

Statement before the Senate Committee on Banking, Housing, and Urban Affairs  
Hearing: AI and the American Dream: Promoting Innovation, Affordability, and  
American Dominance

# **The Evidence on Generative AI and the American Economy**

**Will Rinehart**  
Senior Fellow

June 11, 2026

# The Evidence on Generative AI and the American Economy

By Will Rinehart<sup>1</sup>

## Introduction

The empirical literature on generative AI has grown substantially over the past couple years. Randomized trials, natural experiments, and large-scale adoption surveys now cover a range of different contexts, including customer support, software development, legal analysis, writing, education, and firm-level outcomes. The evidence is worth examining carefully, because the popular debate has run ahead of it.

The clearest pattern across the research is that task-level gains from AI are large and heterogeneous. AI often improves the quality of an output, but not uniformly. Less-experienced workers, lower performers, and firms with fewer preexisting complementary capabilities tend to benefit the most. Yet when tasks fall outside the model's frontier, or when AI encourages multitasking or overreliance, objective quality can stagnate or worsen.

Stepping back from individual tasks complicates the story further. Firm-level and labor-market evidence shows that task-level gains do not automatically translate into measurable productivity or wage effects. Adoption is uneven, complementary investments matter enormously, and the workers most exposed to AI's substitution potential are not the same workers capturing its augmentation benefits. Understanding that gap between task-level results and economy-wide outcomes is the central analytical challenge this testimony takes up.

Four main sections help build this case. First, historical patterns of technology diffusion are reviewed. Then the current evidence on task-level impacts is discussed followed by a section on firm adoption of AI. Finally, the most recent research on wage impacts are unpacked. This testimony concludes with some practical recommendations for Congress on scenario planning and data infrastructure. And since this committee oversees the Commerce Department's Bureau of Industry and Security (BIS), an appendix is included that reviews the basic economics of compute.

## Historical Patterns

Technologies vary in how they change the economy. Some raise productivity for a time as firms adopt new tools. Others reshape entire sectors by forcing new forms of organization, investment, and complementary innovation. Still, others improve the process of discovery itself, changing how quickly new ideas can be found and applied.

Following the framing of Baily et. al (2025), technologies might be thought of as:

- Lightbulbs — Inventions that raise productivity as they diffuse through the economy, but whose effect fades once adoption is complete. They permanently increase the level of output per hour but do not permanently raise the growth rate.

---

<sup>1</sup> The American Enterprise Institute (AEI) is a nonpartisan, nonprofit, 501(c)(3) educational organization and does not take institutional positions on any issues. The views expressed in this regulatory comment are those of the authors. This document is approved for public dissemination. The document contains no business-proprietary or confidential information. Document contents may be reused by the government in developing the AI Action Plan and associated documents without attribution.

- Dynamos — General-purpose technologies (GPTs) that have broad and sustained effects because they require firms to reorganize production, redesign workflows, and develop new business models before their full productivity effects appear.
- Microscopes — Inventions that improve the process of invention itself. By making research more efficient, they can increase the productivity of R&D and potentially raise the economy’s long-run growth rate.<sup>2</sup>

Generative AI seemingly has the characteristics of both a dynamo and a microscope.<sup>3</sup>

Like a dynamo, the technology that converts electricity into mechanical power, AI is a general-purpose technology that can be deployed across industries and combined with new business practices to reorganize work. And like a microscope, it may also improve the process of invention itself. If these early results are correct, then AI may do more than deliver a one-time productivity boost. It could raise the level of productivity across the economy and increase the rate at which new ideas are produced.

But that outcome is not a given. As the authors caution, AI’s “contribution to productivity growth will depend on the speed with which that level is attained and, historically, the process for integrating revolutionary technologies into the economy is a protracted one.” Indeed, the productivity gains from earlier transformative technologies did not appear all at once.

Tractors were first adopted in the Wheat Belt of North Dakota, South Dakota, and Kansas in the 1920s, but it took another two decades for them to become common in the Corn Belt.<sup>4</sup> As Daniel Gross (2017) of the Harvard Business School explained, “The tractor first developed for narrow applications with existing complementary equipment” and “Only later did tractor technology become sufficiently general for its diffusion to be broad based and pervasive.” As he continued,

This pattern of expanding scope is consistent with other historical examples and with economic theory, which suggests that in this context, R&D will naturally progress from specific- to general-purpose variants of an innovation, and that these technical advances will (i) drive the development of additional complementary technologies, and (ii) and directly translate to an increasing scope of diffusion. Lags in diffusion can therefore be the result of holdups and market failures in R&D that stymie the generalization of existing technology.<sup>5</sup>

Electrification followed a similar pattern. Although the transition to electric power began in the 1890s, electricity did not overtake steam power in American manufacturing until around 1920. Still, it wasn’t until the 1940s that the transition was complete.<sup>6</sup> Paul David’s work remains key for understanding this transformation.<sup>7</sup> The factory built around the electric dynamo was “a long-delayed and far from automatic business. It did not acquire real momentum in the United States

---

<sup>2</sup> Baily, M. N., Byrne, D. M., Kane, A. T., & Soto, P. E. (2025). *Generative AI at the Crossroads: Light Bulb, Dynamo, or Microscope?*. Federal Reserve. <https://www.federalreserve.gov/econres/feds/files/2025053pap.pdf>

<sup>3</sup> Ibid.

<sup>4</sup> Gross, D. (2017). *Scale versus Scope in the Diffusion of New Technology: Evidence from the Farm Tractor*. <https://doi.org/10.3386/w24125>

<sup>5</sup> Ibid.

<sup>6</sup> David, P. A. (1990). The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox. *The American Economic Review*, 80(2), 355–361. <http://www.jstor.org/stable/2006600>

<sup>7</sup> Ibid.

until after 1914-17” when regulations had loosened to allow for experimentation with new methods of production.

Before the dynamo, factories had been optimized around large central engines connected to belts that drove the machinery. Electrification made a different kind of factory possible, but only if manufacturers were willing to redesign the physical layout of production. Building on David, economists Andrew Atkeson and Patrick Kehoe firmly established the point.<sup>8</sup> The largest gains came only after firms learned how to reorganize production around electric power.

The lesson for generative AI is straightforward. Productivity growth depends on the slow process of learning how to use the technology well, building the complementary systems around it, and reorganizing economic activity to take advantage of what it makes possible.

Nevertheless, history can help provide a rich set of diagnostic questions about which frictions slow adoption. Adopting the work of Atkeson and Kehoe, for example, policy makers should be asking:

- How much organization-specific knowledge have companies built up with current (non-AI) technologies? Can AI be “bolted on” to existing processes, or does it require fundamental redesign like electricity did?
- Is AI currently only marginally superior to existing approaches, or is the gap larger?
- What institutional barriers and regulatory barriers prevent companies from adopting AI even when it's superior?
- Is AI more or less embodied than electricity was? Electricity required complete factory redesign; does AI require similar organizational restructuring?
- Where is AI “embodied”? In trained models? In data pipelines? In organizational workflows?

Generative AI may ultimately prove revolutionary, but its economic effects will depend on how quickly businesses, workers, and institutions adapt around it.

## Skill and Tasks Changes

Late last year, I testified before the Joint Economic Committee and summarized the empirical literature suggested about AI and skills. The findings were encouraging. Across customer support, software development, legal analysis, writing, and creative work, generative AI was producing meaningful productivity gains. Importantly, those gains often accrued to less experienced or lower-skilled workers, suggesting that AI could equalize skill gaps.

In the months since, additional research has complicated that story in important ways. The basic pattern still holds. AI often raises average performance and disproportionately helps workers with weaker initial skills. But the newer evidence shows that AI is not simply a universal skill equalizer. It can also magnify skill differences when success depends on knowing when to use AI or strategically implementing its suggestions. In other words, AI compresses skill gaps inside the model’s frontier, but it can widen them at the frontier’s edge.

---

<sup>8</sup> Atkeson, A., & Kehoe, P. (2007). Modeling the Transition to a New Economy: Lessons from Two Technological Revolutions. *The American Economic Review*. <https://doi.org/10.3386/w8676>

Brynjolfsson, Li, and Raymond (2023) provided one of the earliest causal estimates of how large language models affect work inside a real firm.<sup>9</sup> By analyzing the staggered rollout of an AI-powered assistant to 5,179 customer support agents, the researchers found that access to the tool substantially increased productivity. Worker productivity rose by 14 percent on average when measured by issues resolved per hour. But the effects were highly uneven. Novice and lower-skilled agents saw productivity gains of roughly 34 percent, while more experienced and higher-skilled agents saw little improvement. Importantly, the study also found that the AI assistant reduced employee attrition, suggesting that workers were more able to handle frustrating customer interactions, or that they had a stronger sense of competence on the job.

This study was the first to show that AI can raise the performance of lagging workers, compress skill differences, and change workplace dynamics. Recent work from Dell’Acqua et al. (2026) both extends these insights and adds important wrinkles.<sup>10</sup> On 18 realistic business tasks that the authors classified as within current AI capabilities, AI users completed 12.2 percent more tasks, worked 25.1 percent faster, and produced substantially higher-quality output. Like Brynjolfsson, Li, and Raymond, this study found that the gains were largest for participants in the bottom half of the skill distribution.

In contrast, “for a complex managerial task selected to be outside the frontier, subjects using AI were 19% less likely to produce correct solutions compared with those without AI, pointing to potential limitations of AI supporting knowledge workers.” Even just a little training mattered. Participants who received a short overview of how to use GPT performed better than those given access to GPT alone. This is one of the most important lessons from the current literature. The returns to AI depend jointly on the task’s location relative to the model’s frontier and on whether people are taught how to use the tool.

In software development, Hoffmann, et al. (2024) exploited a natural experiment created by GitHub’s rollout of Copilot to understand how generative AI reshapes the internal organization of work.<sup>11</sup> When developers gained access to Copilot, their task portfolio shifted. They spent more time writing and modifying code and less time on peripheral project-management tasks. AI effectively pushed workers toward the core of their job.

Two mechanisms seem to explain the shift. First, Copilot increased autonomous work by handing some of the routine. Second, it encouraged exploration by lowering the cost of trying new approaches. Notably, these effects were strongest among lower-ability developers, who saw the largest gains in autonomy and exploratory coding. The result at the end of the study was a workforce whose internal task composition leaned more heavily toward the productive frontier.

Research from three randomized field experiments at Microsoft, Accenture, and an anonymous Fortune 100 company found that developers with access to an AI coding assistant completed about

---

<sup>9</sup> Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). *Generative AI at work* (NBER Working Paper No. 31161). National Bureau of Economic Research. <https://doi.org/10.3386/w31161>

<sup>10</sup> Dell’Acqua, F., McFowland, E., Mollick, E., Lifshitz, H., Kellogg, K. C., Rajendran, S., Kraymer, L., Candelon, F., & Lakhani, K. R. (2026). Navigating the jagged technological frontier: Field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality. *Organization Science*, 37(2), 403–423. <https://doi.org/10.1287/orsc.2025.21838>

<sup>11</sup> Hoffmann, M., Boysel, S., Nagle, F., Peng, S., & Xu, K. (2025). *Generative AI and the nature of work* (Harvard Business School Strategy Unit Working Paper No. 25-021; Harvard Business School Working Paper No. 25-021; CESifo Working Paper No. 11479). <https://doi.org/10.2139/ssrn.5007084>

26 percent more tasks.<sup>12</sup> Likewise, a randomized field experiment covering 7,137 workers across 66 firms found that workers who used the tool spent 31 percent less time on email, increased their concentration time, and completed collaborative documents faster.<sup>13</sup> But there were no significant changes in the amount of meeting time. This suggests that AI shifts workers away from coordination tasks and toward more focused work. But these shifts occurred in tasks that workers had ownership over, not in tasks requiring organizational changes.

AI's impact on legal research and writing seems to be evolving as the models develop. In 2023, Choi et al. conducted a randomized controlled trial with law students using GPT-4 to complete legal analysis tasks.<sup>14</sup> They found that while AI support produced only modest and inconsistent improvements in the quality of legal reasoning, it produced large and uniform gains in speed. Students completed tasks faster regardless of prior ability, but the largest improvements occurred among the least-skilled participants. In follow-up surveys, participants also reported satisfaction with using the tool, suggesting human-AI complementarity might lead to better satisfaction with work.

A newer study by Schwarcz et al. shows that model selection is key.<sup>15</sup> In April 2026, the authors published results from a randomized trial involving 137 law students that were required to complete six realistic legal tasks. The study compared a retrieval-augmented legal AI system with an early reasoning model and found that both tools produced large productivity gains. Across five of the six task types, productivity increased by 50 to 130 percent. Not surprisingly, the quality of legal work improved in at least four of the six tasks. Each of the tools improved performance in different ways. The retrieval-augmented system reduced hallucinated citations by grounding answers in legal sources, while the reasoning model improved the depth and rigor of legal analysis. The lesson here is that AI's effect on skill can vary dramatically, depending on the design of the model, the structure of the tool, and the kind of judgment the task requires.

Generative AI also seems to boost productivity in writing. Noy and Zhang (2023) ran an online experiment with 453 college educated workers who completed writing tasks tied to their occupations.<sup>16</sup> Participants who had access to ChatGPT produced noticeably better work, with average quality increasing by 18 percent, and they finished their assignments much faster, cutting task time by about 40 percent. Similarly, Hauser and Doshi (2024) found that AI can boost writing creativity.<sup>17</sup> By studying nearly 300 people who wrote short fictional stories under a controlled design, participants who received ideas generated by GPT-4 produced stories that independent

---

<sup>12</sup> Cui, Z., Demirer, M., Jaffe, S., Musolff, L., Peng, S., & Salz, T. (2024). *The Effects of Generative AI on High Skilled Work: Evidence from Three Field Experiments with Software Developers*.

<https://doi.org/10.2139/ssrn.4945566>

<sup>13</sup> Dillon, E., Jaffe, S., Immoclica, N., & Stanton, C. (2025, May). *Shifting Work patterns with Generative AI*. Harvard Business School. [https://www.hbs.edu/ris/Publication%20Files/w33795\\_dd1e2857-d195-4333-86ba-6a8953119ed4.pdf](https://www.hbs.edu/ris/Publication%20Files/w33795_dd1e2857-d195-4333-86ba-6a8953119ed4.pdf)

<sup>14</sup> Choi, J. H., Monahan, A., & Schwarcz, D. (2023). *Lawyering in the age of artificial intelligence* (Minnesota Legal Studies Research Paper No. 23-31). *Minnesota Law Review*, 109 (Forthcoming 2024).

<https://doi.org/10.2139/ssrn.4626276>

<sup>15</sup> Schwarcz, D., Manning, S., Prescott, J. J., Barry, P., Cleveland, D. R., & Rich, B. (2026). Ai-powered lawyering: Ai reasoning models, retrieval augmented generation, and the future of legal practice. *Journal of Law & Empirical Analysis*, 3(1), 220–250. <https://doi.org/10.1177/2755323x261427048>

<sup>16</sup> 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>

<sup>17</sup> Doshi, A. R., & Hauser, O. P. (2024). Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science Advances*, 10, eadn5290. <https://doi.org/10.1126/sciadv.adn5290>

evaluators judged as more creative, better written, and more enjoyable. However, there was a downside. Participants anchored their work on the suggestions provided by the model, which raised average quality but narrowed the range of creative outcomes.

There are limits, as Becker, Rush, Barnes & Rein (2025) show.<sup>18</sup> While their study is much more limited, as they conducted a randomized control trial with just 16 experienced open-source developers completing 246 tasks, it added a critical addition. The researchers asked the developers and outside experts how fast they thought the task would be completed. Developers, after completing the task, estimated that the AI tool reduced the complete time by 20 percent, while experts in economics and machine learning predicted 38 percent and 39 percent reduction in time, respectively. However, AI actually *increased* task times by 19 percent. This work makes clear that AI has real limits, especially at the upper end of the skill distribution.

In recent months, the biggest additions to the literature have been a spat of studies on how AI affects education. Bastani et al. (2026) conducted a field experiment with nearly 1,000 high-school students, finding that GPT-4 access raised immediate performance during practice.<sup>19</sup> But when the tool was later removed, students performed 17 percent worse than students who never had access. By contrast, a “GPT Tutor” reproduced the short-run gains while largely mitigating the later declines. In other words, how AI is used matters in either building capability or substituting for it.

Kim et al. (2025) studied 18,904 students on a major EdTech platform, generating 2.1 million student-day observations.<sup>20</sup> Importantly, their AI tutor activated only after students submitted an answer, limiting it to post-solution debriefing, explaining mistakes, explaining solutions, and answering follow-up questions. Treated students completed 36 percent more math problems, spent 3.9 percent less time per problem, and increased their correctness rate by 3.6 percent. Crucially, the authors traced these gains to improved learning efficiency, not just greater practice volume. As they made clear, “Heterogeneity analyses show that the AI tutor disproportionately benefits lower-performing students and does not generate larger gains for students from higher-income areas, suggesting that the technology narrows performance gaps without widening income-based disparities.”

The evidence across customer support, software development, legal work, writing, and education points in a consistent direction. AI can raise average performance and it tends to deliver the largest gains to lower-skilled workers. But there are limits. When tasks fall outside a model's current capabilities or when workers are already at the frontier, AI can hurt performance rather than help it. When workers receive no training on how to use the tools, returns shrink. And when students rely on AI to provide the answer rather than help with understanding, skill formation can suffer. The practical implication is that AI is not a uniform productivity multiplier. Its effects depend on task complexity, model design, worker skill level, and whether the tool has any structure around it.

---

<sup>18</sup> Becker, J., Rush, N., Barnes, E., & Rein, D. (2025). *Measuring the impact of early-2025 AI on experienced open-source developer productivity*. arXiv. <https://arxiv.org/abs/2507.09089>

<sup>19</sup> Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakçı, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences*, 122(26). <https://doi.org/10.1073/pnas.2422633122>

<sup>20</sup> Kim, D., Mitrofanov, D., Wen, Q., & Xu, T. (2025). *Generative AI Can Improve Performance and Engagement without Harming Learning*. <https://doi.org/10.2139/ssrn.5929576>

## Adoption of AI

Stepping up from task changes to a firm level analysis shifts the story. While improvement is often found in individual tasks, the impacts on firms are more muted.

A recent field experiment at the National Bank of Slovakia randomly assigned access to generative AI to workers to test the tech's impact.<sup>21</sup> But before running the experiment, the researchers mapped out all the work onto a standard task-based framework to understand which tasks were altered. The tool proved especially complementary to nonroutine work, where judgment, synthesis, and idea generation are central to performance.

Routine work showed a more complicated pattern. Employees in routine roles experienced some of the largest individual productivity gains, but generative AI was less effective when the task itself required high levels of structure and repetition. This produced a notable mismatch between the workers who benefit the most and the tasks for which the technology is best suited. A simulation exercise suggests that if the organization reallocated workers across tasks to better match these strengths, total output could rise by more than 7 percent.

The experiment also uncovered differences across skill levels. Lower-skill workers saw the biggest improvements in quality, while higher-skill workers mainly benefited through time savings and greater efficiency. Together, these findings provide some of the strongest empirical evidence to date that generative AI reshapes productivity through task-level complementarities rather than simple automation, with important implications for how firms organize labor and how AI may diffuse through labor markets more broadly.

Evidence from Denmark reinforces this picture. Two adoption surveys conducted in late 2023 and 2024 covered 25,000 workers across 7,000 workplaces in 11 occupations, linked to matched employer-employee records.<sup>22</sup> By the time of surveys, AI chatbot use had become widespread, employers were encouraging adoption, and many companies had deployed in-house models as well as training initiatives. Yet, in spite of this investment, the labor market effects were minimal. Using difference-in-differences estimation and employer policies as quasi-experimental variation, researchers found precise zeros across every occupation studied. Average time savings of 2.8 percent, combined with weak wage pass-through, explain the flat outcomes. In other words, productivity gains at the task level did not translate into measurable changes in pay or employment.

A 2026 working paper from the European Investment Bank offers a synoptic view because it covers more than 12,000 non-financial firms across the EU and United States.<sup>23</sup> The researchers instrumented AI adoption among European firms by assigning the adoption rates of comparable US peers. The results show that AI adoption raises labor productivity by 4 percent, a gain consistent with mid-range macroeconomic projections rather than productivity boom scenarios. Critically, the

---

<sup>21</sup> Marsal, A., & Perkowski, P. (2025). *A task-based approach to generative AI: Evidence from a field experiment in central banking*. SSRN. <https://doi.org/10.2139/ssrn.5228176>

<sup>22</sup> Humlum, A., & Vestergaard, E. (2025, April 15). *Large language models, small labor market effects*. [https://static1.squarespace.com/static/5d35e72fcff15f0001b48fc2/t/67fdc07c74f1302536156227/1744683136315/chatbots\\_apr25.pdf](https://static1.squarespace.com/static/5d35e72fcff15f0001b48fc2/t/67fdc07c74f1302536156227/1744683136315/chatbots_apr25.pdf)

<sup>23</sup> Aldasoro, I., Gambacorta, L., Pal, R., Revoltella, D., Weiss, C., & Wolski, M. (n.d.). *BIS working papers no 1325 AI adoption, productivity and employment*. European Investment Bank. <https://www.bis.org/publ/work1325.pdf>

gains stem from capital deepening rather than job displacement. AI augments worker output without reducing headcount, and workers at AI-adopting firms have so far seen higher wages.

The distribution of those gains, however, is uneven. Medium and large firms capture the bulk of the productivity benefits, while micro and small enterprises saw no statistically significant effect. Critically, the productivity dividends also depend heavily on complementary investments. Each added percentage point of spending on software raises the productivity effect of AI adoption by around 2.4 percent. Meanwhile, each additional percentage point on employee training raises productivity by 5.9 percent. Important for this committee, EU countries with less developed financial markets show structurally lower adoption and have not closed the gap that opened in 2023-24, while financially developed EU countries track US adoption rates closely.

In contrast to its most ardent boosters, AI is not tech that yields immediate returns. It is much closer to a general-purpose technology whose payoff depends on complementary assets, process discovery, and capability building.

## Wages and Inequality

While there is widespread worry that AI will lead to greater wage and income inequality, the available evidence simply doesn't support this conclusion.<sup>24</sup> For one, as a previous section underscores, AI often raises productivity, especially for less-experienced workers. But more to the point, AI has nuanced effects that end up substituting for some tasks and complementing others.

Arguably the strongest empirical study to date on how AI is affecting early-career workers comes from economists Erik Brynjolfsson, Bharat Chandar, and Ruyu Chen in a report titled, "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence."<sup>25</sup> The authors use detailed labor market data to isolate how employment patterns have shifted in occupations, which they summarized in a series of six facts:

- There have been substantial declines in employment for early-career workers (ages 22-25) in occupations most exposed to AI, such as software developers and customer service representatives.
- Workers aged 22 to 25 have experienced a 6 percent decline in employment from late 2022 to July 2025 in the most AI-exposed occupations, compared to a 6-9 percent increase for older workers.
- Not all uses of AI are associated with declines in employment. Entry-level employment has declined in applications of AI that automate work, but not those that most augment it.
- Employment declines for young, AI-exposed workers remain after conditioning on firm-time effects like interest rate changes.
- The labor market adjustments are visible in employment more than compensation.
- The above facts are largely consistent across various alternative sample constructions.

So, it's entry level workers who might be affected by AI even as the rest of the job market continues to chug along. However, it remains an open question why just entry level workers are being

---

<sup>24</sup> Harris, J. (2026, April 8). *A.I. may worsen wealth inequality*. New York Times.

<https://www.nytimes.com/2026/04/08/opinion/ai-wealth-inequality-jobs-investment.html>

<sup>25</sup> Brynjolfsson, E., Chandar, B., & Chen, R. (2025). *Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence*. Stanford Digital Economy Lab.

[https://digitaleconomy.stanford.edu/wp-content/uploads/2025/08/Canaries\\_BrynjolfssonChandarChen.pdf](https://digitaleconomy.stanford.edu/wp-content/uploads/2025/08/Canaries_BrynjolfssonChandarChen.pdf)

affected by the tech. I explained my preferred theory in an op-ed in the *City Journal*. Large language models have shifted the fundamental mechanics of how workers and employers find each other. Before generative AI, employers could rely on automated screening tools, while workers had to manually tailor applications for each job, repeatedly entering the same data.

ChatGPT changed this equilibrium by allowing job seekers to apply at scale while also enabling employers to post more low-commitment openings. The result is a noisier labor market. Because AI makes it harder for both applicants and employees to signal genuine intent, there are fewer matches for those with non-differentiated skills. On the other hand, skilled workers don't face this dilemma because they have portfolios, referrals, and networks. This shift could help explain why early-career workers may be struggling in AI-exposed fields even as everyone else is doing fine. They often lack the credible signals that experienced workers possess.

Not all research supports the findings of the Canaries paper, however. Ryan Nunn of The Budget Lab at Yale recently examined the microdata underlying monthly employment reports by workers' AI exposure and found "no strong evidence of impacts as of yet."<sup>26</sup> That conclusion is consistent with earlier work from The Budget Lab, which found no unusual increase in occupational churn. As Nunn explained, "If AI were displacing a substantial number of workers by automating their jobs, one would expect this to result in many displaced workers eventually taking new jobs in new occupations."

The difference between these two studies, as economist Jed Kolko explained earlier this year, often comes down to the type of AI exposure measures that is chosen. As he wrote, "AI exposure or usage could be correlated with other ways occupations differ, such as the extent of over-hiring during the pandemic, or suitability for remote work, or exposure to tariffs, or reliance on immigrants for workers—all of which could also explain why employment patterns have differed across occupations in recent years."<sup>27</sup>

Indeed, it is not even the case that a decline in labor's share of output necessarily means that workers will become worse off. In a scenario gamed out by Autor and Kausik (2026), if AI and automation make the economy substantially more productive, workers could receive a smaller share of a much larger economic pie and still earn higher real wages.<sup>28</sup> The important question is whether it raises productivity enough, and spreads those gains widely enough, to improve living standards. The more plausible near-term picture is uneven adjustment: some entry-level workers may face new barriers to getting started, some experienced workers may become more productive, and the overall effect on wages will depend on how quickly AI-driven productivity gains diffuse across firms, occupations, and workers.

Still, a falling labor share can still generate serious social and political tensions, especially if the gains accrue unevenly or if some workers face painful transitions. People do not experience economic change through aggregate wages or productivity statistics. They experience it through bargaining power, job security, and the fear that gains are being captured by a small number of

---

<sup>26</sup> Nunn, R. (2026, May 7). *Ai is probably not (yet) the reason for labor market weakening*. The Budget Lab. <https://budgetlab.yale.edu/research/ai-probably-not-yet-reason-labor-market-weakening>

<sup>27</sup> Kolko, J. (2026, March 10). *Research on AI and the Labor Market is still in the first inning*. PIIE. <https://www.piiie.com/blogs/realtime-economics/2026/research-ai-and-labor-market-still-first-inning>

<sup>28</sup> Autor, D., & Kausik, B. N. (2026, January 9). *Resolving the automation paradox: Falling labor share, rising wages*. arXiv.org. <https://arxiv.org/abs/2601.06343>

firms. That helps explain why anxiety about AI remains high even though the labor-market evidence is muted.

Advanced AI has created a strange political dynamic. With AI, a startup can now draft marketing copy, write code, design prototypes, and automate back-office tasks at a fraction of the old cost. This has lowered the costs of starting and scaling a business, making it easier for new entrants to test ideas, reach customers, and compete with larger firms. This access has led to a boom in new companies.<sup>29</sup> At the same time, the rise of AI has seemingly concentrated power at the infrastructure layer, since frontier models, chips, and data centers require enormous capital investment. The result is not a simple story of workers versus machines, but a broader debate over who captures the gains from a new general-purpose technology.

## Some Practical Suggestions

Given this uncertainty, Congress should focus on two practical steps: scenario planning and better data.

The point of scenario planning is not to predict exactly what AI will do to workers, families, or financial markets. It is to prepare for several plausible futures at once. In one scenario, AI might diffuse slowly and mostly augment workers. In another, it could create sharp pressure on entry-level jobs while raising productivity for experienced workers. Each scenario implies a different policy response. Congress should therefore build the capacity to track leading indicators, stress-test assumptions, and adjust as evidence changes.

This committee should consider how AI agents might impact finance and banking. AI agents introduce a qualitatively different set of risks to financial markets than earlier automation. Earlier algorithmic trading systems operated within narrow, predefined rules. But AI agents can pursue open-ended objectives, adapt to new information, and interact with each other in ways that are difficult to anticipate in advance. For example:

- If large financial institutions deploy agents built on the same underlying models, those agents may respond to market signals in similar ways simultaneously, amplifying volatility rather than dispersing it.
- Agents can execute complex, multi-step strategies faster than human oversight can follow, compressing the window for intervention when something goes wrong.
- Unlike rule-based systems, agents can fail in ways their designers did not foresee, making pre-deployment stress testing harder and post-hoc explanation more difficult.
- Agents can be instructed to find and exploit gaps in existing rules faster than regulators can close them, shifting the compliance dynamic from deliberate evasion to emergent behavior that was never explicitly authorized.

That requires better data. Today's labor-market statistics are not designed to capture how AI is changing tasks, skills, firm organization, or the flow of ideas across the economy. We often know whether someone is employed, but not whether AI has changed the work they do. Julia Lane's "Industries of Ideas" (IoI) framework is useful here because it treats ideas, workers, firms, skills, and regions as connected parts of the same system. By linking metadata to university records, state

---

<sup>29</sup> Brown, C. (2026, June 9). *Ai use is behind a boom in solo founders, Nasdaq finds*. Axios. <https://www.axios.com/2026/06/09/ai-entrepreneurs-founders-nasdaq>

workforce data, higher education agency data, and job postings, the Iofl approach allows for tracking how workers trained on federal research dollars actually flow into and through the economy. Rather than measuring AI only through broad occupational exposure scores, policymakers should support data infrastructure that can show how ideas move from research into firms, how firms translate those ideas into new products and jobs, and how workers acquire the skills needed to participate. In a world where AI is changing work, the most important policy asset may be the ability to see change as it happens.

## Appendix: Some Economics of Compute

There's a folk theory that the data center craze is fake, a scheme cooked up by the tech oligarchs to burn capital. But the buildouts of data centers are downstream of a much larger trend. The compute market has emerged.

Compute refers to all the hardware needed to train, deploy, and run artificial intelligence models. It has become one of the most important economic inputs of the AI era, sitting between the semiconductor fabs that produce the chips and the AI services that consumers and businesses actually use. Yet the economics of compute remain undertheorized.

Tokens comprise the compute market. Tokens are the chunks of text that AI models process as inputs and produce as outputs. Every token requires compute to process. More tokens means more matrix multiplications, more memory movement, more chip time. Token volume is what drives compute demand, and compute demand is what is filling up those data centers.

While compute can be bought on the open market, AI companies are largely selling a bundle that includes the model and the compute required to run it. In the nomenclature of economics, compute is an intermediate good because it is not consumed as a final good but is used in the production of other services.

In one way, compute is a rental rate for computational resources. When OpenAI charges by the million tokens, or Anthropic prices Claude by model tier, the underlying economic variable is compute in the form of chip time, memory, power, cooling, networking, and the engineering systems needed to make all of that available in real time. Tokens are the meter, but compute is the underlying commodity.

Compute has both stock and flow components. Training compute creates the model, and is a large, lumpy investment in an intangible asset. A company spends enormous sums on chips, data, researchers, engineers, and experimentation to produce a model that can then be deployed. Inference compute is different. It is the continuing flow of computational resources required to turn that model into a service. Every prompt consumes inference compute and every answer requires it.

This is what makes AI so different from traditional software. Software has high fixed costs and near-zero marginal costs. Once Microsoft Word has been written, distributing another copy is almost free. Once Facebook's code has been built, adding another user is relatively cheap. There are still servers, storage, and moderation costs, but the magic of software is that it can be produced once and sold many times.

Artificial intelligence systems are not like software products of prior decades. They are expensive to build and expensive to run but most importantly, the costs scale with the ambition of the model. Every additional user, every additional prompt, every additional generated answer consumes real resources. The marginal cost may fall, and it has been falling quickly, but it is being swamped by demand for compute. AI is therefore closer to a capital-intensive utility than to classic software. It has some of software's scale advantages, but it also has the recurring input costs of an industrial service.

Additionally, users expect responses in seconds and consistent uptime, even under peak demand. In turn, that means providers need spare capacity, redundancy, routing systems, and buffers. A system

built only for average utilization will buckle under real-world demand. Compute must be available where and when it is needed.

AI infrastructure is different from ordinary cloud infrastructure. Traditional cloud workloads can often be delayed, distributed, or smoothed across time. Some AI workloads can be treated that way, especially batch processing and offline training. But consumer-facing and enterprise-facing inference is a real-time service. The result is that AI providers must build for peaks, not just averages. Idle capacity is not necessarily waste, it is part of the product.

All of these features have pushed the industry toward scale. The firms best positioned to compete are those that can secure chips, negotiate power contracts, build or lease data centers, optimize model serving, route traffic across model tiers, and keep utilization high without degrading user experience. A small AI startup can rent compute from the cloud, but it is exposed to price spikes, capacity shortages, and dependence on the infrastructure decisions of larger firms. A hyperscaler can spread those risks across regions, customers, chips, and workloads.

Every major AI company is now facing a tough question: either they vertically integrate or they outsource. The economics of this choice are not obvious. Integration gives a firm control over its supply chain and can lower long-run costs, but it also shifts decision rights in ways that can dull the incentives of non-owning parties to invest. A specialist firm, by contrast, lives or dies by its efficiency in that one function. This logic pushes toward outsourcing when specialized investment is central to production, and in AI infrastructure, it very much is. Data center operations, chip design, and hardware optimization all require deep, specialized human capital that generalist AI labs are not well positioned to cultivate.

That logic helps to explain the deals that have been inked. They are hybrid structures that blend owned capacity, long-term contracts, and cloud rental in proportions that shift with market conditions. Vertical integration offers control over supply and potentially lower long-run costs. Outsourcing preserves flexibility and lets specialist firms absorb the operational complexity of running chips at scale. Frontier AI companies are doing a little of both, which will shape the competitive structure of the AI industry for years.

This also helps explain why China has embraced open-weight AI. U.S. export controls have limited Chinese firms' access to top-end NVIDIA chips, forcing them to rely on downgraded chips like the H800 and to squeeze more performance out of constrained hardware. When you do not have ready access to the most advanced compute and cannot vertically integrate, it makes sense to compete on model diffusion, efficiency, and application rather than on closed frontier scale alone. Open models are a rational response when integration isn't available.

The market structure that follows from this is unusual. Compute has commodity-like features because chip time can be priced, rented, and metered. But it is not a simple commodity. Capacity differs by chip type, model architecture, memory, interconnect, region, latency, energy cost, and software stack. An H100 in one cluster is not perfectly interchangeable with a TPU in another or a Groq chip optimized for a narrow class of inference workloads. Compute is standardized enough to be bought and sold, but differentiated enough to create strategic advantage.

As such, compute is becoming a general-purpose input for cognitive work in the same way electricity became a general-purpose input for mechanical work. Electricity can power a factory, a refrigerator, a server farm, or a lightbulb. Compute can power a chatbot, a coding agent, a drug-

discovery platform, a legal assistant, a design tool, or a military targeting system. Its value depends on the complementary systems built around it.

That is why the economics of compute cannot be separated from the economics of adoption. Businesses aren't buying the best models, they are buying the resources needed to make cheaper customer support, faster software development, better search, automated document review, fraud detection, workflow management, or new products. Compute demand will therefore depend not just on model capability but on organizational change.

This creates a two-sided bottleneck. On the supply side, the AI economy needs chips, data centers, power, cooling, interconnection, and model-serving systems. On the demand side, it needs firms that know how to turn AI services into productivity gains. If supply is constrained, AI services become expensive and unreliable. If adoption is constrained, compute gets built before enough valuable uses are ready. The market will be shaped by the interaction of both bottlenecks.