CARNEGIE
ENDOWMENT FOR
INTERNATIONAL PEACE

APRIL 2026

The AI Labor Debate: Three Views on the Future of Work

Teddy Tawil

Technology and International Affairs Program

The AI Labor Debate: Three Views on the Future of Work

Teddy Tawil

April 2026

About the Carnegie Endowment for International Peace

In a complex, changing, and increasingly contested world, the Carnegie Endowment generates strategic ideas, supports diplomacy, and trains the next generation of international scholar-practitioners to help countries and institutions take on the most difficult global problems and advance peace. With a global network of more than 200 scholars across twenty countries, Carnegie is renowned for its independent analysis of major global problems and understanding of regional contexts.

Carnegie Endowment for International Peace
Publications Department
1779 Massachusetts Avenue NW
Washington, DC 20036
P: + 1 202 483 7600
F: + 1 202 483 1840
CarnegieEndowment.org

This publication can be downloaded at no cost at CarnegieEndowment.org.

TABLE OF CONTENTS

About the Author

Teddy Tawil is a visiting researcher in the Technology and International Affairs Program at the Carnegie Endowment for International Peace, where his research focuses on the economics and geopolitics of artificial intelligence.

He has worked at Sana, an enterprise AI company acquired by Workday, and at the New York City Office of Technology and Innovation, where he contributed to the city's AI strategy. He holds a BA in Ethics, Politics, and Economics from Yale University, where he received the Alfred Twining Hadley Prize.

Technology and International Affairs Program

The Technology and International Affairs Program develops insights to address the governance challenges and large-scale risks of new technologies. Our experts identify actionable best practices and incentives for industry and government leaders on artificial intelligence, cyber threats, cloud security, countering influence operations, reducing the risk of biotechnologies, and ensuring global digital inclusion.

Acknowledgments

The author thanks Jon Bateman, Sam Winter-Levy, Alasdair Phillips-Robins, Zhengdong Wang, David Bloom, Ara Kharazian, Josh Moriarty, Joshua New, Marina Meyjes, Tom Cunningham, Anton Leicht, Sam Manning, Harrison Durland, Alana Brase, Helena Jordheim, Jocelyn Soly, and Amy Mellon for helpful contributions. He is also grateful to Pascual Restrepo, whose course and mentorship helped shape his thinking.

Introduction

Technology has always reshaped work—60 percent of jobs in 2018 were in professions that did not exist in 1940. But new frontier AI applications may greatly speed the pace of change. Some researchers predict that advanced AI systems could function as "drop-in remote workers," upending work across sectors. Will AI systems fully replace workers? If so, which ones and how soon?

Forecasts diverge sharply, because they hinge on many hard-to-predict factors: the speed of progress in AI capabilities, the pace of AI diffusion throughout the economy, the ability of AI systems to substitute for human labor, the ability of firms and workers to adapt to AI, the price and accessibility of various models, and more. These unknowns complicate any effort to develop preparatory policies, even as many believe the next few years represent a critical planning window.

As conflicting predictions pile up and policymakers face growing demands to act, this paper seeks to clarify the emerging grounds of debate. It classifies the most prominent and credible views on AI job disruption into three loose groups and identifies their core assumptions (see figure 1). Drawing on research in labor economics, economic history, and empirical AI evaluation, the paper reviews the best evidence and arguments for each position and shows where and why they diverge.[1]

- **The alarmed** believe that AI will substitute for a large portion of white-collar labor over roughly the next decade.
- **The patient** believe that AI systems will displace and complement labor gradually, over the span of multiple decades if not longer.
- **The excited** believe that AI systems will create more new opportunities for human labor than they eliminate.

Figure 1. Perspectives on AI's Workforce Impact Vary by Magnitude and Direction

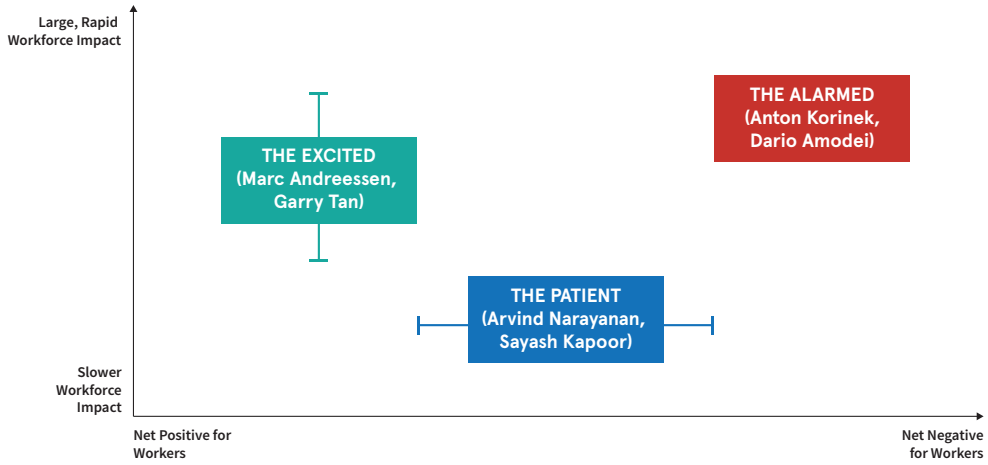

Source: Author's illustration.

A close examination of these camps reveals that their disagreement can be distilled into two core disputes:

- **The pace of AI progress and the importance of barriers to adoption.** The alarmed believe that "scaling laws," large investment, and self-improvement will lead to continued jumps in AI's real-world utility, and that adoption will be relatively seamless. Those in the patient group think that fundamental limitations in the current technology paradigm will bottleneck progress, and that brittleness, hallucinations, and change management will significantly slow enterprise AI adoption.

- **The strength and speed of AI-driven job creation.** Those in the excited camp believe that AI-powered productivity gains will lead to massive economic surplus and further increase the demand for human-centric work, creating large numbers of new jobs. The alarmed are concerned that advanced systems may replace human tasks end-to-end within certain jobs and that competition from AI-first businesses could accelerate layoffs.

By isolating these points of dispute, this paper identifies the critical indicators that policymakers should track and offers guidance on policy that is robust across the different scenarios these groups predict. Given the uncertain outlook, policymakers should prepare for a wide range of possibilities by improving data collection and designing and piloting wage insurance and applied training programs.

Figure 2. The AI Jobs Debate Comes Down to Two Questions

	THE ALARMED Dario Amodei, Anton Korinek, Sam Altman, Mustafa Suleyman	THE PATIENT Sayash Kapoor, Arvind Narayanan, Daron Acemoglu	THE EXCITED Marc Andreessen, Garry Tan, David Autor*
AI Capabilities and Diffusion When will AI complete many work tasks reliably?	**Within 5–10 years** Scaling, investment, agents, and R&D automation will accelerate progress	**>10 years** Brittleness, lack of in-context learning, edge cases, and org friction will slow adoption	**Unsure** Views vary among the group
Strength and Speed of Job Creation How strong are the forces driving new job creation?	**Weak** Demand may not expand fast enough; automation outpaces job replacement	**Moderate** Workers shift toward AI supervision and complementary tasks	**Strong** AI raises demand for complementary work and creates new businesses

Note: Those categorized as excited hold different views on whether large-scale job creation is merely a possibility or the most likely status-quo outcome. Based on this distinction, they can be divided into two groups—conditional optimists and market optimists. David Autor, an economist, is the former; Marc Andreessen and Garry Tan, prominent venture capitalists, are the latter.

The Alarmed

The alarmed believe that AI will cause a rapid collapse in the demand for labor across significant portions of the economy within just a few years. They posit that as generative AI systems become more intelligent and capable of taking actions, they will serve as substitutes for human intellect across many fields, leading to mass unemployment.[2]

This view is popular among some technology executives and investors, and is taken seriously by many economists and even politicians.[3] Anthropic CEO Dario Amodei believes that "AI could displace half of all entry-level white-collar jobs in the next 1–5 years," and he would "bet pretty strongly against" job losses being fully balanced by job creation, as seen in previous periods of technological progress. Economist Anton Korinek has predicted that "if the quest for artificial general intelligence succeeds, we are not looking at another Industrial Revolution" that ultimately rewards workers; rather, it's possible that "labor itself becomes optional for the economy." Former U.S. secretary of commerce Gina Raimondo likewise said that "A.I.-driven mass unemployment is a potential crisis on the horizon," and Senator Bernie Sanders' office wrote a report fearing that AI could "replace nearly 100 million jobs over the next ten years."

The argument of those who are alarmed rests on three premises:

1. **Capabilities progress:** Over the next decade or sooner, AI systems will be able to complete the vast majority of tasks involved in most white-collar professions.[4]

2. **Rapid adoption:** Employers will prefer to use AI systems and lay off large numbers of human workers as a result.

3. **Limited opportunities:** New jobs will not sufficiently replace or supplement those automated by AI systems, leading to large increases in unemployment.

Capabilities Progress

Along many dimensions, AI capabilities have raced ahead in recent years; the alarmed believe this trend is on pace to continue or even accelerate. As semiconductors become more advanced and companies spend huge sums to construct AI data centers, the amount of computational power used to train and run AI systems will increase, enabling more powerful models.[5] Improved model-training algorithms will also mean that AI systems can be trained more efficiently even with the same amount of computing power.[6] And as AI systems are given access to more tools (calculators, search engines, software) and broader context, their effective capabilities will also increase.[7]

This progress is why the alarmed believe that AI systems will soon be able to perform the full depth and breadth of tasks involved in white-collar work. A growing body of research supports their view. In particular, a series of new "applied" evaluations attempts to assess AI systems' capabilities in more realistic scenarios than past benchmarks. By testing models on the kind of complex, situational, multi-step tasks that professionals undertake across service professions, researchers have revealed the growing potential utility of AI on the job.

One such test is OpenAI's GDPVal, a benchmark that evaluates performance at 1,320 common workforce tasks. OpenAI's team of economists analyzed which tasks take up the most time across forty-four of the occupations responsible for the largest share of U.S. GDP. They then translated those high-value tasks into an evaluation designed to measure AI performance on these critical job activities. The tasks in the benchmark were complex and designed by experts: They took an average of seven hours to complete and were written and graded by professionals with an average of fourteen years of industry experience. The newest AI models beat human workers when tested on a subset of 220 such tasks. Expert judges preferred the responses of GPT 5.4, released in March 2026, to human responses or rated them as a tie 83 percent of the time.

Other research corroborates these results, showing AI models exceeding human performance in medical and legal tasks, and rapidly improving at coding. Preliminary results from an MIT working paper suggest that AI systems are improving at a constant rate across a variety of real-world work tasks, and top engineers at OpenAI and Anthropic claim that they have almost entirely automated their coding process.[8] Real-world studies across other fields, surveyed in an article by University of Chicago economist Alex Imas, also show that AI models make human workers more productive across a variety of sectors.[9] In one March 2025 working paper about a randomized controlled trial by the company Procter & Gamble, individual employees assisted by AI matched the performance of entire teams without AI at solving actual product innovation challenges.

There are reasons to think that as models continue to improve, they will not merely approach the level of human expertise but rather could blow past it. For example, some speculate that automation of AI research and development could lead to fast-compounding advances.[10]

Rapid Adoption

The alarmed believe that highly capable AI systems will reduce demand for human labor because in many areas, employers will prefer to "hire" AI employees. AI systems can perform many tasks that would otherwise be assigned to humans at a lower price.[11] They can "work" 24/7 on command, never unionize, and require no healthcare or payroll taxes.[12]

Even with an "AI agent" workforce, in the short term, employers would still enlist teams of humans to manage the automated systems—setting direction, providing feedback, and verifying outputs. Human verification will be particularly important in high-stakes

domains like law, medicine, and financial transactions. Other factors will also impede full automation, such as the difficulty of embodying AI systems via robotics, regulations that protect licensed professionals, and a continuing preference for humans to fulfill certain roles. Therapists, judges, construction workers, and food service workers among others will likely be safe for the immediate future.

However, the top professionals in fields like law are often billed at more than $1,000 an hour, creating immense incentives to implement AI wherever possible. There might still be human lawyers, consultants, and financial professionals, but if their expertise is increasingly commoditized due to AI, the number of junior staff hired each year could decrease significantly.

There may already be evidence of AI-related job loss. A November 2025 pre-print by Stanford researchers found that while employment for non-AI exposed professions and older workers has not changed since the release of ChatGPT in late 2022, "early-career workers (ages 22-25) in the most AI-exposed occupations have experienced a 16 percent relative decline in employment."[13] Economists like Erik Brynjolfsson and Jason Furman also argue that AI has begun to impact productivity statistics. (This is an area of active debate, as other reputable sources dispute the impact of AI on labor markets. Carnegie's Alasdair Phillips-Robins provides an overview of this research in a forthcoming companion article.)

> **The largest hyperscaler tech companies are funding data centers that will use more power than Los Angeles.**

How far could these disruptions go as current models diffuse? Task-based studies estimate the percentage of tasks that AI models could automate across occupations, providing an upper bound. A foundational paper by a team at OpenAI, Wharton, and the Centre for the Governance of AI found that GPT-4 could speed up more than half of tasks for more than 46 percent of workers "when accounting for current and likely future software developments that complement LLM capabilities."[14]

The alarmed believe that progress in capabilities and implementation will not slow down. The largest hyperscaler tech companies are funding data centers that will use more power than Los Angeles, by some measures exceeding all past infrastructure buildouts except railroads in relative scale. AI companies are beginning to commercialize their models more aggressively, and many of the world's smartest people are being paid like star athletes to accelerate progress.[15]

Limited Opportunities

Finally, those in the alarmed camp argue that, unlike previous episodes of technological change, AI may not create many adjacent jobs for displaced workers to shift into. The strength of this argument depends heavily on the trajectory and pace of AI progress.

In his essay "The Adolescence of Technology," Anthropic CEO Dario Amodei lays out several of the reasons many in this camp believe this time could be different. Many earlier technologies improved over decades, affected workers in specific sectors, and still left clear roles for humans, like loading raw materials or operating machines that had replaced physical labor. Frontier AI systems, by contrast, appear to have different characteristics. They have improved rapidly over just a few years, threaten junior-level workers across many fields, and are quickly adapting to overcome many of their current weaknesses. As a result, workers could soon be displaced en masse, with few new openings in their fields and limited pathways to enter new ones.

Those in the excited group offer counterarguments emphasizing how AI could create new jobs, but they are weakened if AI capabilities continue to advance rapidly. Physical and social work tasks will experience less disruption from AI, and humans may continue to specialize in areas where they retain a comparative advantage. But it is not clear that those who become unemployed due to AI will be well positioned to shift into these jobs. And if AI capabilities improve quickly, automated systems may encroach on even these areas. The fear is that we will live in a world with structural technological unemployment and greater inequality. As economist Daniel Susskind puts it, "there might not be enough demand to provide employment for everyone who wants it."

Nor does this scenario require AI to bring a wholly unprecedented level of change. In the past, technological change has at times raised productivity without generating offsetting demand for labor. Consider as a stylized example the automated checkout kiosk. If four cashiers are replaced by four kiosks overseen by one employee, average productivity rises, but employment falls because the technology does not create enough new complementary tasks to absorb the displaced workers.[16] This dynamic is likely to occur in sectors where demand for a good or service is relatively fixed. Hand-weaving in the early nineteenth century provides a historical example from the Industrial Revolution: after the mechanized loom spread, wages and employment in weaving fell sharply, putting many thousands of artisans out of a job.[17]

> **This paints a concerning picture, one in which many millions of Americans and more around the world could be out of a job.**

The alarmed also fear that policymakers will not help workers adjust to these dramatic changes. In the past, policymakers have failed to shield, retrain, or compensate workers effectively as jobs were automated by robots and outsourced due to trade.[18] In fact, current U.S. policy is not just to let AI progress and diffusion happen, but rather, to actively accelerate these processes. Some policymakers view doing so as a national security imperative to beat China, which has taken steps to spread AI across its industrial economy.

This paints a concerning picture, one in which many millions of Americans and more around the world could be out of a job. Software engineers, customer service representatives,

personal assistants, office clerks, project managers, truck drivers, paralegals, and even highly paid lawyers, consultants, and financial analysts could soon face steep competition from AI systems.

The consequences could be destabilizing and would necessitate a proactive government response involving some combination of selectively slowing adoption, creating new labor opportunities, matching displaced workers to those new jobs, and redistributing economic gains from AI.[19]

But not everyone agrees.

The Patient

Drawing from past technologies such as motor engines and the internet, the patient posit that it will take at least several decades for AI to diffuse throughout the economy due to technical barriers and economic frictions. As such, they argue that AI may still significantly boost productivity, but will do so far more gradually; it will not cause demand for human labor to collapse over the next ten to fifteen years.

The patient question whether AI systems will be capable enough to complete complex work tasks, reliable enough to be trusted with important work, and diffused quickly across the workforce even if so.

Their argument rests on three gaps:

1. **Capabilities gap:** In addition to existing deficits, for at least the next decade, AI systems will fail to complete many complex, long-horizon tasks at human level due to limitations in the way they learn compared to how humans do.

2. **Reliability gap:** AI systems cannot be trusted to complete tasks *every time* due to persistent hallucinations and failures that erode productivity gains from AI use.

3. **Adoption gap:** Redesigning workflows and organizations to integrate AI systems is a messy, slow process with many bottlenecks.

Other AI researchers and economists endorse more patient views. For example, in their influential essay "AI as Normal Technology," Princeton computer scientists Arvind Narayanan and Sayash Kapoor write "we think that transformative economic and societal impacts will be slow (on the timescale of decades)." Nobel laureate economist Daron Acemoglu writes that "neither the economic theory nor the data support . . . exuberant forecasts" that "recent advances in generative AI will soon bring extraordinary productivity benefits."

Those in the patient group have different views on the underlying capabilities of AI models and whether they may transform the economy over the long term. A few believe that AI's impacts are greatly exaggerated and that generative AI will prove unhelpful to most workers. In their book *The AI Con*, linguist and computer scientist Emily Bender and sociologist Alex Hanna write that AI "isn't sentient, it's not going to make your job easier, and AI doctors aren't going to cure what ails you." Others harbor greater uncertainty farther into the future. Gary Marcus, a cognitive scientist, states that "most white collar jobs aren't going anywhere that soon" and explains that "the junior people are under some threat [but] . . . that threat is actually exaggerated."

The Capabilities Gap

Those in the patient group question the linear narrative of AI progress favored by the alarmed. Although AI is undoubtedly improving along many dimensions, the patient point to several specific limitations that will challenge workforce AI adoption.

First, they highlight deficits in general planning and reasoning as a major explanation for why AI systems still cannot complete many real-world tasks. ScaleAI's Remote Labor Index (RLI) provides an important example. The RLI tests AI systems at the kind of complicated tasks that a human worker on the gig platform Upwork would need multiple days to complete. ScaleAI's team found that current AI systems perform very poorly. As of March 2026, the best AI system tested, Claude Opus 4.6 Cowork, was only able to complete 4.17 percent of these tasks at a level matching or exceeding the human gold standard. This result and others from evaluations like ARC-AGI 3 suggest that AI models' well-known successes at coding may be an aberration.

AI systems also tend to be brittle, meaning they sometimes struggle to adapt to situations that differ too much from their training data. Ilya Sutskever, a key architect of ChatGPT, has noted that their algorithms "generalize dramatically worse than people." He points to a key limitation: unlike humans, AI systems struggle to map knowledge across concepts and grasp sophisticated causal relationships. Even Sutskever, who previously believed that AI technology would advance fast, has called this problem "very fundamental" and pointed to the need for new approaches.[20]

> **A human worker is capable of extracting rich feedback from nearly every interaction and using that feedback to improve. Current AI models can't do this.**

Beyond those existing deficits, AI models do not learn continuously in the way that humans can. A human worker is capable of extracting rich feedback from nearly every interaction and using that feedback to improve. Current AI models can't do this. While they have limited working memory, the model itself does not actually learn or update as a result of user

feedback. The podcaster and commentator Dwarkesh Patel argues this limitation serves as a critical bottleneck preventing AI systems from learning the tacit knowledge that is important to do many jobs well.

One of the main catalysts of AI progress in recent years has been scaling models by using more computing power and data to train them—but it is not clear how much longer developers can scale up inputs due to limits in power supplies and raw data needed to train models.[21] Moreover, those in the patient group believe that many key limitations are inherent to the prevailing technology paradigm. For example, the writer Timothy Lee argues that current models' tendency to become overwhelmed by long tasks may be fundamental to current LLM architectures.

Breakthroughs are not impossible, but those in the patient group like Richard Sutton, Yann LeCun, and Gary Marcus argue that they are far from guaranteed and will require fundamentally new ideas.

The Reliability Gap

Beyond the models' limited raw capabilities, those in the patient group also point to poor reliability, which Narayanan and Kapoor have defined as "behav[ing] consistently across runs, withstand[ing] perturbations, fail[ing] predictably, or respect[ing] safety constraints." For example, even leading models still hallucinate, or confidently fabricate false information. Narayanan and Kapoor note that their own benchmarking efforts illustrate that "nearly two years of rapid capability progress have produced only modest reliability gains." They find that OpenAI's models from the end of 2025 are not that much more reliable than their models from 1.5 years prior.

Because AI systems struggle with edge cases, they must be continually tested and iterated upon before they can be deployed, particularly in high-stakes domains. This "capability-reliability gap" explains why it took more than two decades to develop reliable systems for self-driving cars and why medical AI systems often succeed in testing but fall short in the real world.

Hallucinations may serve as a particularly important barrier to implementation. Researchers train AI models to be truthful and accurate, but doing so is difficult and can conflict with other goals. Hallucinations can be mitigated through techniques such as Retrieval-Augmented Generation, which helps models identify the most relevant information to draw on when answering a question. They remain an unsolved problem, however, and some computer scientists believe they are fundamentally unsolvable.

In many economically valuable domains, reliability is critical. A single software bug can cause a widespread cloud outage, and one spreadsheet error can lose a company millions of dollars. Even if a model works 99 percent of the time, in many cases 100 percent of its

work still needs to be checked, reducing efficiency gains.[22] One citation to a hallucinated case could cause a legal brief to be thrown out of court.[23] Of course, it is true that humans are not 100 percent reliable, either. But if AI systems make different types of errors that are harder to catch or their errors are more consequential, they still introduce important verification costs.[24]

> **Even if a model works 99 percent of the time, in many cases 100 percent of its work still needs to be checked, reducing efficiency gains.**

At best, this means that humans must extensively check AI's work, reducing its value proposition. At worst, it simply is not worth it to integrate AI into many workflows. Checking over work is often not so different from doing it manually, and a financial model or a computer science project may need to be entirely reworked due to cascading impacts from a single error.[25]

The Adoption Gap

Finally, the speed of AI adoption is likely to be limited by the speed of human skill-acquisition and organizational change, giving rise to an adoption gap.[26] To paraphrase researcher Deena Mousa, we live in a world that requires adapting society's rules to AI and AI systems to society.

Think of the dozens of software programs and legacy IT systems across which enterprise data is scattered, often saved in incompatible, incomplete, or unstructured forms. Think of the vocational apprenticeships through which human workers build essential skills and relationships, typically over the course of years. Think of the new security vulnerabilities that AI models introduce and of how systems will need to be "idiot proofed" to be widely adopted. Think of the unresolved liability questions that arise when AI makes a consequential error, and of the professional licensing regimes in medicine, law, and finance that restrict which tasks it can legally perform at all.[27]

These factors complicate the alarmed group's picture of the "drop-in remote worker." They mean that companies must adapt applications, redesign org charts, fix critical issues, and create new checks and accountability structures to integrate AI into their work. Aided by consultants, in-house AI teams, and startups, this process will occur, but it will likely be slow, taking years or even decades.[28] Most tasks that can be automated easily are already heavily software-assisted or have been turned over to industrial robots and other machines. Many of the others require a large amount of tacit context or benefit from a human touch to be completed well.

This pattern of adoption bottlenecked by industrial organization has been observed before. Factory floors had to be redesigned to make the most of electricity during the industrial era, a process that took more than forty years. It took decades for advances in information technology to be assimilated into the workforce and impact productivity statistics, a phenomenon dubbed the "productivity paradox." The pattern appears to be repeating: Many AI drive-thru pilots have failed and demand for radiologists has grown. This helps explain why in a recent survey, academic economists on average "expect AI capabilities to improve significantly by 2030, but… do not expect this to translate into dramatically different economic outcomes."

Because of the capabilities, reliability, and adoption gaps, those in the patient group generally believe that AI is best viewed as "normal technology," a term coined in an essay by Narayanan and Kapoor. Adherents to the "normal technology" view believe that AI is not so different from previous technologies like electricity and the internet. Over the long term, it could well have a large impact, but disruption will be gradual rather than rapid.

In the patient view, then, dramatic policy interventions like attempts to limit the spread or development of AI are premature and likely counterproductive. Advocates of this view instead endorse a "wait and see" approach with limited government intervention, which could include steps like improving data collection or targeted improvements to social safety programs.

The Excited

Those in the excited group are less concerned with existing jobs than with the new jobs AI will create. They argue that AI's overall effect on labor markets will be significant net job growth, and so society's main task will be ensuring that displaced workers are well-equipped to transition into different roles.[29]

The excited camp can be divided roughly into two groups. Market optimists, many of whom are technology executives, believe that market forces will usher in positive outcomes from AI integration even without policy changes. Those in this group, like venture capitalists Marc Andreessen and Garry Tan, draw on a powerful historical argument: In the past, when new technologies were introduced into the economy, job creation largely overpowered job destruction. Andreessen expresses strong conviction, writing "technology doesn't destroy jobs and never will" and that AI "may cause the most dramatic and sustained economic boom of all time, with correspondingly record job and wage growth."[30]

By contrast, conditional optimists like economists David Autor and Erik Brynjolfsson see job creation overwhelming displacement as a plausible scenario, but far from a sure thing. Whether this plays out, they argue, depends on the institutional context in which AI is

developed, such as the extent to which society incentivizes job-creating uses of AI and helps workers gain new skills. Autor writes that "AI, if used well, can assist with restoring the middle-skill, middle-class heart of the US labor market," though he cautions that this is "not a forecast but an argument about what is possible."[31]

Despite their differences, those who are excited generally believe that AI systems will serve as engines of job creation due to a combination of three different forces:

1. **Task recomposition:** Within jobs, AI systems increase demand for tasks that only humans can do.

2. **Income effects:** Across the economy, efficiency gains from AI will free up income to flow toward other areas, increasing demand for human labor.

3. **AI-driven businesses:** AI systems will enable the creation of new businesses that will also create many human jobs.

The excited have different views on how AI capabilities might advance, which influence the channel of job creation they emphasize. Market optimists who believe that AI systems will soon become highly capable highlight the potential of AI systems to create jobs by boosting income and facilitating the creation of new businesses. Conditional optimists often emphasize task recomposition within existing jobs; many believe the gaps emphasized by the patient group will cause many jobs to have a long tail of difficult to automate tasks that will anchor demand for human labor.

Task Recomposition

Task recomposition occurs when employers change jobs to emphasize tasks that complement AI systems, which generally increases the value of human labor. This outcome is socially desirable because, rather than displacing workers outright, it would likely allow many to remain employed or move into adjacent roles within the same industry.

This can happen in two ways. First, AI can allow workers to shift time away from routine or low-value tasks and toward higher-value activities. Second, new technologies can create new human tasks or roles within existing firms that complement automated systems. Together, these forces can improve quality, reduce costs, and—where demand is sufficiently elastic—expand employment.

Sales, one of the job functions most exposed to AI, illustrates the first channel. AI sales enablement tools like Clay, Apollo, and Gong are designed to change the work of sales representatives by automating certain steps of prospecting and outreach. Their aim is to redirect representatives' time away from low-value activities like untargeted cold calls and mass emails, and toward higher-value work like customizing pitches, refining demos, and managing qualified leads.[32]

Figure 3. Technologies Like Spreadsheets Didn't Replace Workers—They Created New Tasks for Human Labor

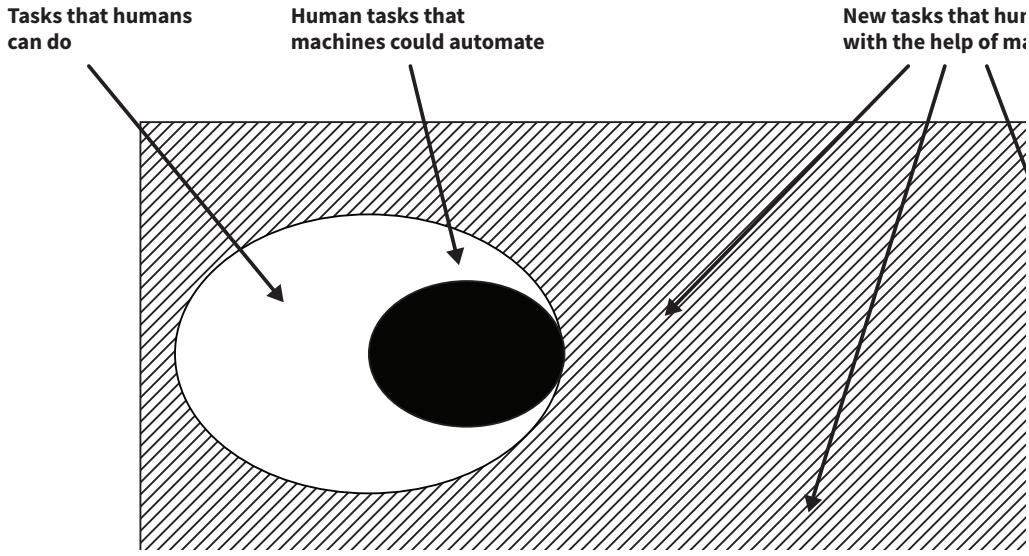

Source: Erik Brynjolfsson, "The Turing Trap: The Promise & Peril of Human-Like Artificial Intelligence," *Stanford Digital Economy Lab*, January 12, 2022, https://digitaleconomy.stanford.edu/news/the-turing-trap-the-promise-peril-of-human-like-artificial-intelligence/. Recreated with author's permission.

The second channel, new task creation, has historical precedent from past industrial technologies like CNC machines, welding robots, and laser cutters. One study suggests that Finnish firms that adopted industrial automation technologies from 1994 to 2018 increased employment by 23 percent, as they used new tools to produce new products like specialized pistons and cater to new market segments rather than replace workers.[33]

These forces also mean that output can be produced at a lower cost or a higher level of quality. So long as demand for the relevant good or service is somewhat price-elastic, lower prices or better products can increase demand enough to create more opportunities for human labor.

Demand is likely elastic for many of the professional services industries that AI is expected to disrupt. As David Autor writes, "demand for healthcare, education and computer code appears almost limitless — and will rise further if as expected AI brings down the costs of these services." Of course, how much demand will increase is an open question; it is not clear that firms would become twice as litigious if legal services cost half as much.[34] Still, many service jobs are likely not similar to agriculture, a field where productivity gains in the late nineteenth and early twentieth centuries so vastly outpaced demand growth that net employment still collapsed.

Whether these forces preserve or expand employment depends on a critical firm-level decision: whether to redeploy workers toward higher-value tasks or simply shed them as productivity rises. The excited scenario requires the former. How often firms choose to redeploy workers depends in large part on how deeply intertwined the automatable and non-automatable tasks are within a given role—which brings a real constraint on full automation that the alarmed sometimes understate.

A financial manager, for example, might both make investment decisions and construct financial models. In theory, they could delegate the modeling to AI while keeping the judgment-intensive decisions human. In practice, however, the manager may still want to engage directly with the model to understand its assumptions, or rely on an analyst who can explain the first draft, answer questions, and weigh in on the investment thesis.

More broadly, most jobs involve context-rich and iterative processes rather than discrete tasks that can be neatly peeled away. Humans will likely remain necessary for at least some parts of these workflows (see figure 4); so, in many cases, AI will reshape jobs around higher-value tasks rather than replacing workers outright.

Figure 4. Some Jobs Will Likely Have a Long Tail of Tasks That are Difficult to Automate

Example: Portfolio Manager, illustrative job tasks*

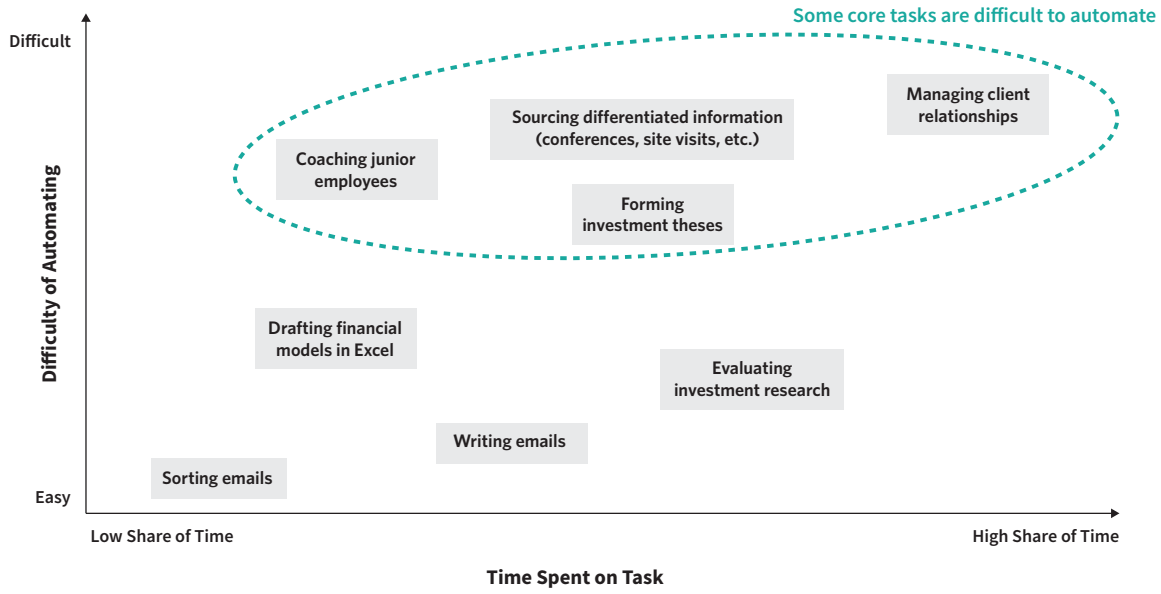

Source: Author's illustration. Illustrative example; task placement is approximate and reflects directional judgment rather than precise measurement.

Income Effects

If AI systems perform jobs more efficiently, then the profits earned by businesses will rise, the prices paid by consumers will fall, or both. As businesses and consumers are left with more money to spend, jobs will be created elsewhere in the economy. Economists call this the productivity effect.

A key question is where that additional income flows. The alarmed group argues that it will likely be reinvested in further automation. AI companies and adopters may capture a substantial share of the gains and use them to fund research, data centers, and enterprise AI adoption, potentially accelerating labor substitution.

But those who are excited counter that gains from AI will not be used only to automate more work. They are instead likely to be spent in large part on human-facing service occupations or physical jobs—construction, hospitality, personal services (such as therapists or trainers), and artisanal crafts. As incomes rise, consumption will likely shift toward high-end goods and services that are difficult to standardize or fully automate, the appeal of which lies in their human qualities.[35] If demand for these products and activities rises faster than productivity, their prices and economic importance may increase as well, meaning the economy could be dominated by education, healthcare, and other high-touch work.

AI-Driven Businesses

AI will also create entirely new categories of jobs and business opportunities. In the short term, deploying AI at scale will require a large new layer of human work in integration, oversight, evaluation, and workflow redesign. The IT boom precipitated a large enterprise software and IT services ecosystem, from firms like Oracle and SAP to services companies like Accenture, TCS, and Infosys. To overcome the reliability, oversight, and integration challenges, AI will likely require a whole ecosystem of human roles in data labeling, quality assurance, model evaluation and monitoring, strategy, and workflow integration.

> **To overcome the reliability, oversight, and integration challenges, AI will likely require a whole ecosystem of human roles in data labeling, quality assurance, model evaluation and monitoring, strategy, and workflow integration.**

Looking further ahead, AI may also enable more novel business models, a point emphasized by economists Ajay Agrawal, Joshua Gans, and Avi Goldfarb, as well as Erik Brynjolfsson. The internet facilitated the rise of new businesses like Amazon, which redesigned the concept of what a retail store could be from the ground up. As mobile devices became widespread, platforms like Uber and Airbnb became major companies with large numbers of contractors,

vendors, and workers. It's not clear exactly what new business models AI will enable, but already, new companies are experimenting with "AI first" structures for law, banking, commercial insurance, and other industries. And we have already seen early evidence that AI-native businesses can outcompete their peers. A working paper by INSEAD and Harvard Business School researchers found that early-stage startups chosen at random to receive training on AI use cases earned 1.9 times higher revenue than peer companies that did not.

The alarmed see this as a cause for concern. In some industries, lean AI-first businesses will displace incumbents and gain market share by minimizing labor costs. Their success could accelerate job cuts across industries and functional areas.

The risk is real, but the possibilities for new work are also substantial. Think of a world where automated coding dramatically reduces the cost of producing software, where every person has a world-class tutor in their pocket, and where custom medical, legal, and financial analysis is accessible on demand. That could lower start-up costs for small businesses and enable a wide range of new companies, possibly including AI-native education, agencies, hedge funds, and compliance and administrative support; custom enterprise software providers; and personalized health or financial guidance platforms. Businesses like these will expand markets, create new types of work, and likely rely on humans for at least some functions. What cannot be automated, or is complemented by human intellect, social skills, decisionmaking, and physical abilities, will become highly valuable.

> **The barriers the patient emphasize become complementarities that anchor demand for human labor, while the capabilities that concern the alarmed become opportunities to expand markets and create new businesses.**

The optimistic argument can be the hardest to articulate and often relies on historical analogies because it is difficult to know what the jobs of the future will look like. But that doesn't mean it's wrong. The excited provocatively flip the arguments of both other camps: The barriers the patient emphasize become complementarities that anchor demand for human labor, while the capabilities that concern the alarmed become opportunities to expand markets and create new businesses. For the excited, the central question is not whether AI automates some tasks, but whether lower costs expand business and existing roles faster than task substitution reduces headcount.

Conclusion

All three camps can draw upon credentialed and articulate advocates, cite evidence from history and current AI development, and tap into deeper political and moral intuitions about the nature of technological change and the appropriate role of government in managing that change. Their opposing arguments put policymakers in a difficult situation, forcing them to make path-dependent decisions about taxation, workforce development, social programs, and education under drastic uncertainty.

These camps are not mutually exclusive—reality could well contain elements of all three. They also could occur in sequence. For example, the patient view may initially prove right as enterprise adoption is slowed by frictions, but if a research breakthrough resolves the key bottlenecks, the fears of the alarmed or the hopes of the excited may abruptly come true.

> **These camps are not mutually exclusive—reality could well contain elements of all three.**

While the camps disagree on many relevant issues, there is common ground. Even most who are excited believe that government and third parties should monitor labor impacts to ensure that limited shocks do not trigger a broader political backlash. And even most of those in the patient group believe that governments should be prepared, as adoption could accelerate in the future.

In particular, the groups are likely to agree with two recommendations; while the proposals lean toward the more cautious posture of the patient group, they are designed as a foundation to detect accelerating adoption and scale up if or when disruption occurs.

1. **Policymakers, industry, and researchers should work together to improve data collection.**

   America's AI Action Plan, the Trump administration's 2025 AI strategy, directed the Department of Labor to establish an AI Workforce Research Hub to gather additional data on AI's labor impacts. This is a positive step, and policymakers should go further to monitor indicators relevant to the different possibilities.

   The U.S. government primarily tracks AI adoption through the Census Bureau's Business Trends and Outlook Survey and its AI Usage Supplement, which asks businesses whether they have adopted AI and whether they plan to over the next twelve months. Companies answer "yes" or "no," but these answers do not provide insight into how or how much they use AI. Companies like Ramp and Anthropic have done important work to supplement this data. But they are companies, not disinterested third parties, and do not have access to the full picture.

As economist Sam Manning and others have suggested, the Department of Labor should investigate the possibility of partnering with AI developers and payroll and hiring platforms to publish more detailed data that includes enterprise use of AI systems through API access. Other stakeholders could help clarify expectations through predictions of these indicators. For example, the Chicago Booth U.S. Economic Experts Panel could ask more questions about labor impacts from AI, while platforms like Good Judgment Open could create forecasting challenges tied to these measures.

To track the debate between groups, policy analysts and economics researchers should monitor three categories of data.[36] First are measures of how capable and reliable AI systems are, providing insight into the readiness of leading AI systems for workforce deployment. These would include evaluations of performance on real work tasks (for example, RLI, APEX-Agents, and Vals AI's applied evaluations, among others), tests of long-horizon reasoning and planning (for example, ARC-AGI 3, Vending-Bench 2, LongCoT, and CRUX), benchmarks of AI reliability (for example, those surveyed in the HAL Reliability Tracker), and experiments showing AI systems' impacts on productivity.

Second, they should monitor measures of AI adoption among first movers, which will precede wider use. These include measures that track enterprise AI spending like the Ramp AI Index and data that compares adoption across workflows and customer segments, in the vein of Anthropic's publications. Junior hiring in highly exposed professions and signals from financial markets are also useful indicators, though they may be noisy.

Third to track are measures that indicate whether AI is augmenting human workers or substituting for their labor. These include measures of wages and output in highly affected jobs and time-use and job posting information to evaluate how skill demands are changing. New research by OpenAI as well as other companies has explored some of the relevant questions.

Combining these data streams could yield even richer insights. Analyzing AI adoption and usage intensity alongside wage and hiring data would allow for more granular understanding of how AI adoption is reshaping the labor market.[37]

2. **Researchers should design and pilot wage insurance and training programs targeted at professions highly likely to suffer job losses.[38]**

Even many in the patient group acknowledge that certain jobs, like customer service representatives, are likely to decline significantly within the next fifteen years.[39] This would be fully consistent with AI as "normal technology," because it is indeed normal for new technology to gradually eliminate certain tasks and roles (such as elevator operators, typists, and so on). Data on AI exposure, usage, and freelance platforms provides guidance—albeit limited—on which professionals are at the greatest risk.

Academic research indicates that when a profession is exposed to automation, those who are displaced—like telephone operators in the 1930s or manufacturing workers who suffered from Chinese competition—struggle to find new jobs and experience worse life outcomes as a result.

Retraining will likely be necessary to connect displaced workers with new opportunities, but designing effective programs will be a tall task for policymakers. Research on an important past U.S. retraining program suggests that its effects faded over time, in part because it taught workers skills that themselves later became obsolete. But it is difficult to say which skills will become obsolete with AI—a skill that is highly complementary to AI systems one day may be automated the next. At the same time, research suggests that programs should focus on STEM and technical/health vocational training, as nontechnical retraining (sales, service, and social science courses, among others) likely fails to improve wages. The problem is, technical professions seem among the most exposed.

With these pitfalls in mind, researchers should design and pilot wage insurance and micro-credential programs. Wage insurance provides cash to displaced workers who find reemployment at a lower wage, and research suggests that it boosts long-term earnings for workers displaced by trade. Programs that offer fast credentials and incentivize on-the-job training are also promising. Policymakers will not be able to centrally plan for the jobs of the future, so the next best thing is to reduce friction for displaced workers to gain skills quickly and learn from experience. In designing these programs, researchers can account for lessons learned from initiatives like coding bootcamps.

At the core, the three camps disagree over a specific set of issues: how quickly AI capabilities will improve on real work, whether reliability and verification costs fall enough for deployment in high-stakes settings, whether firms can redesign workflows fast enough to diffuse these systems widely, and whether new tasks and businesses grow quickly enough to offset substitution.

These are empirical questions, and economists can track specific indicators that provide insight on them to serve as fire alarms. While the three groups disagree on what will likely happen to workers, they do not necessarily disagree on what should happen under different conditions, creating an opportunity for researchers to engage in scenario planning. To make this monitoring framework actionable, economic and policy researchers should identify decision-relevant thresholds that would trigger more aggressive intervention.

The recommendations above, improving data collection and piloting wage insurance and training programs, are deliberately modest precisely because they are robust across all scenarios. Regardless of whether AI complements or substitutes for labor, and how much it does so, better data costs little and prepared workers will best weather the changes ahead.

The debate will be settled not by rhetoric but by evidence. The sooner policymakers strengthen the mechanisms to gather it, the better positioned they will be to act if or when necessary.

Notes

1 Relevant economics literature often caveats the existence of highly advanced AI systems and models potential impacts, while technical writing on AI capabilities often does not fully consider potential barriers to adoption. This paper helps bridge the gap. It is related to some writing on the question of whether AI will cause 'explosive' growth, but is more narrowly scoped to focus on labor.

2 The primary objective of each section of this paper is to lay out that group's worldview, rather than to assess whether they are correct. I mention select counterarguments in each section to provide some balance.

3 Sam Altman, Microsoft AI CEO Mustafa Suleyman, and Elon Musk have expressed similar sentiments. Some in this group make more extreme forecasts, like startup cofounder Matthew Barnett: "I would guess that there is roughly a 1 in 3 chance that human wages will crash below subsistence level within 20 years, and a 2 in 3 chance that wages will fall below subsistence level within the next 100 years."

4 To some, this also includes not just intellectual tasks involved in white-collar work today, but also any tasks that might be part of the production process in the future, including research and development.

5 In the past, this has led to predictable improvements in model performance, and researchers at the organization Epoch AI project that these inputs will continue to scale at a similar pace through 2030.

6 Improved algorithms are in large part responsible for the impressive performance of DeepSeek's R1 model. Because of this progress, the amount of computational power needed to achieve the same level of AI performance decreases by an average of 3x per year.

7 This process is closely coupled with increasing reliance on reinforcement learning, a technique that improves models' ability to reason over long horizons and use tools effectively. These methods are far from perfect, and their limits are discussed in endnote 20.

8 This research is not fully consistent with the alarmed view. While its authors posit that "AI performance is improving rapidly across a wide range of task durations," they also state that improvements have been gradual, meaning that "near-perfect performance will take considerably longer [than 2029]" and providing a possible "window for worker adjustment." Additional research corroborates a trend of fast improvement. Merali found that with each year of progress in AI models, the time models take to complete consulting, data analysis, and managerial tasks decreases by an average of 8 percent. Models have also improved quickly on other benchmarks of professional competence, like Mercor's APEX-Agents.

9 It is worth noting that many of these studies assess the effectiveness of AI tools in improving the productivity of programmers or professions with more routine work tasks (for example, work in call centers or writing advertising copy).

10 This point is a subject of debate among economists, as some believe AI R&D will be difficult to speed up significantly. Northwestern University economist Benjamin Jones, for example, writes that under one framework of how AI may impact R&D, "bottlenecks severely mute the effect of extremely productive AI."

11 Current prices may be artificially low due to sales discounts meant to promote adoption. However, some executives disagree and claim that inference margins are high—that is, that once a model is trained, it is highly profitable, and it costs much less to run than the price at which it is sold to users. Anthropic CEO Dario Amodei stated in February 2026 that the company's "gross profit margins on [. . . inference] are very high because inference is efficient." One third-party analysis estimated that OpenAI's GPT-5 model was unprofitable overall, though that result was driven by the fact that OpenAI replaced it after only four months.

12 They can also be copied at low cost and trained to apply the same consistent rules across a wide range of cases.

13 Additional work demonstrates AI's impact on freelancers and posits that AI is already impacting productivity statistics. Data from the financial services company Ramp shows that, as of February 2026, 47.6 percent of American businesses have paid subscriptions to AI models, platforms, and tools, and a National Bureau of Economic Research working paper estimated from a survey in November 2024 that generative AI tools save users 5.4 percent of their time. Research from OpenAI and Anthropic provides additional insights on how foundation models are used in the workplace.

14 For more on this study and others like it, the Yale Budget Lab published a helpful overview in February 2026.

15 The lesson that many of these firms take from the tech-telecom bubble isn't that a market correction is coming; it's that there is real value underneath the hype and the larger risk is being left behind. Amazon, which almost went bankrupt in 2001, now has a market capitalization in the trillions. AI companies have also announced many strategic partnerships with consulting firms to distribute their products. For example, OpenAI has announced a series of "Frontier Alliances" with a number of top strategy and IT consulting companies. In March 2026, Palantir and Bain & Company expanded an existing partnership first announced in 2025.

16 In practice, evidence on whether automated checkout kiosks have lowered employment in retail stores seems mixed.

17 More recent examples have exhibited a similar pattern. A famous paper by economists Daron Acemoglu and Pascual Restrepo finds that the automation of routine work tasks "has played a defining role in the surge in U.S. wage inequality over the last four decades." The pattern may be repeating. Anthropic's Economic Impacts team wrote in March 2026 that high-skill first adopters "have more successful interactions with Claude than later, less technical adopters," raising the risk of unequal outcomes.

18 During the Industrial Revolution, as technology displaced workers and commoditized their labor, these changes had broader societal effects. Many turned to radical political ideologies from the Luddites to Marx. Some fear that AI will precipitate even worse disempowerment.

19 Those in other groups sometimes make similar recommendations. Daron Acemoglu and David Autor, who I categorize, respectively, as members of the patient and excited groups, also call for selectively incentivizing adoption of "pro-worker" AI systems that complement human workers rather than replacing them.

20 Researchers have used techniques based on reinforcement learning to address these issues, but they have critical limitations. Reinforcement learning forces models to attempt tasks themselves, learn from imperfect behavior, and adjust their behavior to navigate open-ended environments. However, these techniques generally provide a sparse reward signal at the end of a long task. They are somewhat akin to teaching a student math not by instructing them, but merely by grading whether their work is correct. Many experts believe that these methods are very inefficient compared to how humans learn, and they do not work well outside of "verifiable domains" where answers can easily be categorized as right or wrong.

21 Chip production capacity and latency (which serves as a sort of speed limit to AI training) could serve as additional bottlenecks on scaling.

22 Thank you to an anonymous partner at a top law firm, credited for this phrasing.

23 As noted above, the accuracy of AI models also degrades as models are given more context.

24 Notably, frontier AI systems are most likely to hallucinate or fail in precisely the domains where these hallucinations will be hardest to find and fix. The new RL-based AI paradigm is best at training AI in domains where answers are easily checked, like solving a coding challenge or a Sudoku puzzle, as these verifiable rewards can be used as signals for RL. But this means that models will improve more slowly and continue to fail at less easily verified tasks, which are the sorts of tasks assigned to humans anyway because they involve interpreting noisy context, long feedback loops, and balancing complex objectives.

25 In fact, on the margin, managers may prefer to assign work to humans. Humans learn valuable skills and intuitions by doing the work, workers often retain this learning better than AI models do, and managers will likely be blamed less for human errors.

26 Some in the patient group also argue that AI systems are not cost-effective to adopt in many cases. However, in a pre-print paper, Ramp's Ryan Stevens finds that businesses currently replacing freelancers with AI spend a fraction of what they used to on the relevant tasks.

27 Companies like the AI Underwriting Company and General Analysis are working to address these issues by creating new cybersecurity standards and insurance regimes for AI systems.

28 Dwarkesh Patel puts it well, arguing that "people are underrating how much company and context specific skills are required to do most jobs. And there just isn't currently a robust efficient way for AIs to pick up those skills."

29 Some in the excited group emphasize potential distributional benefits from AI, arguing that AI systems could democratize access to specialized fields like law or medicine and spread access to high-quality advice in doing so.

30 Tan argues "as AI makes it cheaper, faster, and easier to do things like analyze MRIs, draft legal documents, and write code, we should expect that the demand for radiologists' treatment plans, lawyers' counsel, and engineers' expertise will broadly increase, not decrease."

31 Economist Erik Brynjolfsson writes "We will see a 'Cambrian explosion' of new products and services as AI lowers the barrier to entry for solving the world's hardest problems. This future, however, is not inevitable. It is a choice."

32 These tools have also created new roles for humans in AI orchestration. And even these companies attempting to automate sales are hiring some human sales reps to utilize their tools and pitch their products.

33 These effects depend on the technology and market context. In the Finnish case, firms appear to have used new industrial technologies to expand into new products and market segments rather than simply automating existing production. By contrast, a well-known study by Acemoglu and Restrepo (separate from the one referenced in endnote 17) finds that adoption of industrial robots from 1990 to 2007 in the U.S. reduced local employment and wages.

34 Though the author's friend working in corporate law joked, "every company will be getting sued for securities fraud twice a week."

35 What David Brooks wrote in 2001 still holds twenty-five years later: "the new elite prefers[…] words [like] authentic, natural, warm, rustic, simple, honest, organic, comfortable, craftsmanlike, unique, sensible, sincere." The human touch is a status symbol.

36 Researchers should also continue to explore questions that move beyond which workers are "exposed" to AI and directly address how they will adapt and other second-order impacts. A January 2026 working paper by economists Sam Manning and Tomás Aguirre notes that many of the workers who might be displaced by AI are well-equipped to adapt (with high incomes, transferable skills, and so on), but that a group of workers in clerical and administrative roles is particularly vulnerable. January 2026 work by International Monetary Fund economists provides a detailed look at how AI is changing demand for skills across emerging and advanced economies.

Relatedly, policymakers and researchers should monitor impacts not only in the U.S. and Europe, but throughout the world. AI adoption remains lower in countries with lower income levels and economists forecast smaller impacts in emerging economies, which have less infrastructure to support AI adoption and where more work involves the physical world. Still, because anyone with a phone or computer can access leading AI models with relative ease, labor impacts are a global issue.

37 The author thanks and credits Ramp's Ara Kharazian, with whom this idea was discussed over email correspondence.

38 As Kharazian notes, this also includes gig workers, who have fewer existing unemployment protections.

39 For example, Narayanan and Kapoor acknowledge "massive job losses in specific occupations" as a "large-scale and systemic" risk from AI, though they do not provide a specific timeline. *Bloomberg* reports that Acemoglu believes that 5 percent of jobs are "ripe to be taken over, or at least heavily aided, by AI over the next decade." Unions have fought to slow implementation in areas like trucking, but it's unclear how effective their efforts will be.

CARNEGIE
ENDOWMENT FOR
INTERNATIONAL PEACE

1779 Massachusetts Avenue NW
Washington, DC 20036
CarnegieEndowment.org