

THE AI LABOR STACK

RETHINKING THE WORKFORCE FROM THE GROUND UP

By Dustin Ferrone, Lindsey Lacey, Christophe Combemale,
Edward Teather, and Kevin Allison

June 2026

Carnegie
Mellon
University



EXECUTIVE SUMMARY

Over the past 18 months, the AI Adoption Initiative (AIAI) has convened policy discussions in over 20 cities on five continents to discuss how government policy can accelerate the adoption of economically useful artificial intelligence (AI). These meetings revealed a broad expert consensus that workforce skills policy will be an important part of successful national AI strategies.

AI is a general-purpose technology (GPT) with the potential to unlock productivity gains by automating work that previously required human input. During past technology transitions, countries that benefitted the most from new general-purpose technologies were those that achieved broad adoption the fastest. At the same time, AI has the potential to simultaneously augment, compete with, or even replace human labor across various occupations.

The need for policies that help workforces acquire skills for capturing the upsides of AI, while simultaneously adapting to potential labor market disruption is clear. However, AI workforce and skills policies are still at an early stage of development. Governments face significant uncertainty identifying which policies will best help their populations seize the opportunity.

On June 24-25, the AI Adoption Initiative and Carnegie Mellon University convened an international group of experts on Carnegie Mellon's campus in Pittsburgh to discuss how governments can best respond to this moment. The choice of venue for this discussion carried deliberate meaning and intent: once built around coal and steel, Pittsburgh has navigated structural economic change comparable in scale to how many places are now thinking about AI. Its renewal - based on robotics, advanced manufacturing, healthcare innovation and education - is a reminder that positive industrial transitions are shaped not by technology alone, but by institutions, skills and local capacity.

This paper introduces an analytical framework for assessing AI workforce skills policies, which will serve as a starting point for the workshop discussion. We argue that there are three distinct types of workers that will be important for the AI transition: 1.) **innovators**, who develop the underlying AI infrastructure and models; 2.) **facilitators**, the technical workers who deploy, integrate, and otherwise enable the use of AI, primarily using software tools; and 3.) **users**, whose work is altered by applying AI across the economy. Together, these workers constitute the **AI Labor Stack**. A country's optimal workforce skills policies will depend on its access to workers at different layers of the stack and its relative labor costs.

A central implication is that facilitators are often the "missing middle" in national AI strategies. In the first wave of national AI strategies that began in 2017, many countries focused on frontier innovation (innovators) or broad AI literacy (users). The next wave will need to pay greater attention to the facilitators who translate AI capability into practical deployment across firms, sectors, and public services.

The Pittsburgh discussion broadly validated the applicability of this framing and suggests that the facilitator layer in particular may be under-emphasized in other policy discussions. Participants highlighted the need for a more granular taxonomy of facilitator profiles, better methodologies for measuring facilitator access by country, region and sector, and more adaptive skills infrastructure capable of responding to changing employer, sectoral and local needs."

Drawing on the Pittsburgh workshop, this paper identifies emerging policy-design principles and a follow-on research agenda: how to train facilitators at scale; how to build continuous AI readiness among teachers and users; how to improve data for policymakers and economists; and how to combine labor-market dynamism with transition resilience.

CONTENTS

Executive Summary	2
Contents	3
Introduction: The AI Labor Stack	4
Tier 1 - Innovators	6
Tier 2 - Facilitators	7
Tier 3 - Augmented & Automated Users	9
AI, labor costs and the global market for services	10
A Framework for Workforce Skills Design for AI	12
From strategy to implementation: applying the AI Labor Stack to policy	13
Workforce Skills Policy Case Studies	17
Innovators: frontier talent and advanced technical capacity	17
Facilitators: the missing middle between AI capability and adoption	18
Users: broad AI literacy, access and complementary skills	19
Full-stack resilience: supporting workers through transition	19
Follow-on Agenda	20
Workshop Participants List	21
Further Reading – Workshop After-Action Report and Companion Paper	23
References	24

INTRODUCTION: THE AI LABOR STACK

AI differs from previous technological advancements in its ability to automate a wide range of high-skill, less-routine cognitive tasks. Despite ongoing improvements in AI capabilities, there is significant uncertainty about how quickly AI will diffuse across economies, and what the labor market impacts of AI will be.

This poses challenges for governments attempting to devise policies to accelerate the adoption and diffusion of productivity-enhancing AI applications, while simultaneously equipping workforces to navigate potential disruptions as computers take on cognitive work previously performed by human workers.

Previous GPTs (e.g., the steam engine, electricity, and the microprocessor) were foundational innovations that spurred complementary advances across industries, reshaping productivity, business models, and the workforce. In these previous cases, diffusion of GPTs depended not only on highly skilled **innovators** who pushed the technology's capability frontier forward; but also on **facilitators**, who helped deploy, integrate, and otherwise enable the new technology's use; and end **users**, whose work was altered by applying the new technology in offices, factories, and elsewhere across the economy.

For example, Ding has argued that the United States's advantages in the Second Industrial Revolution (1870–1914) may be attributable in part to greater investments in technical and engineering education. This included the establishment of mechanical engineering departments at university campuses, which expanded the pool of workers capable of adapting and implementing new technologies like machine tools across diverse manufacturing settings. By training large numbers of facilitators and users who could effectively deploy and use such technologies to solve production problems, the U.S. was able to diffuse the benefits of mechanization more rapidly across its economy than countries that remained more focused on the innovation frontier (Ding, 2024).

AI innovators, facilitators, and users can be visualized as an inverted pyramid, or the **AI Labor Stack**. At the bottom of the stack, a small number of innovators with elite technology skills use large amounts of data and computing power to design and build the underlying AI infrastructure and models. In the middle and upper tiers of the stack, respectively, facilitators help users – who constitute the vast majority of affected workers – use AI to perform work inside a wide variety of firms operating in many different economic sectors.


Figure 1: The GPT Labor Pyramid



These different tiers of the AI Labor Stack correspond to the Cloud Stack framework outlined in the AI Adoption Initiative's October 2024 paper, [*AI, Everywhere, All At Once*](#). That paper argued that while a small number of elite firms might build underlying AI models directly on top of raw compute (the “computer infrastructure” layer in the diagram below), most people would use AI to perform useful work via AI applications, accessed largely through enterprise software running on cloud computing infrastructure. In between this underlying infrastructure

and end-user applications, software-based systems (the “AI tooling” layer) would enable the design and deployment of AI systems that leverage underlying models, domain-specific knowledge and data, and other tools. These systems may use a variety of models, depending on the application.

Figure 2: How the AI Labor Stack maps to the AI Cloud Stack



In the AI Labor Stack, the highest degree of technology-specific expertise is concentrated at the bottom of the inverted pyramid, among a relatively small number of innovators with elite technical skills, while the vast majority of affected workers sit at the top. Each level of this labor pyramid must be available and accessible before AI can diffuse widely and deeply across an economy.

This facilitator layer has often been the “missing middle” in national AI strategies. Many countries have focused either on frontier innovation — compute, models, research talent and start-ups — or on broad AI literacy for citizens and workers. Both are important, but neither is sufficient for economy-wide adoption. The next wave of AI strategies needs to pay greater attention to the technicians, systems integrators, cloud engineers, data specialists, evaluators, sector experts and organizational change leaders who translate AI capability into practical deployment. Without this facilitator capacity, frontier models may remain underused, and broad user populations may lack the support needed to apply AI safely and productively in real workflows.

The Pittsburgh workshop tested this framework with a deliberately international and cross-sectoral group of participants from academia, industry, civil society, public-policy bodies, workforce organizations and local Pittsburgh institutions. Participants broadly found the AI Labor Stack useful as a way of distinguishing between the needs of innovators, facilitators and users, with particular interest in how governments can identify, train and measure access to facilitators.

Below, we describe how innovators, facilitators, and users work together across the AI Labor Stack to drive adoption of AI. We outline how domestic labor costs and the global market in services may further interact to drive workforce impacts as AI diffuses across economies. Based on these two major variables, we then offer a strategic framework for workforce policy design. We offer a series of case studies of effective workforce policies. Finally, we identify a series of open research questions to help guide the workshop’s discussions about how governments can design resilient workforce and skills policies for the AI age.

AI workers (those who develop, facilitate, or use the technology) can be thought of 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.

Users comprise the largest share of workers and will fall on a continuum in their day-to-day use of AI-powered tools. For example, lawyers may use more tools that embed AI than construction workers, or work with tools that use AI more explicitly. However, generally neither lawyers, nor construction workers, nor other users will have primary responsibility for managing, integrating, or deploying those tools.

Globally, a smaller number of **facilitators** must have greater AI expertise than the users, while an even smaller number of **innovators** will build AI models, the underlying software tooling, and the digital infrastructure that powers the use of AI in industry-specific applications.

Whereas innovators may number in the thousands, facilitators that leverage AI tooling and maintain critical infrastructure to help users deploy AI applications may be employed in the millions. These facilitators, in turn, can empower potentially billions of users, who directly apply AI in specific industry contexts to perform economically useful work.

The economic returns to AI technology are ultimately generated by diffusion of the technology from innovators, through facilitators, to end-users at the top of AI Labor Stack. Work in this tier of the inverted pyramid will be performed either by AI-augmented human users, or via automation of tasks traditionally performed by human workers. Understanding how these layers work together, and how the availability and cost of labor at the different layers of the stack may vary by country, can help inform workforce policy design.

Tier 1: Innovators

Among these three tiers, innovators will constitute the smallest labor pool. Workers in this tier have elite technical skills and are likely to be globally mobile. Innovators should also be expected to command a wage premium relative to other roles in this labor pyramid, making them expensive to employ. These experts actively advance the frontier of what is technically feasible with AI technology.

Workers at this tier might build underlying infrastructure, AI models, or software tooling that requires specific programming expertise, but without a specific end customer or end-user application in mind. Products and services produced by innovators could include, for example, new LLMs, agentic frameworks, chips for AI training and inference or memory, lithography machines, or semiconductor fabrication plants.

A relatively small proportion of the world's geography is likely to compete at this elite level of the AI Labor Stack: Taiwan, for example, is expected to grow at its fastest pace in 16 years in 2026¹, at 9.6 percent, driven by AI-related demand for its expertise in fabrication plants, while in London — birthplace of the world's first major AI lab in DeepMind — it is estimated that AI innovator companies could occupy 43 percent of unleased office space currently under development in the city's center by 2033.²

While the advancements innovators achieve may be impactful, the global trade in services means that their absence from a specific economy does not necessarily preclude broader AI adoption across that economy. For example, while the majority of top AI model developers are based in the U.S. and China, Israel and Singapore have seen the greatest level of AI utilization per capita.³ This is a phenomenon with historical precedent; Ding (2024) outlines several historical examples of GPTs being developed by innovators in one country, only to achieve greater diffusion in another. Although, as discussed later in this report, reliance on highly skilled foreign labor can create new dependencies.

¹ Reuters. "Taiwan raises 2026 GDP growth outlook to 16-year high on strong AI demand." 29 May 2026

² CBRE. "London could see 4 million sq ft of AI-led office take-up by 2033." 27 April 2026.

³ Anthropic. "Anthropic Economic Index report: Uneven geographic and enterprise AI adoption." September 2025.

Tier 2: Facilitators

Facilitators are the technicians and engineers who make the deployment and sustained end-use of AI-powered tools possible. 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 end users. Tasks such as data management, validation, quality analysis, and other functions that require deep contextual knowledge may be more likely to be performed in-house than the equally necessary work performed by IT workers or systems integrators. Facilitators will use software-based tooling to create, test, and validate the performance of end-user AI applications, but are less likely to use those applications themselves. Fewer end-use firms will directly hire other facilitators responsible for maintaining data centers, telecommunications infrastructure, or performing other roles divorced from the end user's context.

The common thread tying these facilitator 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 (innovators) nor using the technology to generate a final good or service (users). Many of the services rendered by facilitators may be accessed through the international trade in services. However, this solution introduces other inefficiencies (e.g., system reliability, system latency, language barriers, geopolitical-driven technology or data-access restrictions, and regulatory hurdles).

For companies, accessing facilitators through the global market for services may be most attractive for roles that do not require intimate knowledge of the end user's context. Businesses may be more likely to rely on the domestic labor supply for tasks that do require intimate knowledge of certain types of data or domain-specific or firm-specific operational or business processes, or more nuanced cultural knowledge that might not be as easily accessible outside of a country's borders.

In this tier, we should expect to see some traditional facilitator roles persist (albeit with education and upskilling required to keep up with technological advancements). We should likewise expect the



CASE STUDY: FACILITATORS IN SCIENTIFIC DISCOVERY

Recent work from Anthropic on agents in biology⁴ illustrates why facilitator capacity matters for scientific AI adoption. In a case study using NCBI Virus, a portal enabling access to viral nucleotide and protein sequence data, scientific research agents were asked to retrieve data used by virologists for surveillance and diagnostic assay development. Even strong models did not consistently achieve the accuracy needed for reliable dataset construction. Performance improved dramatically only when researchers added a deterministic retrieval layer that made the underlying biological data infrastructure easier for agents to query and validate.

The lesson is that progress in scientific AI will depend not only on more capable foundation models, but also on the facilitator layer around them: data infrastructure, retrieval tools, domain-specific workflows, validation methods, and experts who can translate scientific intent into reliable agentic systems.

⁴Anthropic. "Paving the way for agents in biology." June 2026.

emergence and increasing prominence of new facilitator roles, as the underlying technology and the software-based tooling that is used to create specific AI applications evolve. For example, the advent of agentic AI systems is already creating demand for human workers who can effectively design and manage multi-agent systems to accomplish specific tasks. The following table provides some examples.

Table 1: Facilitators
Examples

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

In the AI Labor Stack, facilitators form the connective membrane between innovators and users. They absorb and interpret advances from the innovation frontier, then adapt, validate, integrate, and maintain those capabilities so they can be safely and productively used in specific organizational and sectoral contexts.

Recent findings that upskilling of existing ICT workers to help them acquire new AI-specific skills is the strongest single predictor of successful firm-level adoption⁵ further underscores the importance of facilitators for broad AI diffusion. Part of their contribution is to broaden the set of use cases where AI can helpfully be applied. In one analysis of UK task exposure, the Institute for Public Policy Research estimated that “out-of-the-box” LLMs augmented around 11 percent of tasks, but that exposure rose to around 59 percent of working hours when considering “integrated AI” systems connected to proprietary data, databases and task-execution capabilities.⁶

However, based on feedback from the Pittsburgh workshop, the facilitator layer should be understood as a family of roles rather than a single occupational category (or exclusively ICT-focused roles). The following provisional taxonomy reflects themes raised in the discussion and should be refined through future research, country-level workshops and measurement work:

- **Data facilitators:** data architects, data engineers, data annotators, data curators and data-quality specialists who make data usable for AI systems and sector-specific applications.

- **Technical deployment facilitators:** cloud engineers, systems integrators, model evaluators, cybersecurity specialists, forward-deployed engineers and application builders who connect models, tools and infrastructure to real workflows.
- **Sector and domain facilitators:** domain experts who understand how AI can be applied in specific sectors such as manufacturing, scientific discovery, healthcare, finance, public services and SMEs.
- **Governance and assurance facilitators:** AI governance, risk, regulatory, evaluation, assurance and audit specialists who help organizations adopt AI safely, responsibly and in compliance with emerging rules. Cisco’s AI Workforce Consortium analysis reports a 150 percent increase in demand for AI governance skills, underlining the growing importance of this profile.⁷
- **Managerial and cultural facilitators:** middle managers, team leads and internal champions who create permission to experiment, identify new use cases, build trust upwards and downwards within their organization, and translate AI strategy into changed work.
- **Community and training facilitators:** local institutions, workforce organisations, SME advisers, peer networks, public-sector champions and training providers who help users and organisations build the confidence, trust and practical capability needed to adopt AI. These actors may not always be deeply technical, but they perform a facilitator function by translating AI capability into usable practice across communities, firms and public services.

⁵ Ouimette, Teather and Allison. “AI, Everywhere, All At Once”. 2024.

⁶ Jung, C., & Srinivasa Desikan, B. “Transformed by AI: How generative artificial intelligence could affect work in the UK — and how to manage it”. IPPR, March 2024.

⁷ AI Workforce Consortium. *ICT in Motion: The Next Wave of AI Integration*. Cisco, September 2025. The report finds that demand for AI Governance skills increased by 150 percent and AI Ethics skills by 125 percent across the 50 ICT roles analysed.

Tier 3: Augmented & Automated Users

Workers in this tier of the AI Labor Stack 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.

Exposure-driven automation or augmentation can lead to productivity gains at the industry and national level. However, automation may lead to global realignments in comparative advantage, with services that become more AI-intensive being redistributed to countries with high access to capital.

Outcomes for AI users are likely to depend on which of their tasks are exposed (for example, whether a task is feasible to automate or accelerate with AI-powered tools); the share of a worker's time spent performing exposed tasks; and the market demand for their remaining, unexposed, tasks (e.g., how much more valuable work can they perform once an exposed task takes up less of their work day).

Early worker-reported evidence suggests that this time-saving channel is already materializing: in the Federal Reserve's newly published 2025 Survey of Household Economics and Decisionmaking, one in four U.S. workers reported using generative AI at work in the prior month, and 81 percent of those users agreed that it saved them time.⁸

How AI technologies, such as LLMs, 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 of AI applications. As we argue in the following section, a key question for policymakers is whether workers will experience an increase in compensating demand for certain types of labor as AI automates certain tasks, or replaces entire categories of human workers.

When a task requires less human time, affected workers gain greater availability to perform other tasks. The ultimate impact on workers in a particular occupation will depend on how much compensating demand exists for the remaining, unaltered tasks. At one extreme, 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 worker's remaining tasks increases. As AI-augmented workers become more efficient, allowing them to create more products and services, customers demand more of their products or services.

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 at a given level of technology development. For past technologies, there was also a minimum task complexity that was economically viable to automate. For example, meatpacking still disproportionately employs low-wage workers performing manual tasks.

CASE STUDY: AI USERS IN FARMING (CROPX)

CropX⁹ illustrates how AI adoption at the user layer can reach workers far from the frontier AI labs. It offers an easy-to-use agronomic farm management platform and mobile app that helps farmers make better operational decisions using data from their own fields. With sensors and a smartphone, farmers can monitor soil and crop health, and manage irrigation, fertilisation and crop protection without needing to become AI specialists.

This example shows what broad AI diffusion can look like in practice: AI embedded in sector-specific tools that help users improve productivity in everyday work. CropX reports that its platform has supported around 12,000 farms, helping users save 50 percent in water while increasing yields by 20 percent. For policymakers, the lesson is that user-layer adoption depends not only on AI capability, but on adapted tools that are easy to integrate into existing workflows.

⁸ Board of Governors of the Federal Reserve System. Economic Well-Being of U.S. Households in 2025. May 2026.

⁹ About Amazon. "How agricultural firm CropX uses AWS cloud computing for farming." 17 January 2023.

AI, labor costs and the global market for services

The three tiers of the AI Labor Stack do not operate in national isolation. The growth of international trade in services means that innovators, facilitators, and users in one country compete with and supply services to their counterparts in another. A country's relative labor costs are therefore a central variable in workforce policy design. The policy implications differ sharply across each tier of the AI Labor Stack.

The basic dynamic is not new. Beginning in the 1980s, advances in telecommunications allowed firms to route customer calls to centers in India and the Philippines, where labor was cheaper. This triggered a wave of service-sector offshoring that became an important economic ladder for lower-income economies. AI is creating an analogous inflection point, but whereas previous innovations in telecommunications relocated those jobs, AI may eventually automate them away entirely.

Workers in the user layer of the AI Labor Stack are most directly exposed to AI's disruptive effects on global labor markets. For lower-income economies, there is a risk that AI may substitute for service occupations that have historically been offshored, putting pressure on jobs that served as crucial rungs on the development ladder for younger, growing populations.

Pizzinelli et al. (2023) find that AI exposure in lower-income countries is concentrated among professional and managerial workers, not the larger low-skill employment base. In India, the least-exposed 40 percent of workers account for 70 percent of total employment, whereas the same exposure level accounts for only 25 percent of employment in the UK, for example.

For many lower-income economies, the near-term threat of large-scale displacement may be more limited than feared, but so is the near-term opportunity for AI-driven productivity gains to spread broadly. Ganuthula and Balaraman (2025) identify a compounding vulnerability: lower-income countries tend to have both high concentrations of automatable low-skill jobs and weaker infrastructure and human capital to support AI adoption. These countries may face elevated risk from automation while being less positioned to capture the offsetting benefits.

For countries looking to leverage cost advantages in the global AI economy, focusing on the facilitator tier of the AI Labor Stack may be beneficial. Countries with competitive labor costs can build out facilitator workforces, including system integrators, data validators, AI application testers, platform administrators and other facilitator roles, and export those services internationally. This would extend the development model that made call centers viable in the telecommunications era.

However, not all facilitator roles are equally tradeable. Roles requiring deep knowledge of an end-user's specific context, such as monitoring AI systems for locally nuanced bias, enforcing country or other location-specific safety standards, or managing firm-specific data pipelines and workflows, are harder to export. Regulatory barriers, data localization requirements, and latency constraints may add further friction. Countries need to distinguish between facilitator roles that can be oriented toward global service markets and those that are better positioned as drivers of domestic AI adoption.



Both pathways have value. Exported facilitator services generate foreign exchange and diversify economic resilience. Domestically embedded facilitators drive the productivity gains that benefit the broader economy. Upskilling existing ICT workers to acquire AI-specific skills has already been identified as the strongest single predictor of successful firm-level AI adoption (AIAI, 2024) – a starting point available to countries across the income spectrum.

At the innovator level of the stack, high-income countries with strong research infrastructure have an opportunity to consolidate first-mover advantages in frontier AI development. But because innovators are globally mobile, the economic returns from frontier investment are not guaranteed to accrue domestically. High-income countries may also risk a structural gap if they focus solely on innovation. To fully capture the economic potential of AI, they must also invest in the

facilitators and users who translate frontier technology into practical, widespread application.

The AI Labor Stack framing makes clear that workforce policies optimized for a single tier are likely to underperform. All three tiers are important, and present different opportunities, depending on a country's starting position.

Countries with low labor costs and young populations face an asymmetric challenge: their greatest near-term vulnerability is substitution of low-cost human labor with synthetic or machine labor, while their greatest opportunity may lie in building exportable facilitator capacity.

High-income countries, conversely, may face a temptation to over-invest in the innovation frontier at the expense of the facilitator and user tiers that will ultimately determine how much of that frontier gets adopted to perform economically useful work.



A FRAMEWORK FOR WORKFORCE SKILLS DESIGN FOR AI

The above discussion suggests that policymakers should factor in a country's relative access (or potential access) to innovators, facilitators and users as well as relative labor costs when designing workforce policies for AI. The following discussion considers how countries might employ different strategies to address different opportunities and risks, depending on where they sit in this matrix. While not all countries are home to AI-exposed industries, those that fail to take action may limit their competitiveness and policy options in the future, as AI refashions the global competitive landscape.

Countries like the U.S. have high labor costs, but also have access to large numbers of facilitators, placing them at the upper left of the following figure. This is different from a country like India, for example, which has comparatively lower labor costs and (at present) lower numbers of facilitators, placing it in the bottom right quadrant. This likewise explains why a core part of India's AI workforce strategy is to increase the number of facilitators available in its economy. India's **Future Skills Prime Initiative**, for example, provides incentives for ICT workers to participate in a microcredential system that aims to increase the number of workers with the skills required for both facilitator and user roles over time. The system is adaptive, meaning it adjusts the skills it certifies as the underlying technology changes.

Countries with high access to facilitators may be 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.

Low-labor-cost countries with sufficient access to facilitators and a robust outsourcing industry may 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 (and therefore likely harder-to-automate) manual or caregiving occupations.

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, for example, 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).

Conversely, there is potential for reshoring as tasks that previously relied on low-cost labor overseas that could be performed efficiently and cost-effectively domestically with AI tools. This suggests that 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 may be a lack of infrastructure to support and deploy AI systems at scale.

For instance, service exports (such as customer service call centers and tech support) are a major pillar of the economy in the Philippines (OECD 2026), but are highly-exposed to automation due to their high-volume and low-complexity nature. This vulnerability is compounded by access barriers; in the 2025 Government AI-Readiness Index, the Philippines received a score of 0/100 on a compute capacity index (which measures public cloud providers with AI-deployment chips) and 27/100 on a human capital index (measuring formal technical AI training in higher education). High exposure due to reliance on service exports, and low access to facilitators (human capital) and infrastructure (compute) in the Philippines highlight the problems that low-income countries may face in insulating against disruptions from AI.

Figure 3: Policy strategies by National Labor Cost and Skills Availability



From strategy to implementation: applying the AI Labor Stack to policy

The strategy matrix above can help policymakers assess how a country's access to AI facilitators and relative labor costs may shape its workforce policy choices.

The AI Labor Stack framework provides a starting point for assessing which interventions are needed to strengthen innovators, facilitators and users, and what transition support is needed where workers cannot be successfully augmented. This matters, because many existing AI skills policies remain focused on advanced AI skills such as model development and deployment. For countries that cannot realistically compete at the frontier, or whose immediate challenge is economy-wide adoption, investment in the facilitator and user layers are likely to be more important. As was argued in the first AI Adoption Initiative paper, *AI, Everywhere, All At Once*: "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."¹¹

The Pittsburgh workshop broadly validated this framing. Participants welcomed the distinction between frontier AI talent, facilitator capacity and broad user readiness, while emphasising that skills policy will need to become more empirical, sector-specific, local, adaptive and ambitious if it is to support broad AI adoption. The discussion also reinforced that the facilitator layer should not be treated as a residual category between elite AI research and general AI literacy. It includes a range of technical, sectoral, organisational and governance roles that are likely to determine how effectively AI capability is translated into productivity gains across firms, public services and local economies.

¹⁰ Subsidize training programs directly, or promote training by addressing barriers (e.g., child care, elder care, transportation).

¹¹ Quimette, Teather, and Allison. AI Adoption Initiative. "AI, Everywhere, All At Once". 2024

Against this backdrop, nine design considerations follow for adapting skills policy to support workforce transition and AI diffusion across the economy:

1. Build the empirical baseline for workforce and skills policy planning

The workshop repeatedly returned to the need for better data. Economists and policymakers are being asked to plan for AI's effects on wages, productivity, employment, entry-level roles, sectoral demand and geographic labour markets, but current data sources are often too slow, incomplete or poorly structured to support confident forecasting.

Participants argued that governments need a stronger empirical baseline on AI adoption: which firms are using AI, which tasks are changing, which workers are affected, where productivity gains are showing up, and where displacement risks may be emerging. A stronger empirical baseline should also help governments and workforce institutions map the supply and demand for facilitators: which skills profiles are emerging, where they are located, which sectors need them most, and where geographic or sectoral bottlenecks may constrain adoption.

A key theme was that anecdotes are moving faster than data. That creates a risk of overreaction, underreaction or poorly targeted policy. Better labour-market information, more timely employer signals, common data standards and anonymised industry data-sharing could help economists build more useful models and help policymakers design interventions that are responsive rather than speculative. The New Delhi Frontier AI Commitment on the sharing of industry data on AI usage at an aggregated, anonymised level is likely to be an important contribution to this effort, with companies set to report on their progress by the next global AI summit in Switzerland.

2. Use sectoral perspectives to focus skills interventions

AI skills interventions are likely to be more effective when they are tied to sectors with clear adoption pathways, strong labor demand and credible scope for productivity-enhancing augmentation. Evidence from previous sector-focused employment programs suggests that targeted training linked to occupational clusters can generate persistent earnings gains of around 12–34 percent, particularly when combined with wrap-around

services and soft-skills support. For AI, this suggests that governments should identify sectors where AI exposure, demand elasticity and facilitator capacity overlap.

3. Design skills systems around local institutions

AI diffusion is likely to be geographically uneven, making local institutional capacity an important part of AI workforce policy.

The Pittsburgh discussion highlighted several models for building this kind of capacity: existing institutions adapting to the AI moment, as with the ARM Institute and the University of Cambridge's AI for Science programme; or new institutions being built out where they are needed, as in India's roll-out of 500 IndiaAI Data and AI Labs to expand practical data and AI skills capacity through local institutions, including training in areas such as data annotation, data curation, data science and AI technology.

National AI strategies should therefore take account of local adoption systems: employers, education providers, industry bodies, cloud providers, universities, community colleges, local government, workforce agencies and trusted intermediaries. The policy question is not simply how many people can be trained, but whether the right institutions exist, or can be adapted, to connect AI capability to practical adoption.

4. Treat facilitator capacity as an explicit policy objective

The facilitator layer should be treated as an explicit object of policy, not as a residual category between frontier AI talent and broad AI literacy. Facilitators are the workers and institutions that translate AI capability into practical changes in workflows, organisations and sectors. Without them, countries may have access to powerful models and broad user awareness, but still



struggle to turn AI into broad productivity gains.

5. Make adult skills systems modular, updateable and connected to employer demand

The pace of AI development makes it difficult for traditional degree programmes or one-off training courses to keep up with changing employer demand. Governments and employers should therefore support shorter, stackable and employer-recognised pathways that can be combined and refreshed over time.

This applies across the stack. Users may need basic AI literacy, confidence, verification skills and an understanding of appropriate use. Facilitators may need applied AI, workflow redesign, data management, cloud, cybersecurity, model evaluation, assurance, systems integration and sector-specific deployment skills. Innovators may need advanced technical education, compute access, research training and routes into frontier work.

India's FutureSkills Prime provides one example of this logic. It uses an adaptive microcredential model to help ICT workers acquire certifications in AI, cloud computing, cybersecurity, robotics and related emerging technologies, combining industry alignment with national implementation infrastructure. The workshop discussion reinforced the value of this kind of approach, while also suggesting that credentials may need to become more dynamic. In fast-moving domains, credentials may increasingly require refresh cycles, renewal requirements or expiry dates, rather than being treated as permanent signals of competence.

6. Build the human and organisational capabilities that allow AI to be adopted well

AI adoption is not only a technical process. It depends on judgement, communication, ethical reasoning, domain expertise, data interpretation, workflow design and the ability to coordinate human and machine work. These complementary skills matter at the user layer, but they also matter for facilitators who need to help organisations redesign work safely and productively. The workshop placed particular emphasis on middle

managers and organisational leaders. They are often the people who translate strategic intent into changed practice: giving teams permission to experiment, managing risks, building trust, listening to user concerns, and identifying where AI can improve workflows rather than simply being added on top of existing processes. This is especially important where workers are anxious about AI, where adoption requires changes to long-established routines, or where trust and accountability are central to the work.

For policymakers and employers, this means that AI skills policy should not be limited to technical training. It should also build the human and organisational capabilities that allow AI to be adopted well.

7. Prepare the next generation for AI-enabled work

AI raises important questions for how the next generation will learn, work and develop judgement. Students and future workforce entrants will need to understand how to use AI tools, but also how to evaluate outputs, challenge assumptions, verify information, and recognise when human judgement should remain central.

This places new demands on teachers. As AI becomes more capable and more widely used by students, teachers will need a stronger working understanding of the technology: where it can help, where it remains unreliable, how risks such as hallucination, overreliance and sycophancy are evolving, and how classroom practice and assessment may need to change.

Workshop participants discussed the current education environment as a "Wild West". It was noted that while around 60% of teachers in the U.S. have had some training in AI, only around 5% have access to ongoing professional learning on AI.¹² That gap matters because AI is changing quickly: one-off training will not be enough for teachers to understand emerging capabilities, limitations, risks and appropriate classroom uses. This points to the need for continuous professional development so that teachers can help students develop critical thinking, verification skills, problem-framing and "show your working" habits for an AI-enabled labour market.

¹² Langreo, Lauraine. "More Schools Are Providing AI Training for Teachers. Is It Any Good?" *Education Week*, 18 May 2026. The article reports EdWeek Research Center survey data showing that, by winter 2026, 42 percent of K-12 teachers had received no training on using generative AI in the classroom, 22 percent had received multiple training sessions, and 9 percent had received ongoing training.

8. Make mass AI upskilling work for every audience

Governments will need to think big in their approach to AI skills. The scale of the AI transition means isolated pilots will not be sufficient, and many governments, companies and institutions are already setting ambitious targets: national upskilling commitments, large-scale teacher training programmes, local networks of institutes, employer-led training systems and sector-focused adoption efforts. The common lesson is that AI workforce policy will need to move from small programmes to systems capable of reaching large numbers of workers, students, jobseekers and firms.

At the same time, scale alone is not enough. Training millions of people in generic AI awareness will not necessarily produce adoption if the training is not relevant to the needs of employers and workers, or if it is disconnected from the sectors, occupations and institutions where AI will actually be used. National AI skills strategies should therefore be evaluated against two tests: reach, meaning whether they operate at sufficient scale across the workforce; and relevance, meaning whether they meet the needs of employers and workers in the parts of the AI Labor Stack where new capabilities are most needed.

9. Build transition resilience into the policy package

AI workforce policy should not only be about adoption and upskilling. It also needs to include resilience for workers whose roles change or disappear as firms reorganise work around AI.

Some workers will use AI to become more productive in their current jobs. Some will move into new facilitator roles. Others may face disruption. A serious policy framework needs to prepare for all three outcomes. This points toward models that combine labour-market dynamism with worker security: income support, retraining, wage insurance, transition services, active labour-market programmes and pathways back into good work.

The most promising framing is not to freeze labour markets in place, but to help workers move through transition with confidence. Denmark's flexicurity model remains a useful reference point: flexibility for firms, security for workers, and active support to reconnect people to employment through training and reattachment.

Taken together, these considerations suggest that effective AI workforce policy should not be measured simply by how many people are trained in "AI," but by whether countries are building the right capabilities in the right layers of the AI Labor Stack, in the sectors and places where adoption will matter most for their particular economic context.


Policymakers at the national, regional and city level should ask: what evidence is needed to plan under uncertainty; which sectors should be prioritised; which complementary skills will help users benefit from AI; which modular pathways can keep pace with the technology; which regions need facilitator capacity; and what transition systems are needed for workers whose roles cannot be successfully augmented?



WORKFORCE SKILLS POLICY CASE STUDIES

Governments are increasingly experimenting with AI skills and workforce policies, but these interventions should not be treated as a single category of “AI skills policy.” Different policies target different bottlenecks in the AI adoption process. Some aim to expand the small pool of innovators who push the technical frontier. Others focus on the facilitators who translate AI capability into deployable tools. A third set of interventions focuses on the broad base of users who need the confidence, access and complementary skills to use AI effectively in everyday work. These capability-building policies also need to be complemented by labor-market resilience measures that help workers move between roles, occupations and sectors as AI changes demand. We identify current examples of interventions across the AI Labor Stack below:

Figure 4: Example policy interventions



Innovators: frontier talent and advanced technical capacity

At the innovator layer, governments have tended to focus on elite research talent, frontier compute, doctoral pipelines, public research institutes, fellowships and high-skill migration routes. Canada's **Pan-Canadian AI Strategy** is one of the clearest examples of this first-generation model. Launched in 2017, it brought together three national AI institutes — Amii in Edmonton, Mila in Montreal and the Vector Institute in Toronto — and uses the Canada **CIFAR AI Chairs** program to attract, retain and support leading researchers. CIFAR reports that more than 125 leading

researchers are currently supported through this program. This is a useful model of how a country with access to global talent and significant capital can build research excellence, talent density and international credibility around the innovation frontier.¹³

The UK has also targeted the innovator layer of the stack. The **AI Opportunities Action Plan** links compute, talent, fellowships, data and migration as mutually reinforcing parts of the frontier ecosystem. It recommends expanding the AI Research Resource by at least 20 times by 2030; creating 15 new **Turing AI Pioneer Fellowships**; committing funding for a further 25 **Turing AI Acceleration and AI World-Leading Fellowships**; and developing new immigration pathways for graduates from top AI-producing institutions, including those not currently covered by the **High Potential Individual visa** route. **Centres for Doctoral Training** and similar cohort-based models fit the same logic: fellowships attract stars, but doctoral pipelines replenish the talent base over time.

Countries with credible research institutions, capital, compute access and an ability to attract global talent can gain strategic advantages from being closer to the frontier. However, a policy approach focused on the innovator layer of the AI Labor Stack is likely to be insufficient for broader AI adoption.

Facilitators: the missing middle between AI capability and adoption

India offers one of the clearest examples of a facilitator-led approach at national scale, combining a national skilling platform with local delivery infrastructure.

FutureSkills Prime, developed by Nasscom and India's Ministry of Electronics and Information Technology, is a prominent example of adaptive microcredentialing for ICT workers. Designed as a technology skilling hub rather than a traditional university programme, it offers pathways across 12 digital technologies and 10 professional-skill areas, including AI, cloud computing, cybersecurity, robotic process automation, IoT and semiconductors. Learners can progress through competency diagnostics, industry-aligned courses, Nasscom certifications, government-backed incentives,

and routes into internships or job opportunities through Talent Connect. The platform reports more than 2.5 million registered learners and over 3,000 courses and pathways, making it a useful example of facilitator policy treated as national implementation infrastructure rather than as a series of small pilots. The Pittsburgh discussion also highlighted IndiaAI's work to build a wider network of local Data & AI Labs and institutes, described by participants as a network of 500 local training institutions, with a focus on applied data capability including data annotation, data curation, AI technology and data science. Taken together, these initiatives provide an illustration of how national platforms, industry-recognised credentials, local delivery and sector-relevant data skills can be combined to build facilitator capacity at scale.

The **ARM Institute**, based in Pittsburgh, offers a more place-based and sector-specific example of facilitator capacity in advanced manufacturing. Created in 2017 after Carnegie Mellon University won the bid to establish a robotics-focused Manufacturing USA Institute, the ARM Institute operates as a public-private partnership across industry, academia and government. The workshop discussion underscored why this matters for AI adoption: manufacturers need fit-for-purpose training that maps to employer demand, prepares workers to work alongside robotics and AI, and helps firms move from technical possibility to practical implementation. In this sense, ARM is not simply a technology-development institution; it is a local and sectoral adoption institution that connects robotics, AI, workforce training, industrial competitiveness and regional economic renewal.

These examples illustrate that facilitator skills are unlikely to be built effectively through degree



¹³ Notably, after having been the first country to publish a national AI strategy in 2017, Canada's revised national AI strategy, launched in June 2026, has now broadened its focus towards preparing the wider workforce for AI transition.

programmes alone. Countries will need modular, stackable and employer-recognised pathways that help existing ICT workers, technicians, engineers, managers and sector specialists move into implementation-intensive roles without leaving the labour market for years at a time.

Users: broad AI literacy, access and complementary skills

Policies targeting the user layer of the AI Labor Stack should aim to equip the broader workforce with sufficient AI literacy and access, while also cultivating skills that can be complementary to AI capabilities.

In addressing demographic imbalances in AI user uptake, there may be useful lessons from earlier digital inclusion efforts. South Korea's internet-era experience, for example, may be particularly instructive. Alongside infrastructure investment and broadband rollout, the government launched practical internet literacy programs for groups at risk of being left behind, including homemakers, older citizens, farmers and military personnel. In 2000, South Korea launched the "Ten Million People Internet Education" project to provide internet education at mass scale. By October 2002, South Korea had reached 10 million broadband subscribers, with around 70 percent of households connected at speeds above 2 Mbps.

The AI analogy is not exact, but the lesson is relevant: broad uptake of a general-purpose digital technology depends not only on infrastructure, but on access, confidence, training and deliberate inclusion for groups least likely to benefit automatically. Recent data from the U.S Federal Reserve indicates the value of exploring these approaches, noting for example that U.S workers with a graduate degree were more than four times more likely to use AI than those with a high school degree

or less, while those aged 30-44 were nearly 2.5x more likely than those over 60 to use AI (and perhaps more surprisingly, more than 50 percent more likely than their younger colleagues aged 18-29).

Preparing the workforce of the future will of course require a focus on school-age education. Estonia provides a current school-system example for the AI era. Its **AI Leap** initiative is rolling out free AI learning tools first to 20,000 students aged 16-17 and training around 3,000 teachers, with plans to expand to around 58,000 students and 5,000 teachers by 2027. The design is notable because it pairs access with teacher training, digital ethics, self-directed learning and educational equity. This builds on Estonia's earlier **Tiger Leap** experience of embedding digital capability in schools, but updates it for a world in which new labor-market entrants will need to know how to use, evaluate and work alongside AI systems.

Full-stack resilience: supporting workers through transition

The innovator, facilitator and user layers of the AI Labor Stack describe the capabilities countries need to build. But governments also need institutions that help workers move between roles and occupations as AI changes labor demand. Some workers will be able to upskill within their current occupations. Others may move into facilitator roles or more productive AI-enabled user roles. Some however may face displacement as firms reorganise their work around AI. Because the balance between augmentation and automation remains uncertain, transition systems should be treated as part of the core policy architecture of AI workforce policy.

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 of policies designed to address this issue. Its central idea is the "golden triangle": flexibility for firms, income security for workers, and active labor-market policies that reconnect unemployed workers to jobs through training, activation and employment support. It allows employers to hire and fire at will, providing labor market flexibility, but also 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.



FOLLOW-ON AGENDA

The AI Labor Stack has been validated as a useful framework for understanding the workforce-policy challenge, but the Pittsburgh discussion showed that the next phase of work must operationalize the facilitator layer. Countries need to understand which facilitators they have, which they lack, where they are located, which sectors need them most, how they can be trained at scale, and how they can support trusted adoption among users.

This follow-on agenda builds directly on feedback from workshop participants and should be considered as directions for future research, measurement and policy testing:

1. **Develop a fuller taxonomy of facilitator roles and skills:** refine the provisional categories above, identify required skill profiles, and map training pathways and policy levers for each profile.
2. **Design a methodology for measuring facilitator access by country, region and sector:** measure supply, accessibility, sector fit, demand, credential quality, adoption outcomes and trust-based diffusion channels.
3. **Run sector deep dives:** begin with scientific discovery and advanced manufacturing, where workshop discussion suggested that AI adoption depends heavily on domain experts, data infrastructure, workflow redesign and sector-specific institutions.
4. **Develop localized versions of the workshop:** test the AI Labor Stack against specific national and regional contexts, including countries with different labour costs, sector mixes, institutional assets and levels of facilitator access.
5. **Build a teacher AI-readiness agenda:** move from one-off AI literacy toward continuous professional learning, assessment redesign, critical-thinking support and judgement formation for young people entering an AI-enabled labor market.
6. **Strengthen the empirical baseline on adoption, tasks and labor-market change:** combine industry data, job postings, employer surveys, course content, hiring outcomes and public-sector administrative data to help economists and policymakers plan under uncertainty.
7. **Develop a transition-resilience playbook:** identify how governments can combine labour-market dynamism with worker security through income support, wage insurance, retraining, employer incentives, reattachment services and social dialogue.



WORKSHOP PARTICIPANTS

We are grateful to all workshop participants for the generosity, expertise and candor they brought to the Pittsburgh discussion. The convening benefited from an unusually rich mix of perspectives, bringing together voices from government, academia, industry, civil society, workforce organizations, international institutions and local Pittsburgh institutions. Participants brought expertise across AI policy, labor markets, education, skills systems, advanced manufacturing, scientific discovery, governance and technology deployment. Just as importantly, the group reflected a wide range of national and regional contexts, with perspectives from North America, Europe, Latin America, India, Australia and beyond. That diversity of experience was essential to stress-testing the AI Labor Stack framework and shaping the policy insights and follow-on agenda set out in this paper.

This paper now reflects revisions made by the authors in light of insights and feedback from the Pittsburgh discussion. Responsibility for the paper's analysis, framing and recommendations remains with the named authors, and the paper should not be read as representing the views of individual workshop participants or their organizations.

Participants are listed alphabetically by surname.

Participants include those who consented to be listed as of the publication date.

Name	Role / organisation
Kevin Allison	Founder, Minerva Technology Futures
Stijn Broecke	Senior Economist, OECD
Qing Chang	Research Data Scientist, University of Pittsburgh
Christophe Combemale	Assistant Research Professor, Carnegie Mellon University
Pelayo Covarrubias	Presidente, Fundación País Digital
Morgan Frank	Assistant Professor, University of Pittsburgh
Katie Hendrickson	Policy Program Manager, Microsoft
Jennifer Kelly	FA and Senior Director Consultant, Prudential and BNI
Wendi Li	PhD Student, University of Pittsburgh
Marsha Lovett	Vice Provost for Teaching and Learning Innovation, Carnegie Mellon University
Miranda Lutz	Senior Manager, Global Public Policy, Cisco
Jessica Montgomery	Director of Accelerate Science, University of Cambridge
Glenn Parham	Head of Public Sector, Vals AI
Michael Pisa	Economic Opportunity Lead, Google
Kenya Porter	Founder and CEO, EXplore Pittsburgh Tech Week
Alireza Javadian Sabet	Ph.D., University of Pittsburgh
Vincent Sha	Associate Dean of IT and Operations, Carnegie Mellon University
Oskar Nordström Skans	Professor of Economics, Uppsala University
Dr. Peta Jeanne Skujins	Senior Manager, Future Skills Organisation
Cassidy Sugimoto	Dean of the College of Humanities and Social Sciences, Carnegie Mellon University
Kartik Suri	General Manager, Future Skills, India AI
Ed Teather	Global Lead, AI Policy Research and Partnerships, Amazon Web Services
Jason Tyszko	Senior Vice President, Policy and Programs, U.S. Chamber of Commerce
Michelle Wang	Student, Carnegie Mellon University
Teraya White	Academic Program Manager, Carnegie Mellon University

FURTHER READING –

After-Action Report and Companion Paper

This paper was originally prepared as a pre-read for the AIAI-CMU workshop in Pittsburgh on June 24–25, 2026. Following the workshop, it was revised to incorporate insights from the discussion, including a broader understanding of the facilitator layer, the importance of local institutions and organisational change, the need for better evidence on AI adoption and labour-market change, the role of ongoing teacher education, and the importance of transition resilience. The accompanying after-action report provides a fuller and more dedicated account of the workshop, summarising the main themes, areas of debate and priorities for follow-on research that emerged over the two days:



A longer companion paper, which sets out the evidence base, analytical framework and country examples in greater detail, is available via the QR code below:



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>

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>

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

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). [pnas.org/doi/pdf/10.1073/pnas.2419294122](https://doi.org/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>

OECD. (2026). *OECD economic surveys: Philippines 2026*. OECD Publishing. <https://doi.org/10.1787/f0e0c581-en>

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>

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>

