses and workers with financial subsidies for AI software licenses and cloud-computing access could drive localized adoption. Pairing these direct access subsidies with larger investments in workforce retraining, AI literacy programs, and human capital development may be necessary to bridge this new divide.

Rural, low-income, and underserved communities within high-income countries represent a unique case in which resilience-focused strategies could supersede national-level strategies. These communities are often characterized by a concentration of industries with low exposure to AI, either due to task-context barriers to technology deployment (such as home health aids or landscapers) or low cost of labor that makes current automation economically unattractive in some countries (such as grading and sorting agricultural products). Figure 6 plots mean income at the PUMA level (from American Community Survey data) against AI-exposure (using the Gimbell, Kendall, and Kulsakdinun, 2026, PCA weighted AI-exposure index). Findings indicate that there is a positive association between mean income and average AI-exposure.

While low-income communities may be less exposed in the aggregate to AI labor market disruptions, they also have lower access to facilitators. The ability to capitalize on gains from AI for low-income communities is largely dependent on geography and physical access; more




Figure 5: Employment Share Deviation from U.S. National Average, SF Bay Area & San Jose, CA

integrated communities that are in close proximity to high-income areas, or have easy access to transit to these communities, are better positioned. Workers in such communities will benefit from demand for non-AI-substitutable services created from overall economic growth. Policies in these regions should be resilience-focused: robust social safety nets that encourage workforce participation, bridges for workers who are displaced to reskill into jobs with lower exposure and

higher local demand, and addressing barriers (such as lack of childcare) to workforce participation are strategies to build a resilient workforce in these communities. In addition, investments in transit from these communities to communities with direct exposure to economic upside from AI are critical. The “Baumol effect” is a well-documented feature of unbalanced growth across sectors (Baumol 1967; Hartwig and Kramer 2019), in which productivity gains in some parts of the economy drive up wages for workers in other sectors through a combination of increased demand from more productive businesses and workers and from the need of employers to incentivize workers not to enter more productive sectors. While often framed in negative terms relating to increasing costs³, the Baumol mechanism creates a channel for gains from AI to accrue to workers that are not part of the GPT labor pyramid. Less-exposed work, including both low-income work or middle to higher-income trades, often cannot be transmitted digitally, meaning that the Baumol mechanism will have to act through geographic proximity. Thus, transportation access creates access to economic upside, as direct AI productivity gains drive both consumer and business demand for non-exposed occupations.




Figure 6: AI Exposure by PUMA Income

In turn, for isolated low-income communities, workforce participation strategies and robust social safety net structures are critical. These communities are unlikely, due to geographic and transportation constraints, to be able to capitalize on economic growth induced by AI productivity in high-income areas. Instead, reskilling workers into occupations with strong regional demand, and investing in safety nets such as unemployment insurance, are imperative to creating a resilient workforce.

Workforce Development

Understanding adoption patterns is a crucial first step to estimating the labor market impacts of AI tools. While a large body of literature has tried to estimate (both theoretically and empirically) who stands to lose and gain from AI, this is a nascent question with important policy and equity implications. A central tension in the literature concerns whether AI will augment lower-skilled workers, enhancing their productivity (Brynjolfsson, Li & Raymond, 2023; Noy & Zhang, 2023), or displace them by automating the tasks they perform (Acemoglu & Restrepo, 2020; Frey & Osborne, 2017). A parallel tension is whether AI will augment or automate high-skilled work, given its growing ability to perform non-routine cognitive tasks.

We argue that disparate findings in the current literature can be reconciled by considering not only a task's skill level but also its divisibility – the cost of splitting work between two performers. Namely, as tasks become more divisible (regardless of the skill level required to accomplish them), they become more exposed to AI-driven automation (Ales, Combemale, and Krishnan 2025). As tasks become more complex (requiring higher skills, associated with higher difficulty and, generally, higher wages), they become less automatable. Tasks become less complex when divided, creating a tension between complexity and divisibility. Figure 7 provides a framework for evaluating how task divisibility and complexity may affect substitution or augmentation patterns. Tasks studied in the literature concerned with potential AI disruptions to labor and employment, such as customer support responses, tend to be highly divisible and are primarily augmented by AI when workers retain highly complex work (Brynjolfsson et al., 2023; Noy & Zhang,

2023), while tasks with high divisibility resulting in low complexity, such as telemarketing or clerical processing, are expected to be displaced through automation (Pizzinelli et al, 2023). In contrast, low-diversity tasks or non-routine manual work, such as caregiving, are less likely to be automated, consistent with earlier task-based frameworks. Investments in these occupations offer a pathway to redeploy displaced labor into less-exposed positions.

In our framework, divisibility determines the extent to which AI can be applied, while complexity determines whether that application results in automation or augmentation. Occupations with tasks that are not divisible, whether low- or high-skill, are much less exposed to AI and are an important pathway for reskilling AI-exposed workers. Conversely, occupations with tasks that are highly divisible and high-skilled should be targeted for AI upskilling (instead of tasks that are highly divisible and low-skilled) as they are more likely to become augmented by AI, creating productivity gains for workers. Many of these productivity gains are induced by the automation of tasks in an occupation that are quality complements, meaning that the quality or success of one task directly increases the productivity of another (Gans & Goldfarb, 2026). In the highly-divisible and highly-complex task space, when AI handles the highly divisible elements of complex work, humans are freed to focus on high-level judgment and reasoning, multiplying overall output. Simply put, if an AI can perform time-consuming, low-value tasks, it frees up more time for the affected worker to do high-value work.




Figure 7: Task Structure and the Automation or Augmentation of Workers by AI

³ The effect is also referred to as "Cost Disease."

Targeted reskilling and upskilling for workers in occupations vulnerable to automation is one of the most urgent priorities for governments amid AI-driven labor market changes, yet it remains largely underexplored. To date, there is very little empirical evidence on the performance of programs that either (1) upskill workers in AI tools and complementary technologies, which are increasingly in demand across many sectors, or (2) reskill workers into occupations and tasks that are less susceptible to AI-driven displacement. Workforce Innovation and Opportunity Act (WIOA) programs (funding employment, training, and support services to help job seekers succeed in the labor market) in the United States could serve as a model for other governments, with some evidence of increased opportunities and wages for displaced workers (Hyman et al., 2025). However, evidence suggests that workers receiving WIOA-subsidized training who choose to deepen AI-specific skills face a penalty compared to those who pursue general training, indicating that optimal upskilling strategies to capitalize on AI gains remain underexplored.

Workforce development must be informed partly by the starting point of workforce skills: in past transitions, occupations have differed significantly in their ability to transition into new roles based on the content of transferable skills. Figure 8 illustrates this concept with evidence from workers displaced by industrial decarbonization – the horizontal axis of each panel represents the similarity of an origin occupation’s skills to the skills of every occupation with a job opening in the local labor market. We see few job opportunities with high skill similarity for either of the two occupations, but much “faster” returns to employment potential from small changes in similarity requirements for the high-transferability occupation. This result implies that a much greater investment in retraining may be needed to support one occupation than the other.

Similarly, nations undertaking a large-scale re-skilling initiative must evaluate what skills are needed to support that initiative, and to make the trade-off between which occupations are already closest to that skillset (and hence easiest to reskill) versus which occupations may be most vulnerable to technologically induced disruption and most in need of new employment opportunities. Consistent with the labor market resilience perspective in Figure 3, there may be occupations whose exit options are simultaneously subject to AI-powered automation, and a key economic and social objective for national governments is to build




Figure 8: Returns to Training Vary By Starting Occupation (Based on Data from Miles, Combemale, Karplus, and Pistorius 2025)

a talent supply chain that connects this disruption to positive transformation elsewhere in the economy.

New Industries and New Tasks

A historical characteristic of general-purpose technologies is their ability not only to improve productivity broadly but to enable entirely new industries and sets of tasks. The advent of the internet served as a catalyst for the gig economy, paving the road for a set of new careers, including content creators and app developers. Still, many of the most distinctive occupations that emerged from past general purpose technologies matured long after initial technological diffusion (“data scientists” were not added to the U.S. federal occupational taxonomy until 2018), often after the development of platforms that non-technical users could navigate.

New tasks and industries resulting from AI are hard to predict, and they may or may not create demand for human labor. However, by focusing on adoption and diffusion in incumbent industries, we might develop a more principled understanding of the mechanisms that will affect new industries: the economics of task automation, augmentation and compensating demand can help us

to anticipate what type of new work will be performed by humans or AI systems, and hence the potential domain of jobs and needed skills. As workers are augmented by AI (see Figure 7), productivity gains could expand the scale, scope, or customization of tasks in ways that still require human labor. Productivity gains could result in consumer surplus that leads to demand for new tasks. For instance, the first, second and third industrial revolutions drove succeeding rises in leisure and entertainment industries as consumers demanded new tasks with freed time and increased disposable income. In the substitution region of Figure 7, new tasks may emerge that create economic value, but do not increase demand for human labor. For instance, technological advances in computers have created new sets of tasks that were automated from the onset (such as system security, network routing, and binary search). While these new tasks did not increase demand for human labor or create new occupations, they have yielded economic value. Where extensive tasks substitution creates “black boxes” from the perspective of remaining users, continuing productivity improvement may be the domain of the innovator or facilitator layer of the GPT Labor Pyramid, whereas an augmentation scenario may provide more interfaces for users to innovate and experiment.

Labor Strategies for Global Trade

Open questions on adoption, labor market impacts, and training programs all extend to the international context. For high-income countries, AI may accelerate offshoring or reshoring of tasks. AI-powered chatbots have lowered translation costs and barriers to entry in English-speaking markets, potentially enabling some tasks to be performed in countries with lower labor costs. This lever could enable offshoring not only low-skilled tasks (as in the offshoring of telecommunications in the 1980s) but also some high-skilled tasks (such as software development). On the other hand, there is potential for reshoring as tasks that previously relied on low-cost labor overseas could be performed efficiently and cost-effectively domestically with AI tools. Lower-income countries face a “dual vulnerability” problem: large portions of the workforce in emerging economies are employed in sectors highly exposed to automation, and there is often a lack of infrastructure to support and deploy AI systems at scale. For instance, outsourced call centers in India account for approximately 7.5% of the country’s GDP; if these jobs are automatable or reshored to high-income countries,

a key question remains: how to redeploy and reskill the workforce in this sector. At the same time, limited digital infrastructure and shortages of workers skilled in AI tools hinder developing economies’ ability to capture productivity gains from AI.

Many of the least exposed occupations in a country may not belong to industries that produce the most easily tradeable goods and services: for instance, many healthcare jobs are commonly rated as having low exposure to AI, but are also likely to have low tradeability. To participate in global gains from AI diffusion, countries with low exposure, and hence a small AI user layer, face a choice between investing in expanding the user layer (and complementary investments in the facilitator layer) or investing in tradeable goods and services with low exposure, which cannot be substituted for by successful AI-adopting countries. Natural resources may be one example, especially if technological advancements that enhance the productivity of extraction and processing are themselves tradeable (hence importable from other nations). This latter assumption will still depend on a much more well-established standard of global workforce development, namely the need for experts who can apply the new technology – in this respect, maintaining an AI user-base is both an immediate potential productivity boon and an investment in resilience, by creating a pool of talent that may be able to cross-apply AI to new contexts as the technical frontier advances. Countries with strong innovation layers can derive first-order competitive advantage from licensing advanced models, but may derive a much more substantial advantage from the codesign of models with procedures for facilitation and implementation. Innovators of foundational technologies have unique influence over the standards-setting for technology stacks and operations downstream of their technology. Standards can enable others to participate in facilitation, by providing clear benchmarks for buyers to evaluate or coordinate services and complementary technologies – but they can also be used to advantage the particular facilitator composition of a country, for instance by designing standards to leverage skills in which a country has a comparative advantage. AI can also support highly sensitive activities such as semiconductor chip design, and while security concerns are universal, countries with strong AI innovation layers may be better able to identify and mitigate these risks (both from models developed in-country and models licensed from abroad).

The “right” policy approach for economic competitiveness and shared prosperity is made less certain by unanswered questions about adoption patterns, labor market impacts, training and upskilling, and international dynamics. Absent a precise technical roadmap, global competitiveness will

depend on workforce adaptability and resilience, including channels for rapid reskilling and reemployment, especially for national economies competing with AI itself to preserve their service export models.


CONCLUSION

Artificial intelligence represents a General Purpose Technology whose economic impact is not constrained to its region(s) of origin. We can expect the most change to occur in economies capable of widespread deployment and end-use. Historical analogies suggest that the vast majority of productivity gains from GPTs are not captured by the innovators or exporters, but by the countries that effectively deploy them. Each country has a unique industrial and labor composition that influences its susceptibility to AI disruption. This composition also informs which policies a country should adopt in response to the global diffusion of AI. A country’s AI strategy is largely shaped by the share of its economy that is exposed to AI disruption, its relative labor costs, and access to workers with the skills to deploy and maintain AI solutions (AI Facilitators).

Countries like the US have high labor costs and access to an abundance of AI Facilitators (placing it at the upper left of the following figure). Naturally, the optimal policy strategies of the US will differ from those of a country like India, which has comparatively lower labor costs and

access to AI Facilitators. The following table describes a spectrum of policy strategies that might be employed to press a country’s comparative advantages or cover downside economic risk. Countries with deep labor pools of potential AI Facilitators (High Access) are generally better equipped to “outgrow” labor disruptions through investment and growth. However, this requires broad restructuring of the workforce, with some displaced workers reskilling into less-exposed jobs, and remaining users upskilling to make the most effective use of the technology and ensure minimum quality

standards. Conversely, low-labor-cost countries that are net exporters of services face a dual vulnerability and must strategically invest in training programs to either compete directly with high-labor-cost countries or reskill their workforce for less-exposed roles, such as non-routine manual or caregiving occupations. While not all countries are home to AI-exposed industries, those that are cannot sit idle while the rest of the global economy moves forward. Failure to take action now will limit competitiveness and policy options in the future.



⁴ Subsidize training programs directly, or promote training by addressing barriers (e.g., child care, elder care, transportation).

WORKSHOP QUESTIONS

These questions are intended to facilitate discussion in the interactive breakout sessions of the AIAI convening at Carnegie Mellon University, June 24 - 25, 2026.

1. What steps can governments take to measure displacement and transition risk early enough to respond, and what role should public and private partners play?
2. How do we measure actual displacement and transition risk early enough to respond? What information is actionable to your country's government?
3. What role should social dialogue play in enhancing the public acceptance of AI from the User population? What is your government doing to solicit dialogue?
4. How do we think about the exposure of underserved, rural, and low-income communities? How does this exposure differ between these communities in high- and low-income countries?
5. How should education systems build AI literacy from school through adulthood? What is the role of continuing education?
6. How should productivity gains be measured and shared? How does measurement affect the distribution policy that can be taken?
7. Can public sector procurement drive adoption in the private sector?
8. Where does your country sit on our strategy matrix (Figure 9), and what is missing to describe your context?

REFERENCES

- Acemoglu, D. (2024). *The simple macroeconomics of AI* (NBER Working Paper No. 32487). National Bureau of Economic Research. <https://doi.org/10.3386/w32487>
- Acemoglu, D., Anderson, G., Beede, D., Buffington, C., Childress, E., Dinlersoz, E., Foster, L., Goldschlag, N., Haltiwanger, J., Kroff, Z., Restrepo, P., & Zolas, N. (2025). Automation and the workforce: A firm-level view from the 2019 Annual Business Survey. In S. Basu, L. Eldridge, J. Haltiwanger, & E. Strassner (Eds.), *Technology, productivity, and economic growth* (pp. 13–56). University of Chicago Press. <https://doi.org/10.7208/chicago/9780226839097-005>
- Acemoglu, D., Autor, D., Hazell, J., & Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1), S293–S340. <https://doi.org/10.1086/718327>
- Acemoglu, D., & Restrepo, P. (2022). Tasks, automation, and the rise in U.S. wage inequality. *Econometrica*, 90(5), 1973–2016. <https://doi.org/10.3982/ECTA19815>
- AI Adoption Initiative (AIAI). (2025). *AI, everywhere, all at once*. Adopt AI. https://adopt-ai.org/wp-content/uploads/2025/10/AIAI_0325_AI-Everywhere-All-At-Once_Updated-Final.pdf
- Alekseeva, L., Azar, J., Giné, M., Samila, S., & Taska, B. (2021). The demand for AI skills in the labor market. *Labour Economics*, 71, Article 102002. <https://doi.org/10.1016/j.labeco.2021.102002>
- Ales, L., Combemale, C., & Krishnan, R. (2025). Generative AI, adoption, and the structure of tasks. In P. Hacker, A. Engel, S. Hammer, & B. Mittelstadt (Eds.), *The Oxford handbook on the foundations and regulation of generative AI*. Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780198940272.013.0046>
- Anthropic. (n.d.). *Anthropic economic index: Global usage*. Retrieved April 25, 2026, from <https://www.anthropic.com/economic-index#global-usage>
- Autor, D. (2024). *Applying AI to rebuild middle class jobs* (NBER Working Paper No. 32140). National Bureau of Economic Research. <https://www.nber.org/papers/w32140>
- Autor, David H. “Why are there still so many jobs? The history and future of workplace automation.” *Journal of economic perspectives* 29.3 (2015): 3–30. https://pubs.aeaweb.org/doi/pdf/10.1257/jep.29.3.3?mod=article_inline
- Autor, David, David Dorn, and Gordon H. Hanson. *On the persistence of the China shock*. No. w29401. National Bureau of Economic Research, 2021. https://www.nber.org/system/files/working_papers/w29401/w29401.pdf
- Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics*, 118(4), 1279–1333. <https://doi.org/10.1162/003355303322552801>
- Babina, T., Fedyk, A., He, A. X., & Hodson, J. P. (2024). Artificial intelligence, firm growth, and industry concentration. *Journal of Financial Economics*, 155(1), 103–128. <https://doi.org/10.1016/j.jfineco.2023.11.005>
- Baumol, William J. “Macroeconomics of unbalanced growth: the anatomy of urban crisis.” *The American economic review* 57.3 (1967): 415–426.
- Bick, A., Blandin, A., & Deming, D. J. (2025). *The rapid adoption of generative AI* (NBER Working Paper No. 32966). National Bureau of Economic Research. <https://doi.org/10.3386/w32966>
- Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *The Quarterly Journal of Economics*, 140(2), 889–942. <https://doi.org/10.1093/qje/qjae044>
- Choi, E., Bahadori, M. T., Schuetz, A., Stewart, W. F., & Sun, J. (2016). Doctor AI: Predicting clinical events via relation learning. *Proceedings of the 2nd Machine Learning for Healthcare Conference*, 56, 253–266.

Ding, J. (2024). *Technology and the Rise of Great Powers*. Oxford University Press.

Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2024). GPTs are GPTs: Labor market impact potential of LLMs. *Science*, 384(6702), 1306–1308. <https://doi.org/10.1126/science.adj1837>

Ellul, Andrew, Marco Pagano, and Fabiano Schivardi. "Employment and wage insurance within firms: Worldwide evidence." *The Review of Financial Studies* 31.4 (2018): 1298–1340. <https://academic.oup.com/rfs/article-abstract/31/4/1298/4110800?redirectedFrom=fulltext>

Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114, 254–280.

Gans, J. S., & Goldfarb, A. (2026). *O-ring automation* (Working Paper No. 34639). National Bureau of Economic Research. <https://doi.org/10.3386/w34639>

Ganuthula, V. R. R., & Balaraman, K. K. (2025). *Skill-based labor market polarization in the age of AI: A comparative analysis of India and the United States*. arXiv. <https://doi.org/10.48550/arXiv.2501.15809>

G7. (2025, June 17). *G7 leaders' statement on AI for prosperity*. https://g7g20-documents.org/fileadmin/G7G20_documents/2025/G7/Canada/Leaders/1%20Leaders'%20Language/G7%20Leaders%E2%80%99%20Statement%20on%20AI%20for%20Prosperity_17062025.pdf

Georgieff, A., & Milanez, A. (2021). *What happened to jobs at high risk of automation?* (OECD Social, Employment and Migration Working Paper No. 255). OECD Publishing. <https://doi.org/10.1787/10bc97f4-en>

Hartwig, Jochen, and Hagen Krämer. "The 'growth disease' at 50–Baumol after Oulton." *Structural Change and Economic Dynamics* 51 (2019): 463–471.

Hyman, Benjamin G., Brian K. Kovak, and Adam Leive. *Wage insurance for displaced workers*. No. w32464. National Bureau of Economic Research, 2024. https://www.nber.org/system/files/working_papers/w32464/w32464.pdf

Hyman, B., Lahey, B., Ni, K., & Pilossoph, L. (2025). *How retrainable are AI-exposed workers?* (Federal Reserve Bank of New York Staff Reports, No. 1165). Federal Reserve Bank of New York. <https://doi.org/10.59576/sr.1165>

Hui, X., Reshef, O., & Zhou, L. (2023). *The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market* (SSRN Scholarly Paper No. 4527336). <https://doi.org/10.2139/ssrn.4527336>

Katz, L. F., Roth, J., Hendra, R., & Schaberg, K. (2022). Why do sectoral employment programs work? Lessons from WorkAdvance. *Journal of Labor Economics*, 40(S1), S249–S291. <https://doi.org/10.1086/717932>

Mäkelä, E., & Stephany, F. (2025). *Complement or substitute? How AI increases the demand for human skills* (SSRN Working Paper No. 5153230). <https://doi.org/10.2139/ssrn.5153230>

Massenkoff, M., & McCrory, P. (2026). *Labor market impacts of AI: A new measure and early evidence*. Anthropic. <https://www.anthropic.com/research/labor-market-impacts>

McElheran, K., Li, J. F., Brynjolfsson, E., Kroff, Z., Dinlersoz, E., Foster, L., & Zolas, N. (2024). AI adoption in America: Who, what, and where. *Journal of Economics & Management Strategy*, 33(2), 375–415. <https://doi.org/10.1111/jems.12552>

Miles, J., Combemale, C., Karplus, V. J., & Pistorius, P. C. (2025). Simulating regional workforce impacts of decarbonizing integrated steelmaking. *Proceedings of the National Academy of Sciences*, 122(19). <https://www.pnas.org/doi/pdf/10.1073/pnas.2419294122>

Neumark, D., & Simpson, H. (2015). Place-based policies. In G. Duranton, J. V. Henderson, & W. C. Strange (Eds.), *Handbook of regional and urban economics* (Vol. 5, pp. 1197–1287). Elsevier. <https://doi.org/10.1016/B978-0-444-59531-7.00018-1>

Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192. <https://doi.org/10.1126/science.adh2586>

Pizzinelli, C., Panton, A. J., Tavares, M. M., Cazzaniga, M., & Li, L. (2023). *Labor market exposure to AI: Cross-country differences and distributional implications* (IMF Working Paper No. 2023/216). International Monetary Fund. <https://doi.org/10.5089/9798400254802.001>

Schmidt, L., & Contributors. (2024). *How artificial intelligence impacts the US labor market: Evidence from the micro-data* (Working Paper). MIT Sloan School of Management.

Singla, A., Sukharevsky, A., Hall, B., Yee, L., & Chui, M. (2025). *The state of AI in 2025: Agents, innovation, and transformation*. McKinsey & Company. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>

Soares Martins Neto, Antonio; Liu, Yan; Khunara, Saloni; Porras Lopez, Juan Manuel. *Click, Code, Earn : The Returns to Digital Skills (English)*. Policy Research Working Paper;RRR;Digital Washington, D.C.: World Bank Group. <http://documents.worldbank.org/curated/en/099451102172652767>

Venkateswaran, V. (2023). *The impact of AI on the labor market* (SSRN Scholarly Paper No. 4615324). <https://doi.org/10.2139/ssrn.4615324>

Yiu, S., Seamans, R., Raj, M., & Liu, T. (2026). *Worker repositioning and technological change: Evidence from an online labor market* (Wharton School Research Paper). SSRN. <https://doi.org/10.2139/ssrn.4769321>

Zolas, N., Kroff, Z., Brynjolfsson, E., McElheran, K., Beede, D. N., Buffington, C., Goldschlag, N., Foster, L., & Dinlersoz, E. (2020). *Advanced technologies adoption and use by U.S. firms: Evidence from the Annual Business Survey* (NBER Working Paper No. 28290). National Bureau of Economic Research. <https://doi.org/10.3386/w28290>

