



Statement before the Joint Economic Committee
On “Frontier Technologies, Industrial Efficiency, and Pro-Innovation Policies”

Gains, Frictions, and Skill Equalization: How AI Is Reshaping the U.S. Economy

Will Rinehart

Senior Fellow, American Enterprise Institute

Tuesday, November 18, 2025

Introduction

Three years ago this month, OpenAI's ChatGPT 3.5 was released. By January 2023, the service had nearly 100 million users, making it the fastest-adopted consumer technology in history. Since then, the pace of innovation has only accelerated. Both OpenAI and its nearest rival Anthropic have released a suite of products and are actively working to develop AI agents. Meanwhile, a class of new foundational models trained on data other than language like Evo 2 which was trained on the DNA of more than 100,000 species and RT-X, which was trained on robotics data, represent a widening of the frontier.

With such quick advances, it is no surprise that some have predicted this will be “the most significant economic transformation since the Industrial Revolution.”¹ Unlike previous tech which only affected specific sectors like media, retail, or manufacturing, AI has the potential to impact virtually all industries and labor markets by automating cognitive work and enabling new capabilities.

This testimony offers a general overview of how generative AI is reshaping the U.S. economy, workforce, and regulatory landscape. In the next four sections, I make the case that:

- AI should be understood as a potential general-purpose technology, a technology capable of transforming tasks, jobs, firms, and markets, but whose diffusion will be slowed by the same frictions, transition costs, and organizational barriers that characterized earlier breakthroughs like electricity and tractors.
- Empirical research across multiple domains consistently shows that generative AI boosts productivity especially for lower-skilled workers. In other words, AI is a skill equalizer. Yet the gains are uneven, with high-skill workers often benefiting less and, in some cases, even experiencing slower performance with AI tools. Adoption across U.S. businesses is rising but still concentrated in tech-heavy sectors, suggesting the economy remains in the early stages of the diffusion curve.
- The labor market signals are mixed. AI-exposed occupations show employment declines among early-career workers even as older cohorts remain stable. One likely explanation is that this is a story of hiring frictions. AI-driven changes in job search make matches harder for inexperienced workers lacking strong signals.

In the final section, I argue that, as Congress contends with these changes, they should look first to existing legal systems. From my advantage, the real risk comes from poorly designed state laws creating a costly patchwork that stifles innovation. The article concludes by arguing that Congress should adopt narrow federal preemption to prevent a fragmented regulatory landscape and should consider pilot programs and sandboxes to allow new AI tech.

General Purpose Technologies And A Framework

Economists typically make a distinction between specific technologies that fit narrow needs and general-purpose technologies (GPTs) that have a wide range of applications and can reshape broad swaths of the economy. Steam power, electricity, semiconductors, and the Internet are all canonically considered

¹ Korinek, A. (2025). *The Economics of Transformative AI*. NBER.
<https://www.nber.org/reporter/2024number4/economics-transformative-ai>

GPTs.² But not all GPTs have the same economic impact. Electricity famously produced enormous productivity gains once factories reorganized around it but that reorganization took decades. Broadband Internet, by contrast, has been much harder to tie to clear growth effects in the data.³ GPTs are powerful, but their actual economic payoff is inconsistent.

To understand these variations, most researchers tend to root their analysis on the task-based approach. Consider truck driving. A truck driver doesn't just drive a semi. He or she loads cargo, verifies paperwork, secures loads, performs safety checks, and communicates with dispatch. Each employer, be it FedEx, Amazon, or J.B. Hunt mixes these tasks differently depending on its business model. When a firm adopts a new technology, it is almost never replacing an entire job. It is amplifying, altering, or reorganizing specific tasks. That shift alters which workers that the firm needs, how capital is deployed, and how the firm competes.

In other words, new technologies change the tasks that make up jobs. Tasks are bundled into productive combinations to make a job. Jobs combine with capital equipment to form firms. And firms compete within markets. Thus, to understand a new technology, especially a potential GPT like AI, we need to understand how AI changes tasks, and how these shifts reverberate throughout the ecosystem.

Technological change is never free. It requires the buying of new machines, the installation of new software, and the reorganization of production lines. And with the explicit cost comes the implicit cost of the transition period, when the new process is not yet better than the old one. This is the essence of opportunity cost. Firms adopt a technology only when expected gains outweigh these combined costs, which can be substantial.⁴

These frictions and costs can help to explain why GPTs can be slow to diffuse. Tractors, for example, were initially adopted for use in the Wheat Belt of North Dakota, South Dakota, and Kansas in the 1920s. But it took another 20 years for them to become a common technology in the Corn Belt of Iowa, Illinois, and Nebraska.⁵ Daniel Gross (2017) of the Harvard Business School explained,

The tractor first developed for narrow applications with existing complementary equipment, exogenously high demand, and lower [research and development (R&D)] costs, and initial diffusion was accordingly rapid for these applications, but otherwise limited in scope. Only later did tractor technology become sufficiently general for its diffusion to be broad based and pervasive. 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

² Bresnahan, T. F., & Trajtenberg, M. (1995). General purpose technologies 'engines of growth'? *Journal of Econometrics*, 65(1), 83–108. [https://doi.org/10.1016/0304-4076\(94\)01598-t](https://doi.org/10.1016/0304-4076(94)01598-t)

³ Stanley, T. D., Doucouliagos, H., & Steel, P. (2018). Does ICT generate economic growth? A meta-regression analysis. *Journal of Economic Surveys*, 32(3), 705–726. <https://doi.org/10.1111/joes.12211>

⁴ Holmes, T. J., Levine, D. K., & Schmitz, J. A. (2012). Monopoly and the incentive to innovate when adoption involves switchover disruptions. *American Economic Journal: Microeconomics*, 4(3), 1–33. <https://doi.org/10.1257/mic.4.3.1>

⁵ Gross, D. (2017). *Scale versus Scope in the Diffusion of New Technology: Evidence from the Farm Tractor*. <https://doi.org/10.3386/w24125>

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.⁶

Electricity also took time to diffuse. Though the transition to electricity started in the 1890s, it didn't overtake steam power until around 1920. It then took another 20 years for the transition to be complete.⁷ Paul David (1990) is indispensable for understanding this change since it summarizes a decade of economic history on the transition from water and steam power to electrical power, the dynamo.⁸ The rise of the electric factory with their whirling dynamos, as David explained it, was "a long-delayed and far from automatic business. It did not acquire real momentum in the United States until after 1914-17" when regulation changes allowed entrepreneurs to experiment with new methods of production.

Electrification took decades because factories had perfected steam-powered layouts that they were reluctant to abandon. In follow on research partially based on David, Andrew Atkeson and Patrick J. Kehoe explained that manufacturers were reluctant to close existing plants and lose this knowledge for what, initially, is only a marginally superior technology.⁹ Learning how to use new technologies efficiently took decades. New technologies or knowledge are physically embedded in specific forms, like machines, factories, or organizations. Realizing their benefits requires building or redesigning those material systems.

These historical lessons help explain why generative AI is attracting such attention from economists and policymakers. Like past general-purpose technologies, AI today is not a single product but a range of different applications, each flexible, and increasingly capable of doing complex tasks. This generative capacity allows the technology to substitute for and complement cognitive labor in ways that earlier digital tools could not. Generative AI is not just faster software. It can produce ideas, draft solutions, and reshape workflows, provided that firms can overcome all of the same frictions, transition costs, and organizational barriers that will come with AI.

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 clearly 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?

⁶ Ibid.

⁷ 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>

⁸ Ibid.

⁹ 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>

If generative AI is to function as a true general-purpose technology, its impact must extend beyond impressive demos and headline-grabbing capabilities. It must alter the underlying mechanics of production, which is the focus of the empirical research reviewed in the next section.

Skill And Tasks Changes

Empirical research has begun to clarify how generative AI is reshaping work and productivity. These studies offer a more grounded picture than the speculative debates a decade ago. One consistent finding is that generative AI can produce substantial productivity gains in a variety of settings like customer support, legal work, and software development, among others. But this is especially true for low skilled workers. A table detailing the current literature is attached.

Brynjolfsson, Li & Raymond (2023) provided one of the earliest causal estimates of how large language models reshape work inside real firms.¹⁰ This paper analyzed the staggered rollout of an AI-powered assistant for 5,179 customer support agents to find that access to the tech substantially increases worker productivity. Measured by issues resolved per hour, the average increase was 14 percent. The effects, however, were far from uniform. Novice and lower-skilled agents saw dramatic gains, with productivity rising roughly 34 percent, while experienced, higher-skilled agents saw little change.

The study also found that the availability of the AI assistant reduced employee attrition. Workers who used the tool were less likely to quit, which is a pattern consistent with reduced stress, fewer frustrating interactions, or a stronger sense of mastery over daily tasks. So it seems that not only did generative AI make agents more productive, it also appeared to make the job itself more manageable and rewarding.

This paper matters because it demonstrates not just efficiency gains, but distributional effects. AI can boost the output of lagging workers and can meaningfully improve workplace dynamics. But this pattern is not just confined to customer support. In software development, legal reasoning, professional writing, and creative work, generative AI is consistently altering which tasks workers perform and how skill advantages translate into output.

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.¹¹ 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.

¹⁰ 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>

¹¹ 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>

In the legal domain, AI access seems to speed up the process. Choi et al. (2023) conducted a randomized controlled trial with law students using GPT-4 to complete legal analysis tasks.¹² 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.

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.¹³ 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.¹⁴ 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 evaluators judged as more creative, better written, and more enjoyable. Once again, the strongest gains appeared for participants who were less naturally creative. 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.¹⁵ 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%, while experts in economics and machine learning predicted 38% and 39% reduction in time, respectively. However, AI actually *increased* task times by 19%. This work makes clear that AI has real limits, especially at the upper end of the skill distribution.

The domain of education remains a blackbox. It is known that students and teachers are using generative AI at rapidly rising rates, but rigorous causal studies remain rare. Roldán-Monés (2024) is an exception. In this randomized controlled trial, students prepping for a university debating competition were randomly assigned GenAI to help prepare. In a result that bucked the trend, high-ability students actually benefit more from generative AI than lower-ability students.¹⁶

¹² 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>

¹³ 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>

¹⁴ 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>

¹⁵ 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>

¹⁶ Roldán-Monés, T. (2024). *When GenAI increases inequality: Evidence from a university debating competition* (Working paper). EsadeEcPol. https://www.esade.edu/ecpol/wp-content/uploads/2019/09/2409-ChatGPTRoldan_ecpol.pdf

One area that seems especially promising involves workers who are neurodivergent. A recent study from the United Kingdom's Department for Business and Trade found that neurodiverse employees were significantly more enthusiastic about AI assistants than their neurotypical peers.¹⁷ They reported higher satisfaction with the tools and were about one quarter more likely to recommend them to others. Interviews from people with ADHD, dyslexia, and autism reveal similar patterns.¹⁸ Many report that AI systems help them break down and organize complex tasks, improve the clarity of their written communication, and manage cognitive load in ways that traditional workplace tools do not. In other words, generative AI could meaningfully expand the range of workers who can thrive in certain roles, raising important questions about inclusion and accessibility.

Still, there is very little evidence on the best way to train people to use these tools effectively. Most firms are experimenting in real time, often relying on informal peer learning rather than structured training programs, and the education sector has only begun to test targeted curricula. Without clearer data, it is difficult to know whether current approaches are capturing the full potential of the technology or leaving significant value on the table. Developing rigorous research on AI training and skill acquisition will be essential for understanding how quickly generative AI can diffuse through the workforce and which workers will benefit the most.

Taken together, the emerging evidence points to several broad lessons about skill development in the age of generative AI. The technology consistently delivers its largest gains for workers with weaker initial skills, both by accelerating routine tasks and by improving the quality of their output. It also appears to reshape task portfolios, pushing many workers toward more productive activities and away from administrative or peripheral work. At the same time, the benefits are uneven. High performers who rely on deeply ingrained heuristics or finely tuned workflows often see smaller improvements, and in some cases AI tools can even slow them down.

Still, we still know relatively little about how people learn to use these tools effectively, how organizations should structure training, or how the technology might enable entirely new forms of work for groups such as neurodivergent employees. But the central lesson is that generative AI is a skill equalizer. It strengthens some workers, challenges others, and requires a deeper understanding of how people acquire and deploy skills in environments where machines can produce ideas and perform reasoning steps alongside them.

How Companies Are Reacting

The U.S. Census Bureau's Business Trends and Outlook Survey (BTOS) provides the most comprehensive real-time picture of AI adoption across the American economy. Released biweekly, the BTOS captures responses from approximately 1.2 million businesses through a rotating panel design. The survey's AI supplement, included in recent cycles, measures the use of machine learning, natural language processing, virtual agents, and voice recognition technologies, and is available by sector, state, and the 25 most populous metropolitan areas.

¹⁷ Department for Business and Trade. (2025). *The evaluation of the M365 Copilot pilot in the Department for Business and Trade*. <https://assets.publishing.service.gov.uk/media/68adbe409e1cebdd2c96a19d/dbt-microsoft-365-copilot-evaluation.pdf>

¹⁸ Curry, R. (2025, November 8). *People with ADHD, autism, dyslexia say AI agents are helping them succeed at work*. CNBC. <https://www.cnbc.com/2025/11/08/adhd-autism-dyslexia-jobs-careers-ai-agents-success.html>

Since September 2023, the share of U.S. businesses using artificial intelligence in producing goods or services has more than doubled, rising from 3.7% to 9.9% as of September 2025. While this represents consistent growth over the two-year period, the vast majority of businesses, some 83% of them, still report no AI usage in their operations. Meanwhile, about 7% of businesses indicate they do not know whether they use AI technologies. That data is displayed below.

U.S. Businesses Using AI

Share of U.S. businesses reporting AI use in production of goods or services. Examples of AI include machine learning, natural language processing, virtual agents, voice recognition, etc.




Chart: Will Rinehart • Source: BTOS • Created with Datawrapper

AI adoption varies dramatically across U.S. industry sectors. As expected, the information sector leads adoption with 29.9% of businesses reporting AI use as of September 2025, up from 13.9% in September 2023. This sector includes publishing, software, telecommunications, and data processing, all industries on the cutting edge of the AI race. Professional, scientific, and technical services follows at 19.8%, and finance and insurance at 17.8%.

In contrast, most sectors remain in single digits. Traditional industries show particularly low adoption including construction (3.0%), retail trade (4.2%), and accommodation and food services (2.9%). The disparity suggests AI adoption is concentrated in knowledge-intensive and technology-oriented industries, while labor-intensive and consumer-facing sectors have been slower to integrate these technologies into their operations.

U.S. Businesses Using AI By Sector

Share of U.S. businesses reporting AI use in production of goods or services by industry sector. Examples of AI include machine learning, natural language processing, virtual agents, voice recognition, etc.

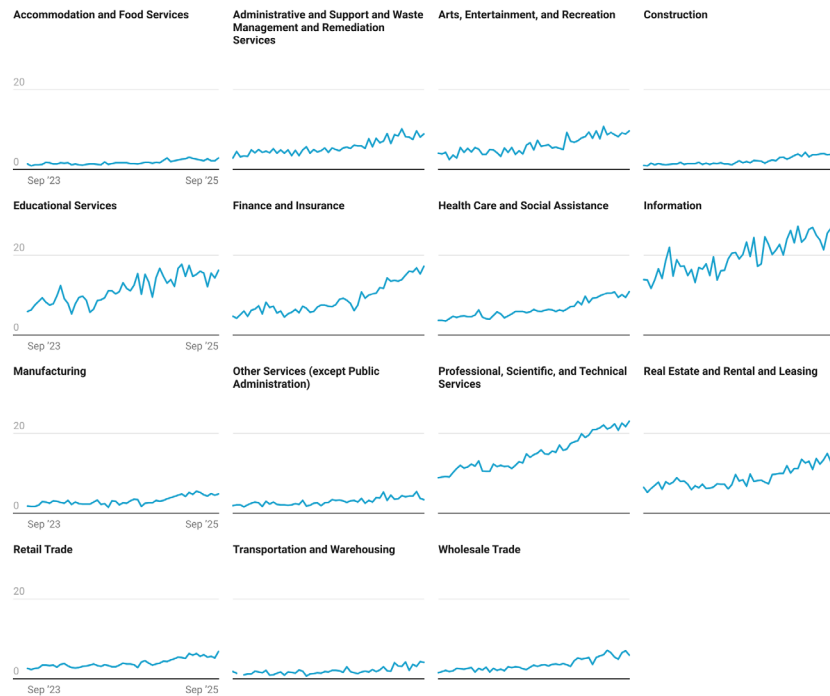


Chart: Will Rinehart • Source: BTOS • Created with Datawrapper

While aggregate adoption remains uneven, new experimental evidence shows what happens inside firms once generative AI is actually deployed. A recent field experiment at the National Bank of Slovakia randomly assigned access to generative AI to workers to test the tech's impact.¹⁹ 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%.

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.

¹⁹ 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>

The Labor Market

In late summer 2025, a chart from the Federal Reserve Bank of New York went viral for showing that recent computer engineering grads had the third worst market prospects.²⁰ It followed alongside a spat of reports about the dismal job market for recent college graduates. The episode quickly became a focal point for worries that AI might already be reshaping early-career opportunities in unexpected ways.

As Connor O'Brien noted in a careful review of the underlying data, the chart's alarming implications stem largely from very small subsamples within the American Community Survey (ACS). As he explained,

While the ACS is itself a large survey, the sub-samples highlighted in the chart — young college graduates with very specific majors — are quite small. As a result, the confidence intervals around these estimates are huge.

Take the estimated 7.5 percent unemployment rate for computer engineering majors, for example. Using special weights provided by the ACS, we can calculate the confidence interval on this estimate. We find that we can say with 95 percent confidence that the unemployment rate for new computer engineering graduates is somewhere between four percent and 11 percent. For perspective, the most we can feel comfortable saying is that the job market for those graduates looks like something between the absolute height of the Great Recession and the booming job market of 2019, just before the pandemic. That's it.²¹

Two weeks after that came out, economists Erik Brynjolfsson, Bharat Chandar, and Ruyu Chen published what is arguably the strongest empirical study to date on how AI is affecting early-career workers, a report titled, "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence."²² Their analysis brings much-needed clarity to a debate that had been dominated by anecdotes and viral charts. 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% decline in employment from late 2022 to July 2025 in the most AI-exposed occupations, compared to a 6-9% 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.

²⁰ *The labor market for recent college graduates*. Federal Reserve Bank of New York. (2025).

<https://www.newyorkfed.org/research/college-labor-market#--:explore=outcomes-by-major>

²¹ O'Brien, C. (2025, August 13). A viral chart on recent graduate unemployment is misleading.

<https://agglomerations.substack.com/p/a-viral-chart-on-recent-graduate>

²² 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

- The above facts are largely consistent across various alternative sample constructions.

Given all of the available evidence, 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 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.

I'm not the only one thinking about this problem through the matching framework. Economist Jack Meyer recently mapped out the Beveridge curve of two different cohorts, college graduates from 2010 to 2020, and graduates from 2021 to 2025 ,to show how generative AI may be degrading the efficiency of the hiring process itself.²³ The Beveridge curve plots the relationship between job vacancies and unemployment, and when it shifts outward, it signals that the labor market is becoming worse at turning openings into hires. Meyer argues that this is exactly what we are now seeing for recent graduates as his chart below explains.



These findings point to a labor market that is not collapsing under the weight of automation, but one that is becoming noisier, slower, and less efficient at matching young workers to opportunities. But the

²³ Meyer, J. (2025, November 9). *AI might not be taking your job anytime soon*. AI Might Not Be Taking Your Job Anytime Soon. <https://jackbmeyer.substack.com/p/ai-might-not-be-taking-your-job-anytime>

implications extend well beyond hiring. The forces reshaping entry-level jobs are also bringing about calls for regulatory action. The next section examines how policymakers are already responding—and how that response is shaping the emerging AI economy.

How Congress Should Respond

These dynamics matter for policy. As states rush to legislate on AI, they are doing so in a landscape where employers, workers, and innovators are all grappling with rapid technological change. The wrong regulatory moves could intensify the very frictions that are already making it harder for young workers to find their footing.

It isn't as though the United States lacks a regulatory architecture capable of responding to AI. As analyst Adam Thierer has emphasized, we already have countless overlapping regulatory regimes to discipline new technologies without Congress having to build an entirely new system from scratch.²⁴ Consumer protection law, intellectual property, torts, product recall authority, insurance markets, and agency adaptation all operate as real guardrails. These mechanisms are already shaping how AI is deployed, and they are often far nimbler and more targeted than broad mandates.

The status quo is not perfect, but it is simply not the case that we live inside a regulatory vacuum. Before rushing to legislate, Congress should recognize that much of the regulatory scaffolding is already in place. Instead of constructing sweeping new frameworks, policymakers should clarify these laws and extend their authority where appropriate. Overreacting now may do more harm than good. Indeed, the single biggest risk to AI innovation is the proliferation of state laws that undermine the innovation they claim to protect.

Colorado's experience with its 2024 AI Act is illustrative. Initially, Colorado estimated minimal fiscal impact for its pioneering AI Act, claiming that compliance could "be accomplished with existing appropriations."²⁵ Then, over successive fiscal notes, the cost estimate rose to nearly \$650,000 per year just before passage.²⁶ After the Colorado General Assembly passed the bill, Governor Jared Polis signed the bill into law, but he did so with a scathing critique of it in a signing statement:

I am concerned about the impact this law may have on an industry that is fueling critical technological advancements across our state for consumers and enterprises alike. Government regulation that is applied at the state level in a patchwork across the country can have the effect to tamper innovation and deter competition in an open market.²⁷

Since the passage of the bill, its cost has skyrocketed and problems with implementation have become

²⁴ Thierer, A. (2023). *Flexible, pro-innovation governance strategies for artificial intelligence* (R Street Policy Study No. 283). R Street Institute.

²⁵ Legislative Council Staff. (2024, April 12). *Fiscal note for SB 24-205: Consumer protections for artificial intelligence*. Colorado General Assembly.
https://leg.colorado.gov/sites/default/files/documents/2024A/bills/fn/2024a_sb205_00.pdf.

²⁶ Legislative Council Staff. (2024, August 12). *Fiscal note for SB 24-205: Consumer protections for artificial intelligence*. Colorado General Assembly.
https://leg.colorado.gov/sites/default/files/documents/2024A/bills/fn/2024a_sb205_f1.pdf.

²⁷ Polis, J. (2024, May 17). *Signing statement on SB24-205*. Office of the Governor, State of Colorado.
<https://drive.google.com/file/d/1i2cA3IG93VViNbZxu9LPgbTrZGqhyRgM/view>.

apparent. In a special session, Colorado Budget Director Mark Ferrandino now estimates that it will cost between \$2.5 million and \$5 million annually to implement.²⁸ Colorado Governor Jared Polis believes it may reach \$6 million per year. Lawmakers have been debating a major overhaul of this bill due to its flawed construction.²⁹ Not surprisingly, Polis was also one of the few governors to support the proposed AI moratorium in Congress earlier this year because it would “give Congress time to figure this out and create a true 50-state solution to smart AI protections for consumers while driving innovation.”³⁰

Yet, Polis isn’t the only policymaker warning against a rush to regulate AI. In testimony before the Connecticut General Law Committee, Daniel O’Keefe of Connecticut’s Department of Economic and Community Development Building cautioned that “Introducing a novel, untested regulatory framework creates ambiguity in our statutes and courtrooms, and may inadvertently create opportunities for discriminatory activity that is already illegal to evade best efforts at enforcement.”³¹ Continuing, he explained:

Moreover, even the most carefully crafted regulations, with the very best of intentions, can create an unintended chilling effect for businesses and innovators. Faced with the legal obligation to interpret, periodically reinterpret, and continuously meet extensive reporting obligations under the unique laws of only one jurisdiction, most businesses will default to shifting their operations elsewhere. Compliance with the new regulations proposed by this bill would cost innovative companies substantial amounts of time and money, and would constitute an especially prohibitive barrier for those considering establishing operations in Connecticut.³²

Connecticut’s internal estimation produced eye-watering numbers, estimating that compliance with the act is likely to cost \$3 million or more for the budget year.³³ In both Colorado and Connecticut, what began as an effort to ensure responsible AI quickly devolved into a bureaucratic tangle that diverts public funds and private capital away from innovation and funnels it toward compliance and litigation risk.

Since so many states have been proposing AI bills, I’ve been experimenting with large language models (LLMs) to conduct quick cost compliance estimates. In each piece, I prompted the leading LLMs to act as a compliance officer, then read the laws, and estimate the hours needed for first-year implementation and ongoing compliance. Using this method, compliance with New York’s RAISE Act could cost companies between 1,070 and 2,810 hours in the first year, effectively requiring a full-time employee to

²⁸ Benson, J., & Heck, Z. (2025, September 5). *Colorado gives businesses breathing room before AI act takes effect*. Taft Privacy & Data Security Insights. <https://www.taftlaw.com/news-events/law-bulletins/colorado-gives-businesses-breathing-room-before-ai-act-takes-effect>.

²⁹ Sealover, E. (2025, May 7). *Late-night legislative drama attempts to extend deadline for implementation of artificial-intelligence regulations*. The Sum & Substance. <https://tsscolorado.com/late-night-legislative-drama-attempts-to-extend-deadline-for-implementation-of-artificial-intelligence-regulations/>.

³⁰ Ferenstein, G. (2025, June 2). *Governor Polis signed Colorado’s restrictive AI law, but supports a federal moratorium on similar legislation*. Reason Foundation. <https://reason.org/commentary/governor-polis-signed-colorados-restrictive-ai-law-but-supports-a-federal-moratorium-on-similar-legislation/>.

³¹ O’Keefe, D. (2025, February 26). *Testimony on Senate Bill 2, An Act Concerning Artificial Intelligence*. Connecticut General Assembly, General Law Committee. <https://www.cga.ct.gov/2025/gldata/TMY/2025SB-00002-R000226-OKeefe,%20Daniel,%20Commissioner-DECD--TMY.PDF>.

³² Ibid.

³³ Office of Fiscal Analysis. (2024, August 12). *Fiscal note for SB 2: An Act concerning artificial intelligence*. Connecticut General Assembly. <https://www.cga.ct.gov/2025/FN/PDF/2025SB-00002-R000603-FN.PDF>.

keep up with the law.³⁴ California's AB-1018, which didn't pass the legislature, suggests individual firms could face between \$2 million and \$6 million over the course of a decade.³⁵ Congress should consider ways that they could help build these models for their own use to quickly estimate cost compliance, including a pilot project at the Congressional Budget.

Regardless, the cost to innovation in complying with countless states will act as an invisible tax, draining resources from research, hiring, and product development to navigate inconsistent reporting regimes. Indeed, when President Trump announced the AI Action Plan, he pointed out this problem:

If you are operating under 50 different sets of state laws, the most restrictive state of all will be the one that rules. So you could have a state run by a crazy governor, a governor that hates you, a governor that's not smart or maybe a governor that's very smart but decides that he doesn't like the industry and he can put you out of business.³⁶

Congress does not need to micromanage AI development, nor should it attempt to build a sweeping new administrative superstructure. What it *must* do is prevent a regulatory arms race among states. Following this, leaders in Congress should be thinking through what a narrowly tailored federal preemption bill would entail. It would provide the clarity and uniformity the economy needs.

Meanwhile, lawmakers should look for narrowly tailored ways to *enable* AI to enter the economy where it can unlock clear productivity gains. Recently, I wrote about especially in domains like contracting, where legal rules artificially block efficiency improvements.³⁷ Congress does not need to rewrite contract law or upend state authority, but it could create carefully scoped safe harbors, pilot programs, or regulatory sandboxes that allow firms to test AI-assisted drafting, evidence generation, or compliance monitoring.

Conclusion

Generative AI is already reshaping how Americans work, how firms organize production, and how young workers enter the labor market. The evidence shows remarkable productivity gains, especially for lower-skilled workers. At the same time, AI adoption remains uneven across industries. Congress does not need to construct an entirely new regulatory superstructure to meet this moment. In fact, doing so would likely create more problems than it solves. The real danger lies in allowing a patchwork of conflicting state AI laws to calcify into a de facto national regime that stifles innovation.

³⁴ Rinehart, W. (2025, June 10). *New York's RAISE Act misses AI safety risks*. City Journal. <https://www.city-journal.org/article/new-york-raise-act-artificial-intelligence-safety>.

³⁵ Rinehart, W. (2025, September 9). *The hidden price tag of California's AI oversight bill*. Exformation. <https://exformation.williamrinehart.com/p/the-hidden-price-tag-of-californias>.

³⁶ Hendrix, J. (2025, July 24). *Transcript: Donald Trump's address at 'Winning the AI Race' event*. Tech Policy Press. <https://www.techpolicy.press/transcript-donald-trumps-address-at-winning-the-ai-race-event/>.

³⁷ Rinehart, W. (2025, October 29). *Will AI agents make the perfect contract?*. Will AI Agents Make The Perfect Contract? <https://exformation.williamrinehart.com/p/will-ai-agents-make-the-perfect-contract>

Attachment

Empirical Studies on AI Impacts

Authors	Title	Summary		
Sarkar (2025)	'AI Agents, Productivity, and Higher-Order Thinking: Early Evidence From Software Development'	<ul style="list-style-type: none">Analyzed the impacts of an AI-assisted programming platform during its rollout.More experienced workers are more likely to accept agent-generated code.One standard deviation in work experience corresponds to 6% higher accept rates.Software output increased by 39% after the agent became the platform's default code generation mode.Experienced workers are more likely to develop plans in their initial messages to agents. These results suggest that abstraction, clarity, and evaluation may be important skills for workers.	Hauser and Doshi (2024)	<ul style="list-style-type: none">Generative AI enhances individual creativity but reduces the collective diversity of novel content <ul style="list-style-type: none">Conducted an online experiment with nearly 300 people who were asked to write short, eight-sentence stories. Participants were divided into three groups: People who wrote stories without any AI help, people who could request one story idea from GPT-4, and people who could request up to five different story ideas from GPT-4. A separate group of 600 evaluators read and rated the stories without initially knowing which ones had AI assistance.Stories written with AI assistance were rated as more creative, better written, and more enjoyable than those written without AI. Having access to multiple AI ideas (up to five) produced the biggest improvements. Less naturally creative writers benefited the most from AI assistance.Stories written with AI assistance were more similar to each other than stories written by humans alone. Writers using AI tended to anchor their stories to the AI's suggestions, making them less unique.
Kumar, Khare, Sharma, Kumar, Saini, Yadav, Jain, Rana, Verma, Meena & Edubilli (2025)	'Intuition to Evidence: Measuring AI's True Impact on Developer Productivity'	<ul style="list-style-type: none">Analyzed how 300 engineers across multiple teams integrated an in-house AI platform (DeputyDev) that combines code generation and automated review capabilities.Pull request reviews saw a 31.8% reduction in time. Adoption scaled from 4% engagement in the first month to 83% peak usage by the sixth month, stabilizing at 60% active engagement.Top adopters achieved a 61% increase in code volume pushed to production. The AI tool accounted for an 28% increase in code shipment volume.	Brynjolfsson, Li & Raymond (2023)	<ul style="list-style-type: none">Generative AI at Work <ul style="list-style-type: none">Studied the staggered introduction of a generative AI-based conversational assistant using data from 5,179 customer support agents.Access to the tool increases productivity, as measured by issues resolved per hour, by 14% on average, including a 34% improvement for novice and low-skilled workers but with minimal impact on experienced and highly skilled workers.Our results suggest that access to generative AI can increase productivity, with large heterogeneity in effects across workers.AI assistance improved customer sentiment, increased employee retention, and may lead to worker learning.
Becker, Rush, Barnes & Rein (2025)	'Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity'	<ul style="list-style-type: none">Conducted a randomized controlled trial (RCT) with 16 developers completing 246 tasks to understand how AI tools affect the productivity of experienced open-source developers.Before starting tasks, developers predicted that allowing AI will reduce completion time by 24%. After completing the study, developers estimated that allowing AI reduced completion time by 20%. Meanwhile, experts in economics predicted 39% shorter times and machine learning experts predicted 38% shorter times.Surprisingly, this study finds that allowing AI actually increases completion time by 19%.	Choi, Monahan & Schwarcz (2023)	<ul style="list-style-type: none">Lawyering in the Age of Artificial Intelligence <ul style="list-style-type: none">Conducted a randomized controlled trial with law school student to study the effect of AI assistance on human legal analysis.Access to GPT-4 only slightly and inconsistently improved the quality of participants' legal analysis but induced large and consistent increases in speed.AI assistance improved the quality of output unevenly: the lowest-skilled participants saw the largest improvements. On the other hand, AI assistance saved participants roughly the same amount of time regardless of their baseline speedIn follow up surveys, participants reported increased satisfaction from using AI to complete legal tasks and correctly guessed the tasks for which GPT-4 were most helpful.
Maršál and Perkowski (2025)	'A Task-Based Approach to Generative AI: Evidence from a Field Experiment in Central Banking'	<ul style="list-style-type: none">Randomly assigned generative AI access to central bank employees at the National Bank of Slovakia.Generative AI access leads to large improvements in both quality and efficiency for the majority of participants.While workers in routine jobs experience larger individual performance gains, generative AI is less effective for the routine task content of their work.The divergence between task-level and worker-level productivity effects under generative AI is substantial. Simulation results indicate that reallocating tasks across workers to reflect these shifted comparative advantages could raise organizational output by roughly 7.3%.Low-skill workers benefit most in terms of quality while high-skill workers benefit in terms of efficiency. Our findings provide empirical support on generative AI and task-level complementarities, with important implications for how generative AI will impact workers, organizations, and labor markets more broadly.	Noy & Zhang (2023)	<ul style="list-style-type: none">Experimental evidence on the productivity effects of generative artificial intelligence <ul style="list-style-type: none">Conducted an online experiment with occupation-specific, incentivized writing tasks to 453 college-educated professionals, randomly exposing half of them to ChatGPT.The average time taken decreased by 40% and output quality rose by 18%.Workers exposed to ChatGPT during the experiment were 2 times as likely to report using it in their real job 2 weeks after the experiment and 1.6 times as likely 2 months after the experiment.
Dillon, Jaffe, Immorlica & Stanton (2025)	'Shifting Work Patterns with Generative AI'	<ul style="list-style-type: none">Conducted a 6-month randomized field experiment with 7,137 workers where half received access to generative AI tools integrated with email, documents, and meetings.AI primarily impacted tasks workers could change independently.Workers using AI tool frequently spent 31% less time on email (3.6 fewer hours weekly). Documents were completed moderately faster but there was no significant change in meeting time.	Peng, Kalliamvakou, Chon, & Demirel (2023)	<ul style="list-style-type: none">The Impact of AI on Developer Productivity: Evidence from GitHub Copilot <ul style="list-style-type: none">Conducted a controlled experiment using GitHub Copilot to implement an HTTP server in JavaScript as quickly as possible.The group with access to Copilot completed the task 55.8% faster than the control group.
Schwarcz, Manning, Barry, Cleveland, Prescott & Rich (2025)	'AI-Powered Lawyering: AI Reasoning Models, Retrieval Augmented Generation, and the Future of Legal Practice'	<ul style="list-style-type: none">Conducted a randomized controlled trial with upper-level law students to complete six legal tasks using a RAG-powered legal AI tool (Vincent AI), an AI reasoning model (OpenAI's o1-preview), or no AI.AI assistance significantly boosts productivity in five out of six tested legal tasks, with Vincent yielding statistically significant gains of approximately 38% to 115% and o1-preview increasing productivity by 34% to 140%.There were strong effects in complex tasks like drafting persuasive letters and analyzing complaints. o1-preview improved the analytical depth of participants' work product but resulted in some hallucinations, whereas Vincent AI-aided participants produced roughly the same amount of hallucinations as participants who did not use AI at all.	Cui, Demirel, Jaffe, Musloff, Peng & Salz (2025)	<ul style="list-style-type: none">The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers <ul style="list-style-type: none">Conducted randomized controlled trials with 4,867 developers at Microsoft, Accenture, and an anonymous Fortune 100 company. A random subset of developers got access to an AI-based coding assistant.Though each experiment is noisy, the analysis reveals a 26.08% increase (SE: 10.3%) in completed tasks among developers using the AI tool.Less experienced developers had higher adoption rates and greater productivity gains.
Hoffmann, Boysel, Nagle, Peng & Xu (2024)	'Generative AI and the Nature of Work'	<ul style="list-style-type: none">Exploited a natural experiment arising from the deployment of GitHub Copilot "to investigate the impact of AI technology on the task allocation of software developers within a quasi-experimental regression discontinuity design."Access to Copilot shifts task allocation towards coding activities and away from non-core project management activities.There are two underlying mechanisms: an increase in autonomous rather than collaborative work, and an increase in exploration activities rather than exploitation. The main effects are greater for individuals with relatively lower ability.		
Caplin, Deming, Li, Martin, Marx, Weidmann & Ye (2024)	'The ABC's of Who Benefits from Working with AI: Ability, Beliefs, and Calibration'	<ul style="list-style-type: none">Conducted a controlled experiment to understand how ability and belief calibration jointly determine the benefits of working with AI.AI improves performance more for people with low baseline ability. Holding ability constant, AI assistance is more valuable for people who are calibrated, those that have accurate beliefs about their own ability. People who know they have low ability gain the most from working with AI.Eliminating miscalibration would cause AI to reduce performance inequality nearly twice as much as it already does.		