Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.
ISSN 1936-2854 (Print)
ISSN 2767-3898 (Online)

Artificial Intelligence and the U.S. Economy: Current Evidence

Nick Carollo, Leland D. Crane, Daniel Garcia Molina, and Joseph Nichols

2026-074

Please cite this paper as:
Carollo, Nick, Leland D. Crane, Daniel Garcia Molina, and Joseph Nichols (2026). "Artificial Intelligence and the U.S. Economy: Current Evidence," Finance and Economics Discussion Series 2026-074. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2026.074.

NOTE: Staff working papers in the Finance and Economics Discussion Series (FEDS) are preliminary materials circulated to stimulate discussion and critical comment. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. References in publications to the Finance and Economics Discussion Series (other than acknowledgement) should be cleared with the author(s) to protect the tentative character of these papers.

Artificial Intelligence and the U.S. Economy: Current Evidence[1]

Nick Carollo, Leland Crane, Daniel Garcia Molina, and Joseph Nichols

Abstract

We assess the current evidence on how AI-related investment and the use of AI are affecting the U.S. macroeconomy. Specifically, we document the boom in data center capital investment, the current rate of adoption of AI tools in the economy, the impact of the AI-related investment surge on prices, and discuss whether AI tools have affected the labor market and productivity yet. Driven by current and expected future demand for AI services, investment in data centers, both the physical structures and semiconductor chips and other IT equipment inside the structures, has climbed sharply. AI adoption is growing rapidly both among consumers and firms, though most firms are currently not using AI very intensively or at all. AI model-producing firms have also grown extremely fast, and now account for a noticeable fraction of software industry gross output. Prices have climbed in sectors that provide direct inputs to the data centers, such as memory chips. Effects on the labor market so far appear small and somewhat mixed. Similarly, the effects of AI on aggregate productivity growth also appear to be small so far, and the recent pickup in productivity likely reflects factors that predated the diffusion on generative AI.

[1] We benefited from helpful comments and suggestions from many Board staff, including Stephanie Aaronson, Dave Byrne, Tomaz Cajner, Paul Lengermann, Logan Lewis, Andrew Paciorek, Ivan Vidangos, Joshua Blonz, Jake Orchard, and Daniel Villar. Adi Zytek provided valuable assistance in the preparation of the charts. Most of the analysis in this paper was completed in June 2026. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Federal Reserve Board or the Federal Reserve System.

1. Introduction

Rapid progress in AI model capabilities and increased adoption of the technology have raised interest in the effects on the economy, both now and in the future. In this paper we assess the current imprint of AI developments on macroeconomic outcomes in the United States.[2]

Technological progress and innovation are normal features of modern economies. General purpose technologies (GPTs), however, are different from normal innovations in that they contribute significantly to aggregate productivity growth for many years. The characteristics of a GPT are (1) wide diffusion, (2) significant related knock-on innovations, and (3) continued improvement in the core technology. As discussed in Baily et al. (2025), AI appears to have all three characteristics: Adoption is high and growing; related products like coding assistant interfaces, search tools, and chips are undergoing increased innovation; and the core AI models are improving rapidly. In addition, AI tools appear already to be directly speeding up general research and development, potentially further raising productivity and technological progress.[3] Although it has been only three and a half years since ChatGPT was launched, and there is considerable uncertainty about the magnitude and timing of the economic effects of AI, it seems likely that it will have large and persistent productivity effects.

To date, we have seen AI's imprint in the U.S. economy in several areas. Investment in data centers and their related equipment has surged and will likely step up further, as the buildout of many announced data centers intensifies. Meanwhile, AI-related revenues and adoption have both grown rapidly. With strong demand we have seen the prices for some goods important for data centers surge recently, and they may rise further for a time as demand outstrips production capacity. Software prices have also risen, which may partly reflect the incorporation of AI features. These effects have been large enough to boost headline inflation.

So far, the impact of AI on labor markets and productivity is harder to discern, especially at the aggregate level. There is evidence that AI is having some effect on early career workers and some narrow occupations, but most of the transition is likely still ahead of us. At the same time, U.S. productivity growth has risen in recent years, but that change appears mostly to predate the current AI innovations. Productivity growth may further increase over the next several years as data center buildout continues and workers and firms learn to make more effective use of the technology. In addition, labor markets might see a significant transformation in response to widespread adoption of AI, as

[2] Throughout our focus is primarily on *generative* AI, which is largely responsible for the increased attention to the field.

[3] See for example the recent apparent resolution of the Navier-Stokes Millenium Prize problem (OpenAI 2026, Clay Mathematics Institute 2026), the development and use of AlphaFold for protein research (Jumper et al. 2021), and the use of AI coding tools to accelerate AI research itself (Toner et al. 2026).

demand for certain types of workers shifts, the task composition of jobs evolves, and aggregate labor demand changes. By changing cost structures, higher productivity could also have an ameliorating impact on inflation. However, these future effects of AI are beyond the scope of this paper.

The rest of the paper is organized as follows: Section 2 provides a high-level overview of the AI economy focusing on the major AI firms. The investment in physical infrastructure that powers AI applications, along with some of the underlying risks, is discussed in section 3. Section 4 covers the revenue of AI providers and the adoption of AI tools in the wider economy. Section 5 briefly covers recent data on price effects, while sections 6 and 7 discuss the state of the literature on observed and anticipated labor market and productivity effects of AI. Section 8 concludes.

2. The AI Economy

Recent AI developments have been centered on generative AI, and large language models (LLMs) in particular. The leading LLMs typically require massive resources to train and deploy. In the U.S. training and development are dominated by a handful of companies, or "AI labs" like OpenAI, Anthropic, Google DeepMind, Meta, and SpaceX/xAI. There are also open-source models—many developed in China—whose capabilities are advancing quickly and seem to be gaining market share, though the degree of actual adoption in the US is still low.[4]

The AI labs use large quantities of computing capacity. This capacity is developed by the labs themselves or purchased or leased from other firms—in particular the hyperscaler cloud providers such as Amazon and Microsoft. Data center capacity is used for *research, training,* and *inference*. Research involves experiments with new data and algorithms to find improvements (better capabilities, methods to train at lower cost, methods to serve models to customers at lower cost, etc.) relative to the previous generation of models. Research activity often takes the form of many small-scale experimental training runs to test new methods, with the successful methods incorporated into the final large-scale training run for a new model. Research activity has been highly productive; the available benchmarks suggest that models are improving steadily (and by some measures accelerating) on both cost and capabilities. For example, Gundlach et al. (2026) estimate that algorithmic improvements are making models three times more efficient each year, and METR (2026) documents that new models are able to complete

[4] March 2026 polling data from Epoch shows that fewer than 6% of AI-using consumers use open source models. Ramp AI data shows that less than 6% of covered businesses use the platforms associated with open source models, a figure that has risen modestly since the start of the year. However, data from low-cost routing platforms like OpenRouter and Vercel show large recent gains in market share for open models. These platforms tend to serve sophisticated AI-intensive users and may be an early signal for broader adoption trends.

longer and longer software development tasks. While it is difficult to translate measured capabilities into real world productivity (see sections 5 and 6), the persistent growth in proficiencies and falling costs have led to significant real-world adoption (see section 4).

Training is the process of actually fitting a new model. LLMs require large upfront costs to fit; the initial training process for a single state-of-the-art LLM can use tens of thousands of interconnected GPUs in a single data center running for weeks or months, along with large quantities of CPU time and other inputs. Once fit, the model is more or less frozen and does not really learn from subsequent interactions, though this is an area of intense research.[5]

Inference is the process of serving the trained model to customers, which is relatively low-cost per instance but requires running many instances to meet demand. Typically, inference means the user sending a prompt to the model and the model generating a response. The labs sell inference to consumers (e.g., ChatGPT subscriptions) and businesses. Some businesses in turn build products that add value to the basic inference service for sale to other businesses and consumers (e.g., Microsoft Copilot uses OpenAI models but adds integrations to interact with Microsoft products). Research and training costs can together be considered "model development" costs: These are fixed in that they do not depend on the quantity of AI products/services sold, while inference costs are the marginal cost of serving additional customers.[6]

Given the complexity of the AI ecosystem, it is useful to show where AI-related expenditures should appear in BEA's national accounts if properly measured:

- **AI model development** (the research and training discussed above) is recorded under business fixed investment (BFI) as *R&D investment* by the AI firm. This is the same treatment as standard software developed for sale to others.
- **Data center investment** is recorded as BFI under *nonresidential structures investment* (for the building) and *equipment investment* (for the chips and other equipment). As discussed in section 3, there is significant imported content in equipment for data centers.
- **Domestic sales of AI products**
    - **Sales to consumers** are recorded as *PCE*.

[5] LLM-powered tools sometimes provide a limited form of learning by using "scaffolds": tools and instructions that allow the LLM to essentially write notes to itself which are then loaded into future conversations. But this approach limits the information that the tool can reliably recall.

[6] Research, model training, and inference costs are all significant. One analysis estimated research to be the biggest expense for OpenAI in 2024, with inference coming in second and final training third.

- **Sales to businesses** may be recorded under BFI as *software investment* by the purchasing firm; alternatively some portion of sales may be recorded as *intermediate expenditures* on data hosting and cloud services.[7]

- **Exports of AI products** at this stage are dominated by strong demand for U.S. high-tech goods, which are used in foreign data centers to meet AI demand. While some sales of AI subscriptions could be recorded in services exports, others may be captured elsewhere, as sales by foreign affiliates of U.S. multinationals.[8]

On the income side of the national accounts, higher productivity would show up as either higher wages or higher profits for workers and firms benefiting from the use of AI, though there may also be reallocation of surplus between different types of workers and firms, as well as between workers and firms. Both on the expenditure side and the income side there is considerable uncertainty about how AI activity will be recorded in practice, and it is possible some AI activity will go unmeasured, perhaps to a greater degree than is typical, given how it is produced and sold.[9] Indeed, the 1990s IT boom raised a number of similar measurement issues that led to material revisions subsequently; see Landefeld and Fraumeni (2001).

3. Data Centers and Investment

The dramatic increase in demand for computing power from AI firms has triggered a huge increase in plans for new domestic data centers. These plans totaled almost $300 billion in 2024 and over $1 trillion in 2025, as shown in the left panel of figure 1.[10] Construction of some large data centers is underway, and we are also seeing a complementary wave of investment in high-tech equipment. We estimate that about one-

[7] For the treatment of software as an investment good see the discussion in the BEA's NIPA handbook on pages 6-3 and 6-4 here. See also Bontadini et al. (2026) for more on how AI enters the national accounts.

[8] Under the residence principle, subscriptions supplied by foreign affiliates of U.S. AI companies are not U.S. exports, they are production of the country in which the affiliate is located. However, associated intrafirm payments to the U.S. parent may be recorded as U.S. exports or income receipts.

[9] For example, the BEA may be missing a substantial part of so-called "own-account" equipment investment, specifically the work that data center operators are doing to assemble servers themselves rather than purchasing assembled units. Rapid progress in AI model capabilities may also make measurement of quality improvements and prices difficult.

[10] Our data on plan values appear to reflect total spending costs—that is, both outlays for structures and equipment. This conclusion is based on our comparisons of the values of the largest projects with public announcements.

third of the investment in data centers goes to the structure itself while most of the remainder is high-tech equipment.[11]

Brandsaas et al. (2025) provide a framework that shows how project-level plans can be used to generate spending forecasts using assumptions about future plans, abandonment rates, and construction timelines[12]. The assumptions about abandonment rates and construction timelines are estimated from historical project-level data on data centers. Figure 2 shows a forecast based on announced data center plans as of May 2026 under different assumptions about the inflow of new data center plans going forward: a high-end scenario in which plans double the recent pace, a low-end scenario in which new plans shrink to a quarter, and a medium scenario in which the inflow is similar to the recent pace[13]. In this middle scenario, this method suggests that the level of total measured data center-related investment will step up considerably, from about $100 billion in 2024 to over $880 billion by the end of 2027, exceeding 2 percent of GDP. Tracing the overall effect of the data center buildout on GDP is complex for many reasons. One is that a sizeable share of the data center equipment is sourced from abroad, so the increase in investment is offset by the higher imports, with our estimate of the import share to be about 50 percent.[14] Another is that the massive surge in data center spending is likely to have some spillovers into other types of investment. For instance, since data centers are very energy intensive, their growth should boost investment in power-related structures, and we have already seen a jump in new power plans near the end of 2025, as shown in the right-hand panel of figure 1. However, it's also possible that the surge in data center investment crowds out some non-AI-related investments. One primary channel for the crowding out is if the surge in the costs for the high-tech equipment reduces some high-tech spending by non-AI firms. But we have seen little evidence to date of non-AI investment growth falling below its long-run average pace.

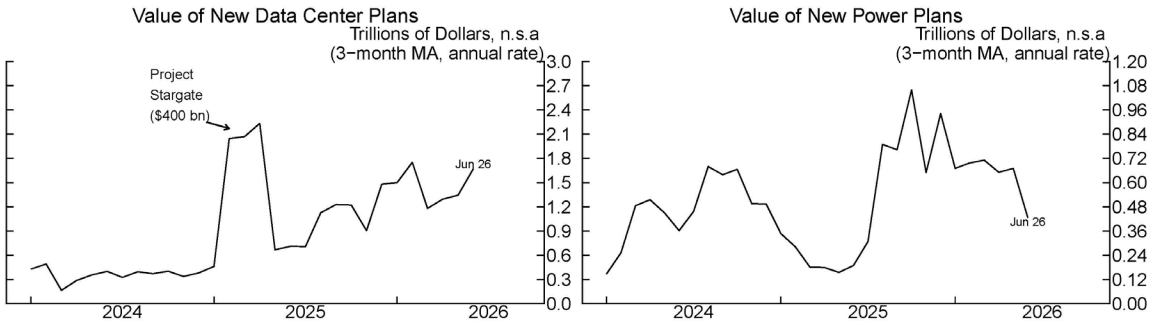
[11] We have also seen a significant increase in the output of domestic manufacturing production in sectors exposed to data center construction, based on the industry exposure measures in Braun and Kamal 2026.

[12] Note that the approach in Brandsaas et. Al (2025) is to generate path of investment working forward using the proposed dollar value of newly announced projects in the Dodge Construction database, as opposed to the approach in Van Nieuwerburgh (2026) which generates a path of investment working backwards projects in a different database for proposed delivery of computing power, measured in GWs.

[13] Updates of the analysis in Brandsaas et al. 2025 are posted at https://github.com/eirikbrandsaas/DataCenterPublic.

[14] Some analysis has assigned significantly higher import shares to the high-tech equipment, sometimes reaching 80 percent. These estimates are usually based on the higher share of imports of high-tech equipment for consumer goods or intermediate uses instead of the lower share we see for capital investment goods which is around 60 percent. These estimates also do not include the impact of the capitalized services component, which is nearly all domestic and lowers the import share to around 50 percent.

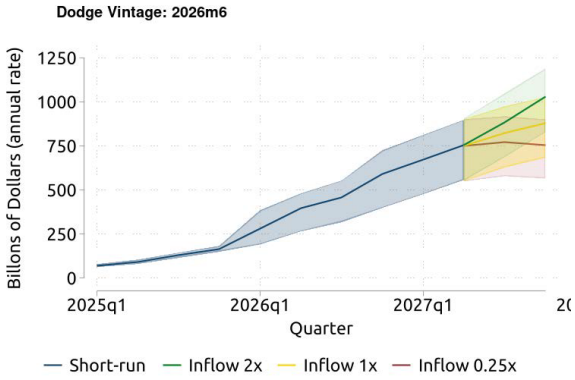
Figure 1



Source: Dodge Data & Analytics.
Notes: Plans are plotted at the date they first appear in Dodge data; dollar values reflect the most recent project estimates. Arrow denotes Project Stargate, a joint venture led by Open AI to build AI infrastructure, with a planned investment of $400 billion in the Dodge data (value not annualized).

Source: Dodge Data & Analytics.
Notes: Plans are plotted at the date they first appear in Dodge data; dollar values reflect the most recent project estimates.

Figure 2

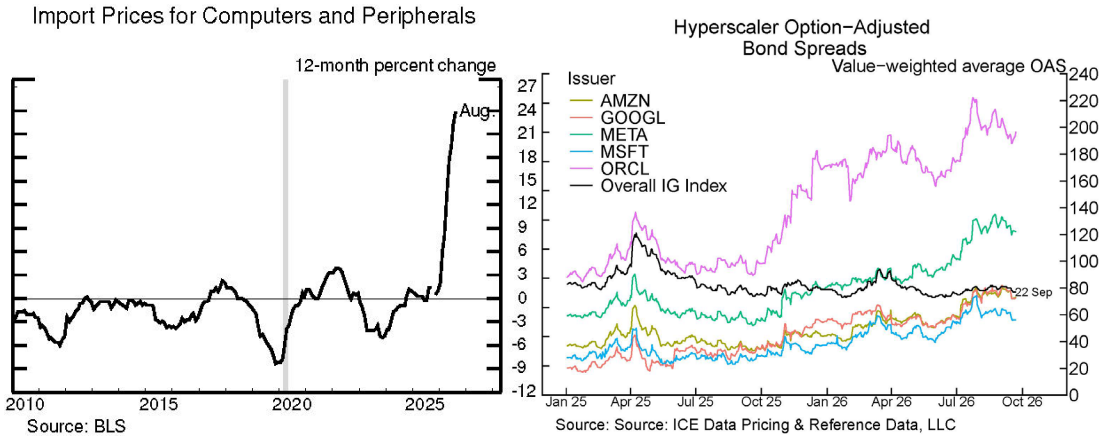


See "Estimating Aggregate Data Center Investment with Project-level Data" by Brandsaas et al. for details. The views expressed in this paper are solely those of the authors and do not necessarily reflect the opinions of the Federal Reserve Board or the Federal Reserve System.

The overall magnitude of the eventual buildout using this forecasting framework is of course highly uncertain and depends on highly debatable assumptions about abandonment rates and time to build. Available data and recent developments are sending mixed signals. On one hand, plans and equipment investment have come in strong over the past year. At the same time, there are increasing signs of supply frictions and increased local opposition to projects over the same period, suggesting that planning times will be longer for these newer projects or that they will never start at all. One reason to doubt that all of the plans will come to fruition is that these data centers require many inputs, including land, power, water, labor, and materials (in particular high-tech equipment), all of which must be financed using internal cash or external funds, which are increasingly being raised through debt. Shortages of any of these inputs would create supply frictions, driving up the cost of new data centers and delaying their completion. Among the challenges:

- Data centers are subject to state and local zoning laws that determine where they can be built. Recent concerns about the impact of the new centers on the demand for power and water have led some state and local governments to restrict data center projects.
- Data center operations seem likely to be held back by constraints related to power generation and lengthy grid interconnection queues. Updates to the electric grid are typically much slower to plan, approve, and build than the data centers themselves. The ability to locate in markets with more regulatory flexibility (e.g., Texas) or excess power availability (e.g., rural areas) may help mitigate this risk, and some projects may also build dedicated power generation co-located with data centers and, at least initially, unconnected to the local grid. There are limits, though, to both of these mitigation approaches.
- Shortages of labor and equipment for data center construction are another potential headwind. The left panel of figure 3 shows that import prices for computers and peripherals are surging, and there are reported shortages of chips and natural gas turbines (with delivery times increasing from two to more than five years).

Figure 3



- Hyperscalers initially financed data center investment out of their significant preexisting cash flows. Over the past year they have pivoted to debt financing, both via traditional corporate bonds and more complex off-balance sheet structures secured by long-term lease agreements. The surge in the volume of AI-related debt-issuance has increased concern that this could represent another channel for crowding out. While this is possible, spreads for non-hyperscalers have been mostly flat over the same period and there are a range of other factors, including but not limited to the potential impact of AI on overall economic growth, likely pushing up long-run yields in general. Oracle (purple line in the right panel of figure 3), which has the highest leverage of all hyperscalers, has

seen a significant increase in bond spreads over the last year, while bond spreads for the others have steadily inched up closing the gap with the overall investment grade spread (black line). Investors still seem confident in the overall hyperscaler business model, but they are clearly monitoring and responding to differences and changes in potential risk. A shock to investor confidence could result in a rapid deterioration in these bond spreads that would in turn limit the ability of these firms to finance further data center expansion. The ability of the hyperscalers to recoup their investments and satisfy their debt service payments is contingent on the continued adoption of AI in the broader economy and the revenue stream of the AI model providers. We turn to these topics in the next section.

4. AI Revenue and Adoption

4.A Revenue

The AI labs collect revenue from sales of AI products to consumers and businesses. Invariably these products access the LLM via an application programming interface (API): The model stays on the servers of an AI lab or a cloud provider, receiving messages and data from customers and sending replies. Some AI lab revenue comes from subscriptions to products like ChatGPT and Claude Code. These products are web pages or local software that provide an interface and tool set for the AI model. AI labs and their cloud partners also sell direct API access to the AI, especially to businesses that want to build their own products around the models.

The complexity of the AI ecosystem makes it hard to measure total AI-related sales and revenue. For example, Microsoft and OpenAI have a complex business relationship, so summing AI-related revenue for both would result in double counting. However, the fact that some of the largest AI labs (OpenAI, Anthropic, and Google DeepMind) are upstream of a large share of generative AI activity means that we can estimate a *lower bound* by focusing on these firms. Based on public disclosures, we estimate the annualized revenue of these three frontier labs was at least $107 billion in 2026Q3.[15] We emphasize (1) that these estimates are based on unaudited claims and so should not be treated as highly reliable, and (2) this is likely a lower bound, because we don't have good estimates of DeepMind's relevant revenue, we are only focusing on a subset of the AI labs, and we don't have figures for the most recent months. The lower-

[15] Statements put OpenAI's revenue run rate at $40 billion and Anthropic's at $65 billion around August. A lower bound of $2 billion for DeepMind is based on a claim of 8 million Gemini Enterprise seats (from the 2025Q4 earnings call) and a minimum price of $21 per month per seat. Revenue run rates, which are annualized monthly revenue for the most recent month, are potentially more useful than 12-month revenue when growth is extremely rapid.

bound $107 billion figure already amounts to about 19 percent of gross output in the software sector, and about 0.22 percent of private industry gross output.[16]

While AI lab revenue is significant, the labs are not currently consistently profitable. Anthropic disclosed an adjusted operating profit for 2026Q3 but doesn't expect a full year of profitability for some time, while OpenAI doesn't project profitability until 2030.[17] The firms have significant costs, which can be broken down into (1) inference costs, (2) training and research for currently operational models, (3) training and research for the next generation of models, and (4) overhead. There is little hard information on this breakdown, but both OpenAI and Anthropic have recently stated their gross margins on inference are roughly 30-40 percent.[18] In other words, roughly 65 cents out of every dollar earned goes to the direct inference computing costs of processing the prompts (primarily electricity, amortized data center costs, and payments for rented compute), leaving 35 cents to cover research, training, and overhead. The labs would be profitable if the gross margin could cover the training and overhead costs listed above.

Importantly, the labs are rapidly scaling up research and training efforts, so even if the cash flow from a given model covers all costs for that model, net cash flow will likely be negative as the next generation of larger models is researched and trained. Eventually, profitability would have to come from some combination of higher gross margins (e.g. falling inference costs and/or higher prices for AI products), larger scale (so fixed overhead and research/training costs can be spread over a larger volume of inference) or lower research/training costs. The high growth in demand for AI products and continued efficiency gains in computing make each of these possibilities plausible, but by no means inevitable. If demand for the tools falls (if, say, businesses decide that the productivity gains are overstated) or cheap open-source models are able to undercut the labs, then the current business model may not succeed. Longer-run profitability will depend in large part on breadth and intensity of adoption, which we turn to next.

4.B Adoption

The data show that AI adoption is high and growing, but often limited to narrow applications or less frequent use. AI is used both by households and businesses. On the household side, recent survey evidence shows about 50 percent of adults using AI outside of the workplace (Bick et al. 2025), most often for gathering information and writing

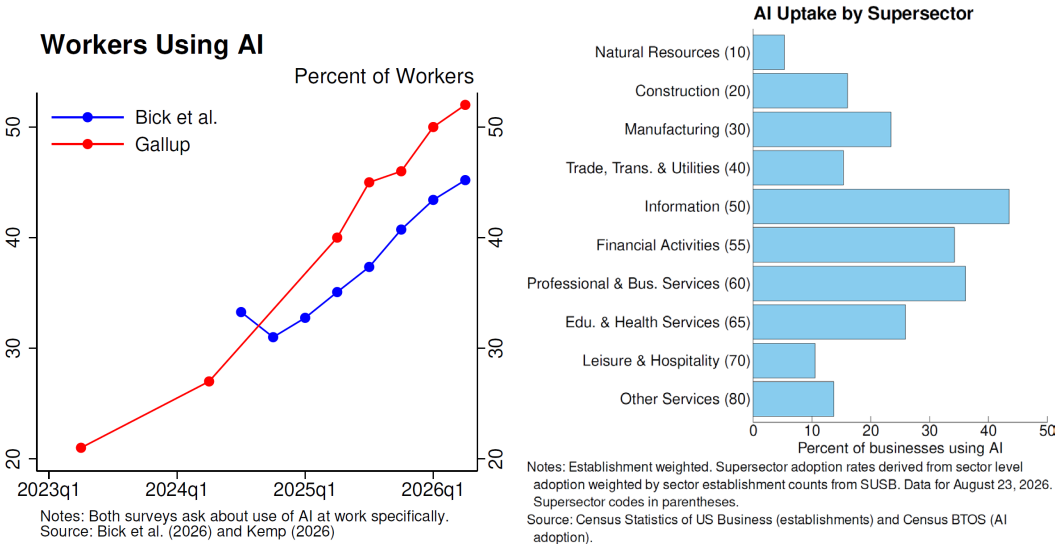
[16] These figures are based on BEA 2024 gross output by industry ($550 billion for software, $50.7 trillion for total private industry), which does not subtract out the value of intermediates and is therefore comparable to revenue estimates. AI firms may also be classified as data processing and hosting; their annualized revenue accounts for 25 percent of that industry's gross output. 2024 is the most recent year with detailed industry data.

[17] See here and here for Anthropic's recent statements, and here for OpenAI.

[18] See here for OpenAI, here for Anthropic. Gross margins can be highly variable and sensitive to the exact method of calculation. For example, Anthropic also recently claimed an 80 percent gross margin.

(Chatterji et al. 2025). Turning to business adoption of AI, surveys of firms and of workers agree that AI use in the workplace is fairly high and increasing but generally restricted to a limited number of applications. Multiple surveys of workers show more than 40 percent of workers currently using AI in the workplace, a number that generally rose over the course of 2024 and 2025 (figure 4, left panel). On the other hand, survey evidence shows that only 10 percent of workers are using AI every day at work, so AI is not yet a core part of most workers' daily routine, though younger and more educated workers tend to use AI more often (Bick et al. 2025).[19]

Figure 4



Business surveys show more variability, with the Census Business Trends and Outlook Survey (BTOS) putting firm-side AI adoption at about 20 percent, Baslandze et al. (2026) estimating 58 percent, and Yotzov et al. (2026) estimating 78 percent; other estimates generally fall somewhere in this range. There is evidence that the exact wording of questions can be an important factor in surveys of measured adoption, as may be the sampling frame (e.g., BTOS surveys all firms while the other surveys are concentrated on larger firms) and employment- vs firm-weighting.[20] Surveys also document heterogeneity in adoption, with larger and younger firms more likely to adopt AI. More technology-oriented sectors like information, financial activities and professional & business services have higher adoption rates (see figure 4, panel B). Indeed, surveys of software developers and related occupations suggest that adoption of

[19] Similarly, the Board of Governors' 2025 Survey of Household Economics and Decisionmaking (SHED) finds that about 20 percent of workers used generative AI in the previous month.

[20] See Crane et al. (2025) and Allen (2026) for reviews of recent surveys and discussion of the differences, and Bonney et al. (2026) for more on the Census BTOS specifically.

AI coding tools is very high, with more than half using the tools every day. The high adoption rates among coders and the rapid progress in coding tools suggest that software and coding may become an important channel for business AI adoption. For example, businesses may turn to their IT teams to develop new automated workflows or other innovations, rather than wait for AI firms to provide occupation- or industry-specific tools.

While adoption has grown significantly, Bonney et al. (2026) show that AI-adopting firms tend to only apply it in a couple of business areas at a time (e.g. IT and HR, or marketing and strategy), rather than across the entire organization. R&D, a business area where AI might provide a particularly significant boost to more fundamental innovation and productivity gains, has relatively low uptake. Business areas like IT, marketing, finance, and strategy seem to be the most common AI applications, with less use in HR and actual production (Bonney et al. 2026, Baslandze et al. 2026). More dramatic changes—such as using AI as part of their core production process, engaging in wholesale reorganization, or fully automating some tasks—might be necessary to sustain any productivity boost over the long term. The establishment of new AI-centered businesses could also be an important channel for longer-run productivity improvements. We consider the current and future economic effects in the next sections.

5. Effects on Consumer Prices So Far

The data center buildout has resulted in price increases for some consumer products that use similar hardware. In particular, prices for memory (DRAM), data storage equipment (flash drives), bare circuit boards, and electronic assembly (loaded circuit boards) have all moved up sharply in recent months. In addition, PCE prices for "video, audio, photographic, and informationA processing equipment and media" grew almost 8 percent over the 12 months through March, after falling 3 percent per year, on average, over the preceding three-year period. A notable contributor to inflation in this component is "computer software and accessories", which has seen surging prices since the end of last year (Barbarino et al. 2026)—perhaps as software companies pass on the costs of new AI features.[21] The surging price for data storage equipment also shows up in this component of PCE prices.

Up to now, these price increases have had a noticeable but limited effect on topline PCE prices, but the effect may continue to increase over time. Recent research has found that core PCE categories in the 90th percentile of data-center exposure have had around a 5 percentage points higher inflation (since January 2025) than the median exposed core PCE category (Leonard, Minton, and Orchard, 2026). Another possible effect of AI on

[21] Adjusting price indexes for quality changes is difficult; ideally the cost of new software features should not be passed into the price index. But if measurement issues prevent the statistical agency from fully accounting for quality changes, then this error may cause an increase in measured inflation.

prices is through the cost of electricity, which has received considerable public attention. There is some evidence of wholesale prices increasing, but it's relatively early for there to be any conclusive evidence on the ultimate size and timing of price effects.[22] As mentioned in the introduction, it is also possible that as AI shows through into higher productivity it could also have a more salutary effect on inflation, but it is too soon for such dynamics to be apparent.

6. Labor Market Effects

6.A. A framework for thinking about the labor market effects of AI

A useful and increasingly common approach to analyzing the effects of technological change on the labor market is the "task-based" framework, which assumes that output is a function of a set of tasks that can be performed by either capital or labor (Acemoglu and Autor 2011). Capital may be tangible (e.g., machinery) or intangible (e.g., software). At the task level, AI tools may either automate tasks (allowing them to be performed with capital rather than labor) or augment them (increasing output per unit of labor input). In the simplest framework, automated tasks would likely see labor demand decrease, implying lower wages and employment, while augmented tasks would likely see labor demand increase, implying higher wages and employment; though this is not guaranteed.

In practice, most jobs—or occupations—require the same worker to complete a bundle of different tasks. Demand for an occupation can therefore change in a number of ways depending on which tasks are automated or augmented. For example, if AI can outperform humans in most of an occupation's core tasks, it is likely to substitute for human labor in that occupation. However, if AI automates only a fraction of an occupation's tasks, it may complement human labor because workers can dedicate the time they save to their remaining tasks. In this case, AI could increase workers' productivity and, potentially, labor demand (Hampole et al. 2025).[23]

Beyond its effects on demand for specific tasks and occupations, AI may also affect demand for industries' output through higher productivity and lower output prices,

[22] Alvarez et al. (2026) found a statistically significant but quantitatively very small effect of data center investment on local electricity prices. An analysis from The Economist found that through 2024 there was no correlation across U.S. states between changes in data center capacity and changes in electricity prices. In contrast, Kay et al. (2026) and an analysis from Bloomberg found larger impacts on prices in wholesale markets in response to increased data center demand.

[23] How AI affects wages and employment for an occupation may also depend on the type of tasks that AI automates, specifically the degree of expertise they require to complete. Autor and Thompson (2025) point out that technologies that substitute for expert tasks are expected to lower wages but increase employment, while technologies that substitute for inexpert tasks are expected to raise wages and lower employment.

which in turn could increase labor demand in those industries.[24] Additionally, AI could also increase the pace of innovation, which could lead to labor demand for entirely new occupations, or an economy-wide change in labor demand. At this point, there is enormous uncertainty about the relative importance of these channels and how they may evolve over the next several years.

6.B. Measuring AI exposure

Empirical research on the labor market effects of AI operationalizes the task-based framework by identifying which work activities could be feasibly performed using AI applications, then aggregating these task-level exposure scores to occupations or industries.[25] These measures show that a substantial share of workers may see at least some of their tasks affected by AI automation or augmentation, though actual adoption is likely to lag behind theoretical exposure (Massenkoff and McCrory 2026).

One popular measure created by Eloundou et al. (2024), for example, uses an AI model (as well as human annotators) to assess whether access to LLMs or LLM-powered systems could reduce the time required for a human to complete a given task by at least 50 percent while preserving or improving quality. This exposure index implies that 46 percent of jobs could have at least half of their tasks affected by LLMs and complementary software in the near future. Handa et al. (2025) develop an alternative exposure index using data from millions of conversations with Claude. Their measure implies that 36 percent of occupations have at least a quarter of their tasks exposed to AI. Over half of the user interaction patterns they identify, however, are more consistent with task augmentation than automation.[26]

The literature generally agrees that more educated and higher-wage occupations are, on average, more exposed to AI, both in terms of potential automation and augmentation effects.[27] This contrasts with earlier waves of technological innovation such as industrial robotics and computer software which disproportionately automated "routine" tasks concentrated in the middle of the income distribution. That said, there is significant disagreement across metrics about which specific occupations are most highly

[24] The price elasticity of demand for the industry's output plays an important role in determining whether productivity gains raise or lower labor demand. If the elasticity of product demand is inelastic, for example, labor demand could decrease as productivity rises.

[25] See Frey and Osborne (2017); Brynjolfsson, Mitchell, and Rock (2018); Webb (2020); Felten, Raj, and Seamans (2023); Eloundou et al. (2024); Handa et al. (2025).

[26] Similarly, responses to the 2026 AI supplement to the Census Bureau's Business Trends and Outlook Survey (BTOS) show that 44 percent of firms that use AI reported augmenting worker tasks, while only 5 percent reported pure task substitution (Bonney et al. 2026).

[27] Manning and Aguirre (2026) note that AI exposure is positively correlated with workers' adaptive capacity. That is, workers in occupations highly exposed to AI also tend to have greater liquid financial resources, live in denser local labor markets, and have more transferable skills, on average, suggesting that they may have relatively strong means to manage an AI-induced labor market transition (provided that such disruptions occur gradually).

exposed to AI (Crane and Soto 2026). Different AI models, for example, make different assumptions about their own capabilities, and therefore produce different task exposure rankings (Yin et al. 2026). Further, AI exposure can also change over time as the underlying technology improves, meaning that tasks that are not currently exposed to AI may become exposed in the future. Advances in agentic AI or integration with robotics, for example, could affect tasks that are currently outside the technological frontier.

6.C. Empirical evidence on labor market effects

Recent empirical work suggests that AI may have already reduced demand for some highly exposed occupations. Crane and Soto (2026), for example, find that employment in computer programming-intensive occupations—though still increasing—has decelerated sharply since the introduction of LLMs in late 2022. Similarly, while job postings for software developers on Indeed have edged up since the middle of last year, they remain about 20 percent below their pre-pandemic level, consistent with weaker demand for these roles than before the introduction of LLMs. AI may also have contributed to recent employment declines in other highly exposed occupations such as customer service representatives and administrative assistants (Boesler 2026). However, factors unrelated to AI, including payback from overhiring during the post-pandemic tech boom and rising interest rates, may have contributed to these patterns as well.

Within occupations and firms, there is some evidence that reductions in labor demand have been concentrated among early-career workers. Using data from the payroll processor ADP, Brynjolfsson et al. (2025) find that workers between the ages of 22 and 25 in AI-exposed occupations experienced a 16 percent decline in employment following the introduction of LLMs relative to more experienced workers in the same occupation and workers of all ages in occupations less exposed to AI.[28] Hosseini and Lichtinger (2025) and Tucker (2026) reach similar conclusions using complementary datasets and research designs. Additionally, these studies find that AI exposure has had a larger effect on early-career employment than on wages, and that these adjustments are driven primarily by reduced hiring rather than an increase in separations. That said, the Job Openings and Labor Turnover Survey (JOLTS) shows that layoffs in the information industry (a sector highly exposed to AI) have picked up this year, even as economy-wide layoffs remain relatively low.

There is also some evidence that AI has led to labor reallocation within and across firms. A survey conducted by the Federal Reserve Bank of New York, for example, showed that firms were more likely to report retraining workers in occupations exposed to AI rather than laying them off, consistent with within-firm labor reallocation (Abel et al. 2025). In aggregate, the Yale Budget Lab found that since the introduction of LLMs

[28] Entry-level positions may be especially vulnerable to AI as these workers often perform relatively simple, codifiable tasks within the occupation. Experienced workers, by contrast, may have accumulated idiosyncratic knowledge or firm-specific expertise that is more difficult to substitute with AI.

in late 2022 the occupational composition of the labor market has been changing a bit faster than it did following the introduction of personal computers in the 1980s or the internet in the 1990s, though it is difficult to attribute this pattern entirely to generative AI (Gimbel et al. 2025).

All told, the literature suggests that AI is already influencing labor market outcomes for certain occupations and likely for early-career workers in particular. That said, the extent to which AI is showing up in aggregate employment and unemployment is less clear, as the groups most affected by AI at this point make up a relatively small share of total employment. This, however, could change as business adoption continues to grow and deepen. One survey, for example, found that few senior business executives reported any effect of AI on their firms' employment over the past three years, but, on average, U.S. executives anticipate that AI will lead to a 1.2 percent reduction in employment at their firms over the next three years (Yotzov et al. 2026).

7. Productivity Implications

There is increasing evidence that access to AI tools improves worker efficiency, speed, and accuracy at certain tasks.[29] These studies tend to find the largest effects for the least experienced workers, suggesting that access to AI may help inexperienced workers catch up to their peers (conditional on their continued employment). That said, they also focus on tasks such as coding, writing, and customer service, which are activities especially well-suited to generative AI applications. The extent to which these task-level productivity effects scale within firms, or across the economy, remains an important open question.

Although aggregate productivity growth has been relatively strong in the last few years, the contribution of generative AI thus far is likely small (Soto et al. 2026). Few executives surveyed by Yotzov et al. (2026), for example, report that AI has influenced their firms' productivity over the past 3 years. At the same time, there is suggestive evidence that certain publicly traded firms that are more likely to be AI users—and may be using AI more intensively than the typical U.S. firm—have experienced stronger growth in real revenue per employee compared to firms with a lower likelihood of AI use (Kadyrzhanova and Tuzun 2026). Going forward, the productivity effects of AI could broaden out as adoption continues to grow and deepen: Yotzov et al. (2026) report that U.S. executives expect that AI will increase their firms' productivity by 2.3 percent, on average, over the next 3 years.

The full effects of AI on aggregate productivity will likely take considerable time to materialize, as they depend not only on firms adopting the new technology, but also on

[29] See, for example, Brynjolfsson, Li, and Raymond (2025); Cui et al. (2024); Dell'Acqua et al. (2023); Peng et al. (2023); Noy and Zhang (2023).

reorganizing their production processes and investing in complementary assets, both tangible and intangible (such as training and management practices). If firms temporarily divert resources from current production toward these innovation efforts, then *measured* productivity may even decline in the early years of adoption. McElheran et al. (2025), for example, document J-shaped returns following the adoption of (pre-LLM) AI technologies among American manufacturing firms from 2012 to 2021. Other production bottlenecks, regulatory barriers, and institutions could also dampen the productivity effects of AI adoption in the near term.[30]

Over a longer horizon, there are a wide range of estimates for the potential effect of AI on aggregate productivity growth. These depend heavily on assumptions including how fast it will be adopted by firms, how much it will influence task-level productivity, and whether it will lead to an increase in the productivity of research and development. Acemoglu (2025), for example, estimates that AI will increase cumulative TFP growth by only 0.7 percent over the next decade. Using similar frameworks with more optimistic assumptions, however, Filippucci et al. (2025) estimate a TFP growth effect of 0.3 to 0.9 percentage point per year over the next decade and Aghion and Bunel (2024) arrive at a median estimate of 0.7 percentage point per year. Models that incorporate R&D effects and innovation spillovers in addition to task-level efficiency gains suggest that the productivity effects of AI could be even larger (Baily et al. 2023; Bontadini et al. 2025).

8. Conclusion

We are only a few years into the AI boom but have already seen significant impacts on some areas of the U.S. economy, if not yet on others. The direct capital investment in the data centers needed to train and operate AI models seems likely to boost GDP for several years, even under conservative assumptions. The surge in investment has driven up prices in a limited number of related sectors. Demand for new AI products has resulted in AI adoption growing rapidly both among consumers and firms, though most firms are currently not using AI very intensively or at all. Nevertheless, growing adoption has already pushed the revenues of the foundational AI companies to about 19 percent of the entire software industry's gross output. Effects on the labor market so far appear to have been small and somewhat mixed, but may be larger going forward. Similarly, productivity appears to have moved higher in recent years but that change mostly predates the diffusion of generative AI. Looking ahead, the macroeconomic effects of AI could become much more pronounced if, for instance, data

[30] Jones and Tonetti (2026), for example, show that even when most tasks can be automated, output growth is constrained by the remaining tasks that must be accomplished using human labor. Demirer, Musolff, and Yang (2026) provide empirical evidence consistent with this hypothesis in the context of software development: AI tools, including coding agents, significantly increase coding activity, but these gains translate only partially into the deployment of final software and its use.

center investment demand were to outstrip the capacity to supply needed materials and labor for a prolonged period of time, if AI starts to boost productivity more substantially and ultimately proves disinflationary, or if its impact on labor market dynamics were to broaden.

References

Abel, Jaison R., Richard Deitz, Natalia Emanuel, Benjamin Hyman, and Nick Montalbano. 2025. "Are Businesses Scaling Back Hiring Due to AI?" Federal Reserve Bank of New York: Liberty Street Economics. https://libertystreeteconomics.newyorkfed.org/2025/09/are-businesses-scaling-back-hiring-due-to-ai/.

Acemoglu, Daron. 2025. "The Simple Macroeconomics of AI," *Economic Policy*, 40 (121): 13-58.

Acemoglu, Daron, and David Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." In *Handbook of Labor Economics*, Vol. 4.

Aghion, Philippe and Simon Bunel. 2024. "AI and Growth: Where do We Stand?" Federal Reserve Bank of San Francisco. https://www.frbsf.org/wp-content/uploads/AI-and-Growth-Aghion-Bunel.pdf.

Allen, Jeffrey S. 2026. "Monitoring AI Adoption in the US Economy." *FEDS Notes*. Board of Governors of the Federal Reserve System, April 3. https://www.federalreserve.gov/econres/notes/feds-notes/monitoring-ai-adoption-in-the-u-s-economy-20260403.html.

Alvarez, Fernando E., David Argente, Joyce Chow, and Diana Van Patten. 2026. "Data Centers and Local Economies in the Age of AI: A Shift-Share Approach"**.** NBER Working Paper No. 35194. National Bureau of Economic Research. https://doi.org/10.3386/w35194

Autor, David H. and Neil Thompson. 2025 "Expertise," *Journal of the European Economic Association* 23(4): 1203-1271. https://doi.org/10.1093/jeea/jvaf023

Baily, Martin Neil, Erik Brynjolfsson, and Anton Korinek. 2023. "Machines of Mind: The Case for an AI-Powered Productivity Boom." Brookings Institution, https://www.brookings.edu/articles/machines-of-mind-the-case-for-an-ai-powered-productivity-boom/.

Baily, Martin Neil, David M. Byrne, Aidan T. Kane, and Paul E. Soto. 2025. "Generative AI at the Crossroads: Light Bulb, Dynamo, or Microscope?" *Finance and Economics Discussion Series* 2025-053. Washington: Board of Governors of the Federal Reserve System. https://doi.org/10.17016/FEDS.2025.053.

Barbarino, Alessandro, Anthony Diercks, and Stephen Miran. 2026. "Measurement of "Computer Software and Accessories" Inflation." *FEDS Notes*. May 22,2026. https://doi.org/10.17016/2380-7172.4089.

Baslandze, Salomé, Zachary Edwards, John R. Graham, Ty McClure, Michael Sparks, Brent Meyer, Sonya Ravindranath Waddell, and Daniel Weitz. 2026. "Artificial Intelligence, Productivity, and the Workforce: Evidence from Corporate Executives." *Federal Reserve Bank of Atlanta Working Paper* 2026-4, March. https://doi.org/10.29338/wp2026-04.

Bick, Alexander, Adam Blandin, David J. Deming, and Nathalie Gazzaneo. 2025. *Generative AI Adoption Tracker*. Harvard Project on Workforce. https://www.genaiadoptiontracker.com/.

Bick, Alexander, Adam Blandin, David J. Deming, Nicola Fuchs-Schündeln, and Jonas Jessen. 2026. "Mind the Gap: AI Adoption in Europe and the U.S." NBER Working Paper No. 34995. National Bureau of Economic Research. https://doi.org/10.3386/w34995.

Braun, Robin and Fariha Kamal. 2026. "When Protection Raises Costs: Evidence from the 2025 Tariffs on U.S. Manufacturing." mimeo.

Bonney, Kathryn, Cory Breaux, Emin Dinlersoz, Lucia Foster, John Haltiwanger, and Aditya Pande. 2026. "The Microstructure of AI Diffusion: Evidence from Firms, Business Functions, and Worker Tasks." *Center for Economic Studies Working Paper* CES-26-25, U.S. Census Bureau, April.

Bontadini, Filippo, Carol Corrado, Jonathan Haskel, and Cecilia Jona-Lasinio. 2025. "AI as an Innovation in the Method of Innovation: Implications for Productivity Growth in the US and Europe," https://spiral.imperial.ac.uk/server/api/core/bitstreams/aa4caa79-9ed3-4da2-b8f8-8281b022e249/content.

Bontadini, Filippo, Carol Corrado, Jonathan Haskel, and Cecilia Jona-Lasinio. 2026. "Is Software Eating the World? Measuring the Progress and Diffusion of AI." Working paper, National Bureau of Economic Research, Cambridge, MA. https://conference.nber.org/conf_papers/f234032.pdf.

Boesler, Matthew. 2026. "US is Starting to See Heavy Job Losses in Roles Exposed to AI." *Bloomberg*. May 15, 2026. https://www.bloomberg.com/news/articles/2026-05-15/us-is-starting-to-see-heavy-job-losses-in-roles-exposed-to-ai.

Brandsaas, Erik, Daniel Garcia, Robert Kurtzman, Joseph Nichols, and Adelia Zytek. 2025. "Estimating Aggregate Data Center Investment with Project-level Data." *FEDS Working Paper*. https://doi.org/10.17016/FEDS.2025.109.

Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock. 2018. "What can machines learn, and what does it mean for occupations and the economy?" AEA Papers and Proceedings, 108, 43–47. https://doi.org/10.1257/pandp.20181019

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen. 2025. "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence." https://digitaleconomy.stanford.edu/app/uploads/2025/12/CanariesintheCoalMine_Nov25.pdf

Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond. 2025. "Generative AI at work," *Quarterly Journal of Economics*, 140(2): 889-942.

Chatterji, Aaron, Thomas Cunningham, David J. Deming, Zoe Hitzig, Christopher Ong, Carl Yan Shan, and Kevin Wadman. 2025. "How People Use ChatGPT." NBER Working Paper No. 34255. National Bureau of Economic Research. https://doi.org/10.3386/w34255.

Clay Mathematics Institute. 2026. "Navier-Stokes Announcement." September 11, 2026.

Crane, Leland, Michael Green, and Paul Soto. 2025. "Measuring AI Uptake in the Workplace." *FEDS Notes*. Board of Governors of the Federal Reserve System, February 5. https://www.federalreserve.gov/econres/notes/feds-notes/measuring-ai-uptake-in-the-workplace-20240205.html.

Crane, Leland D. and Paul E. Soto. 2026. "AI and Coder Employment: Compiling the Evidence" *FEDS Notes*. Board of Governors of the Federal Reserve System, March 2026. https://doi.org/10.17016/FEDS.2026.018

Cui, Zheyuan, Mert Demirer, Sonia Jaffe, Leon Musolff, Sida Peng, and Tobias Salz. 2024. "The Effects of Generative AI on High Skilled Work: Evidence from Three Field Experiments with Software Developers." SSRN Working Paper. https://dx.doi.org/10.2139/ssrn.4945566.

Dell'Acqua, Fabrizio, Edward McFowland, III, Ethan Mollick, Katherine Kellogg Lifshitz-Assaf, Saran Rajendran, Lisa Krayer, Francois Candelon, and Karim Lakhani, 2023. "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality," Working Paper no. 24-013, Harvard Business School, Cambridge, MA.

Demirer, Mert, Leon Musolff and Liyuan Yang. 2026. "Writing Code vs. Shipping Code: Productivity Effects Across Generations of AI Coding Tools." NBER Working Paper No. 35275. National Bureau of Economic Research. https://doi.org/10.3386/w35275.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2024. "GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models," Science, 384(6702), 1306–1308. https://doi.org/10.1126/science.adj0998

Felten, Ed, Manav Raj, and Robert Seamans. 2023. "How will language modelers like ChatGPT affect occupations and industries?" arXiv preprint arXiv:2303.01157

Filippucci, Francesco, Peter Gal, and Matthias Schief. 2025. "Aggregate Productivity Gains from Artificial Intelligence: A Sectoral Perspective." SSRN Working Paper. https://dx.doi.org/10.2139/ssrn.5260318.

Gimbel, Martha, Molly Kinder, Joshua Kendall and Maddie Lee. 2025. "Evaluating the Impact of AI on the Labor Market: Current State of Affairs." The Budget Lab at Yale, https://budgetlab.yale.edu/research/evaluating-impact-ai-labor-market-current-state-affairs

Hampole, Menaka, Dimitris Papanikolaou, Lawrence D.W. Schmidt, and Bryan Seegmiller. 2025. "Artificial Intelligence and the Labor Market." NBER Working Paper No. 33509. National Bureau of Economic Research. https://doi.org/10.3386/w33509

Gundlach, Hans, Jayson Lynch, Matthias Mertens, and Neil Thompson. 2026. "The Price of Progress: Price Performance and the Future of AI." *arXiv*, https://doi.org/10.48550/arXiv.2511.23455.

Handa, Kunal, Alex Tamkin, Miles McCain, Saffron Huang, Esin Durmus, Sarah Heck, Jared Mueller, Jerry Hong, Stuart Ritchie, Tim Belonax et al. 2025. "Which Economic Tasks are Performed with AI? Evidence from Millions of Claude Conversations," arXiv preprint arXiv:2503.04761.

Hosseini, Seyed M. and Guy Lichtinger. 2025. "Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Resume and Job Posting Data." SSRN Working Paper. https://dx.doi.org/10.2139/ssrn.5425555.

Jones, Charles I. and Christopher Tonetti. 2026. "Past Automation and Future AI: How Weak Links Tame the Growth Explosion," https://web.stanford.edu/~chadj/JonesTonetti_Automation.pdf.

Jumper, John, Richard Evans, Alexander Pritzel, et al. 2021. "Highly Accurate Protein Structure Prediction with AlphaFold." *Nature* 596: 583–589. https://doi.org/10.1038/s41586-021-03819-2.

Kay, Owen and Robert Reaser and Reid Taylor. 2026. "Processing Power: The Effect of Data Centers on Wholesale Electricity Markets." Federal Reserve Bank of Dallas Working Paper. https://doi.org/10.24149/wp2606

Kemp, Andy. 2026. "Frequent Use of AI in the Workplace Continued to Rise in Q4." *Gallup Workplace*, January 25. https://www.gallup.com/workplace/701195/frequent-workplace-continued-rise.aspx.

Landefeld, J. Steven, and Barbara M. Fraumeni. 2001. "Measuring the New Economy." *Survey of Current Business* 81(3): 23–40.

Leonard, Emma, and Robbie Minton and Jacob Orchard, 2026, "The Inflationary Impacts of the Data-center Build-out" Working Paper.

Manning, Sam J. and Tomás Aguirre. 2026. "How Adaptable Are American Workers to AI-Induced Job Displacement?" NBER Working Paper No. 34705, https://doi.org/10.3386/w34705.

Massenkoff, Maxim and Peter McCrory. 2026. "Labor Market Impacts of AI: A New Measure and Early Evidence," https://www.anthropic.com/research/labor-market-impacts.

McElheran, Kristina, Mu-Jeung Yang, Zachary Kroff, and Erik Brynjolfsson. 2025**.** "The Rise of Industrial AI in America: Microfoundations of the Productivity J-curve(s)," U.S. Census Bureau Working Paper CES 25-27.

Mertens, Matthias, Adam Kuzee, Brittany S. Harris, Harry Lyu, Wensu Li, Jonathan Rosenfeld, Meiri Anto, Martin Fleming, and Neil Thompson. 2026. "Crashing Waves vs. Rising Tides: Preliminary Findings on AI Automation from Thousands of Worker Evaluations of Labor Market Tasks." *arXiv*, https://arxiv.org/abs/2604.01363.

METR. 2026. "Task-Completion Time Horizons of Frontier AI Models." Last updated May 8, 2026. https://metr.org/time-horizons/.

Van Nieuwerburgh, Stijn, "Financing the AI Buildout". forthcoming in the Brookings Papers of Economic Activity, Fall 2026, Available at SSRN: https://ssrn.com/abstract=6444363 or http://dx.doi.org/10.2139/ssrn.6444363

Noy, Shakked and Whitney Zhang. 2023. "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence," *Science* 381: 187-192.

OpenAI. 2026. "On the Navier–Stokes Millennium Prize Problem." September 8, 2026.

Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer. 2023. "The Impact of AI on Developer Productivity: Evidence from GitHub Copilot," *arXiv preprint arXiv:2302.06590*.

Soto, Paul E., Mason Thieu, and Jeffrey S. Allen 2026. "The AI Buildout and the Economy: Publicly Available Data to Assess AI's Impact," FEDS Notes. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/2380-7172.4119 .

Toner, Helen, Kendrea Beers, Steve Newman, Saif M. Khan, Colin Shea-Blymyer, Evelyn Yee, Ashwin Acharya, et al. 2026. *When AI Builds AI: Findings From a Workshop on Automation of AI R&D*. Center for Security and Emerging Technology. https://doi.org/10.51593/20250027.

Tucker, Lee C. 2026. "You're (not) Hired: Artificial Intelligence and Early Career Hiring in the Quarterly Workforce Indicators," Census Bureau Working Paper CES-26-27.

Webb, Michael. 2020. "The Impact of Artificial Intelligence on the Labor Market." https://www.michaelwebb.co/webb_ai.pdf.

Yin, Michelle, Hoa Vu, and Claudia Persico. 2026. "How (un)Stable Are LLM Occupational Exposure Scores? Evidence from Multi-Model Replication," NBER Working Paper 35110, https://doi.org/10.3386/w35110.

Yotzov, Ivan, Jose Maria Barrero, Nicholas Bloom, Philip Bunn, Steven J. Davis, Kevin M. Foster, Aaron Jalca, Brent H. Meyer, Paul Mizen, Michael A. Navarrete, Pawel Smietanka, Gregory Thwaites, and Ben Zhe Wang. 2026. "Firm Data on AI," NBER Working Paper 34836, https://doi.org/10.3386/w34836.