ECONOMIC INNOVATION GROUP

The American Worker Project

Understanding AI and Productivity

Chad Syverson

August 2026

At the microeconomic level, productivity is one of the best predictors of businesses' survival and growth prospects, the wages their workers earn, and the prices their customers pay. At the macro level, to quote Paul Krugman, "Productivity isn't everything, but in the long run, it is almost everything." (And I am not sure the proviso is necessary.)

Sustained long-run growth in income per capita requires productivity growth. It essentially sets the speed limit on how fast material standards of living can grow. Faster productivity growth means faster per capita income growth. Slower productivity growth means bad news.

But what about its relationship with AI? Well, it is still early, but a balanced assessment of the data indicates we do not yet know for sure. A few case studies have offered convincing evidence that AI has raised productivity in specific business settings. At a broader, economy-wide level, the combination of AI's novelty and the vagaries of measurement mean we cannot yet be certain if it is having an effect, or how big any such effect would be (though I discuss below what can be seen and what it might mean).

Regardless of what is in the data now, though, it is not hard to imagine AI and productivity could have a very important relationship in the future. The reason is simple: AI is plausibly the newest general purpose technology.[1] A general purpose technology has many applications across all sorts of production settings. Think the internal combustion engine, electricity, or the computer.

Past general purpose technologies spurred productivity accelerations that facilitated faster growth in economic well-being. The questions at hand are whether AI will do so too, how large and long-lasting that acceleration will be, and how the resulting productivity growth would be distributed through the economy and its people (for instance, how will its impact be spread across sectors, capital versus labor owners, and the income distribution).

What the Data Say Now

Let us take a look at current (U.S.) economy-wide productivity. Like most economies in the world, the U.S. had been experiencing a productivity growth slowdown since the middle-to-late 2000s. Labor productivity (output per worker-hour) grew throughout the 2010s by about 1.5 percent per year, roughly half of its growth rate during the productivity boom of 1995–2004.[2] Figure 1 shows labor productivity since 2013. As can be seen, the slow growth of the 2010s continued to the onset of the Covid pandemic. There was then a large, up-and-then-down blip in productivity during the height and eventual ebb of the pandemic. This blip was probably mostly a spurious measurement phenomenon due to composition and utilization effects rather than a true change in the economy's productive capacity. After the blip unwound, however, productivity from mid-2022 on has maintained a faster-than-2010s trajectory involving annual growth of about 2.2 percent.

Could this acceleration be due to AI? Perhaps. The timing leans against AI being the sole initial cause. The acceleration began when AI investments were still quite small relative to the size of the economy. Additionally, there were well-documented increases in economic dynamism (labor market churn and business formation) during the pandemic emergence whose timing matches the acceleration's start. On the other hand, total AI-related infrastructure investments grew very quickly and by 2025 were unmissable even in aggregate investment statistics. Business adoption of AI tools and AI-enabled processes have sped up as well. And of course even if other factors started the acceleration, it could well be that AI is starting to contribute now. Regardless of AI's current effect, the longer the aggregate productivity acceleration continues, the more plausible it is that AI is an important driver.

As for the magnitude of any effect, a sustained increase from 1.5 to 2.2 percent annual productivity growth would be substantial (after a decade, GDP per capita would be 7 percent higher than otherwise), but hardly unprecedented. The 1995–2004 productivity boom saw annual productivity growth of nearly 3.0 percent per year, and other past general-purpose-technology-related productivity boosts saw labor productivity growth in excess of 2.5 percent for a decade or longer.

Figure 1: U.S. Labor Productivity, 2013–26

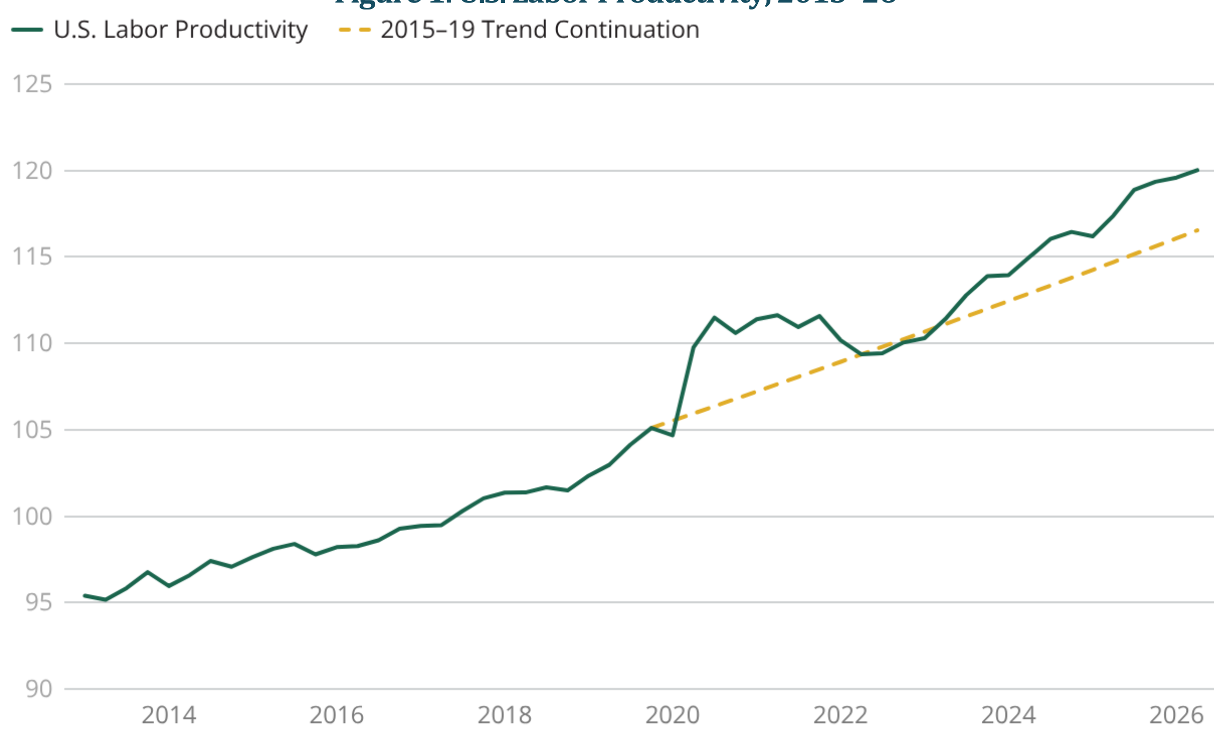

Source: U.S. Bureau of Labor Statistics; Author's Calculations.

Another approach to examining whether AI is currently affecting productivity is to look across sectors. For many reasons, the intensity of AI use already differs substantially across sectors. The highest adoption is in the information, financial services, and professional services sectors. At the other end of the spectrum are mining, construction, and accommodation and food services. If productivity growth has risen more in high-AI-adoption sectors than in low-adoption ones, this would be a sign of AI pushing productivity gains already.

Figure 2 shows this comparison. It plots the post-2020 change in a sector's contribution to economy-wide productivity growth (that is, the 2019–25 contribution minus the 2007–19 contribution) versus a measure of the sector's AI adoption rate. A sector's contribution to economy-wide productivity growth is the product of two numbers over the respective period: the sector's average annual productivity growth and its average share of total worker-hours in the economy. Essentially, this metric size-weights the sector's productivity growth, so between two sectors experiencing the same productivity growth rate, the larger sector increases economy-wide productivity growth by more. A contribution value of 0.1, for instance, implies that a sector's productivity growth raises economy-wide productivity growth by 0.1 percent. To measure a sector's AI adoption rate, I use the data from the Census BTOS survey administered in 2025.[3] The measured adoption rate reflects the fraction of companies (weighted by firm employment) that have adopted AI in their business practices. If the relationship between these two values is positive, the sectors that adopted AI more intensively by 2025 saw their productivity growth rise more.

The results, shown in Figure 2, are mixed. The correlation between the two values is positive, so higher-adopting sectors saw a larger pickup in size-weighted productivity growth. What is less clear is whether one can reject the possibility that this relationship is only due to statistical chance. Taking in all the data, standard statistical methods say the answer is no: the positive relationship is weak enough, relative to the number of sectors, that one cannot rule out chance as its source.

That said, this result is substantially driven by the fact that the retail sector saw a large productivity acceleration across the two periods even though its AI adoption rate is rather low. (That is why its dot is up and to the left on the figure.) We are treading on thin statistical ice if we simply ignore this observation and concentrate on the remaining sectors. However, to live dangerously for a moment, if we were to focus on all sectors other than retail, the positive correlation between a sector's productivity acceleration and its AI adoption intensity is more than twice as large as with retail included, and it is considerably less likely that this relationship would be due to mere chance alone.

Figure 2: Sector Productivity Growth Contributions vs. AI Intensity

Percentage-point change in each sector's contribution to labor productivity growth, 2019-25 vs. 2007-19

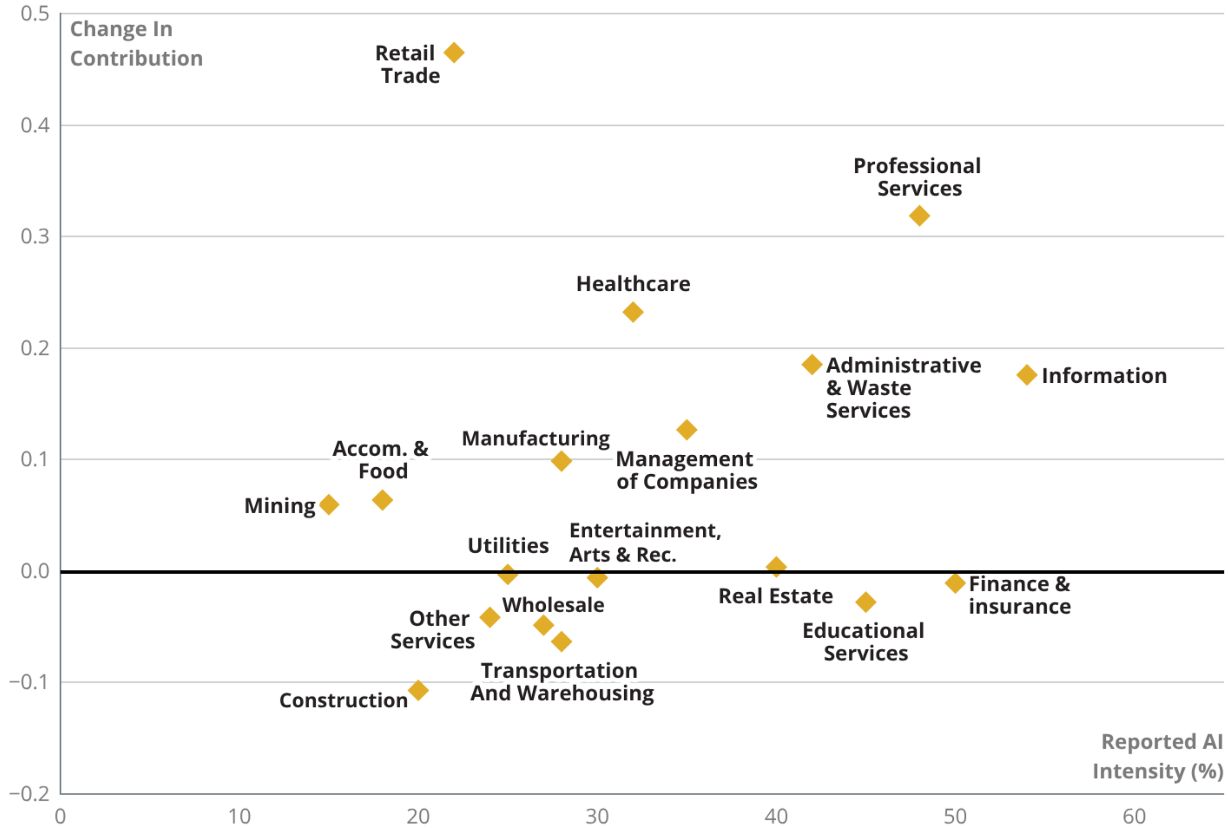

Source: U.S. Bureau of Economic Analysis; U.S. Bureau of Labor Statistics; U.S. Census Bureau, Business Trends and Outlook Survey; Author's Calculations.

Employment Effects

Any conversation about AI's productivity effects inevitably leads to speculation about its employment effects, often with a particular focus on AI-induced unemployment.[4]

Is unemployment a likely outcome if AI substantially boosts productivity growth? Well, recall that productivity is the ratio of outputs to inputs. If productivity rises, fewer inputs would be needed than before to make the same output. If that is the entire story, then yes, productivity growth would imply less employment, because fewer workers would be needed to make the same output.

But that is clearly not the entire story, for several reasons.

For one, that scenario assumes output is held fixed. This is highly unlikely to be the case after a productivity growth surge. Because productivity growth means fewer inputs are required per unit of output, the costs of making goods and services fall. Lower costs in turn mean lower prices, and lower prices mean people want more of the good (that is, demand curves slope down, as an economist would say).[5] That extra output demand requires more workers to produce, counteracting the employment effect from productivity gains requiring fewer workers per unit output. Does making this extra output need so many more workers that employment will actually rise when productivity goes up? Not necessarily; the net employment change depends on the relative size of the two opposing effects. However, it does mean that if productivity, say, doubles, employment will not fall by half. Additionally, there is a related effect if productivity growth also facilitates the creation of entirely new products (which, historically, it always has). Consumers will want these new products too, creating new demand for labor.

A second consideration is that products are made with multiple inputs. AI is one type of input and workers another (or another several types, as different types of workers can be thought of as different kinds of inputs). AI's employment effects will depend on how workers are used alongside it in the production process. If workers, or certain types of workers, are substitutes for AI, then AI adoption will reduce demand for them. On the other hand, if they are complements to AI, AI adoption will increase demand for them.

Third, AI's productivity effects are unlikely to act evenly across industries and sectors. Sectors with larger employment-reducing effects from productivity growth will shed more workers, but if demand for products made by sectors with slower productivity growth remains high (or perhaps is driven still higher as consumers shift spending share away from sectors with rapidly decreasing costs and prices), slower-productivity-growth sectors could absorb workers displaced from faster-productivity-growth sectors. This sort of cross-sectoral shift due to differential productivity growth has plenty of historical precedents. Advanced economies have seen enormous, broad-based shifts away from agriculture and toward manufacturing first, and later, away from manufacturing and toward services. There have been many other shifts operating at narrower sectoral levels.

Finally, any mass unemployment effects will depend on how fast productivity grows and how quickly the economy can reallocate resources in response to the productivity surge. Past GPT-driven productivity booms have raised annual productivity growth by 1 to 1.5 percent for a decade or two. The economy has been easily able to absorb any aggregate employment effects of these shifts. If AI were to raise productivity growth by 5 to 10 percent per year for a few years, on the other hand (a rate claimed by some boosters but that has not been observed before at an economy-wide level), that may outstrip the economy's flexibility and absorptive abilities.

So, do the early data offer any hints as to which combination of these multiple considerations about AI's employment effects may apply going forward? Early empirical studies of AI applications in production actually offer minimal evidence either way. At least over the first couple years of its use, AI has had remarkably little employment effect, positive or negative.[6] While that may seem on its face to be a non-informative result, the very fact that there have been no obvious employment effects yet does say something about the combined effects to this point.

Productivity J-Curve

One phenomenon at play in measuring AI's productivity effects is the productivity J-curve. The J-curve implies that when a new general purpose technology emerges, there will be a period — possibly of considerable length — during which resources are spent but with little incremental observed output. As a result, measured productivity is depressed relative to true productivity, and measured productivity growth may slow or even turn negative during the new technology's initial diffusion. Later, however, the mismeasurement reverses itself and there is a period where measured productivity growth is overstated.

Here is how that works. The undermeasured-then-overmeasured productivity pattern arises because fully harnessing general purpose technologies requires a lot of intangible investments. Companies cannot simply plug and play AI and experience the new technology's productivity benefits. To use AI effectively, companies must make many complementary investments like training workers, restructuring organizations, designing new business processes, and so on.[7] These investments require real resources to build. Conceptually, they are capital investments — that is, outputs of the production process that will be themselves valuable in the future for making more output.

These intangible investments are rarely measured as such, however. Instead, standard accounting systems treat them as invisible in terms of outputs while still counting the resources spent to produce them as business expenditures. For example, paying a team of employees to redesign an organization for AI efficacy builds intangible capital. Production of capital goods should be counted as output, and if those employees instead built a new office building, standard accounting would count that building as investment output. But intangibles like organizational capital here are not tracked in most accounting systems. As a result, output is undercounted, implying productivity (output per unit of inputs) is undercounted too.

That output undermeasurement causes the initial dip in the J-curve. However, it is just one of two measurement effects resulting from the failure to track intangibles in productivity measurement. The second arises because once installed, those intangible investments become part of the intangible capital stock. Like all capital, they are inputs used to make output (as a more effective organizational structure would, in the example above). However, being intangible, they are not measured as such. Once these intangible inputs are put to use, they become productive inputs. But because they remain unmeasured, the input denominator in the productivity ratio is understated. This leads to measured productivity being overstated, which creates the hump after the dip in the J-curve. We have understatement of productivity growth early in the diffusion of a general purpose technology and overstatement later.

This phenomenon in the context of AI raises the obvious questions of how large and how long the periods of productivity growth under- and over-measurement will be. Based on work I did with coauthors Erik Brynjolfsson and Daniel Rock, we found other technologies' J-curves could have understatement periods that extend well beyond a decade, and cumulative productivity mismeasurement of double-digits percent. As a result, it is worth remembering even persistently slow measured productivity growth in the face of broad adoption of AI is not an indication that AI isn't having (unmeasured) productivity effects now and that measured productivity will accelerate later in response. Where are we on the AI J-curve now? It is too early to know that with much precision (and it might depend on how dangerously you want to live when you interpret the results in Figure 2).

What Next?

As even this brief discussion has conveyed, AI is a technology with enormous potential to shape our economic futures. Economic theory points us in directions to look to try to predict something about those futures, and empirical work has given us a very early — if fuzzy — look at them. If AI adoption continues at the rapid rate it has seen over the past couple years, we will quickly learn more about how large AI's productivity effects are and how broadly they spread.

Chad Syverson is the George C. Tiao Distinguished Service Professor of Economics at the University of Chicago Booth School of Business, where he has been a faculty member since 2001. Syverson also serves as Deputy Director of the Becker Friedman Institute for Economics at the University of Chicago Booth.

Explore the Economic Innovation Group's American Worker Project here.

[1] Coincidentally (or perhaps prophetically?), the common abbreviation in economics for general purpose technology is GPT, the same as for generative pre-trained transformer, the underlying technology of chatbots.

[2] I focus the discussion on labor productivity because of its simplicity and familiarity. Single-input productivity measures like labor productivity can, however, be influenced by changes in input mix (like capital intensity) in ways that aren't related to the economy's overall ability to produce outputs. A more direct measure of the production frontier is total factor productivity (TFP), the amount of output produced per unit of combined inputs (labor, capital, and intermediates). That said, the labor productivity patterns discussed here hold for TFP as well.

[3] These AI adoption data are summarized in 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." CES Working Paper Number 26–25.

[4] AI industry leaders have contributed to this discussion in alarming ways. Anthropic CEO Dario Amodei predicted, "[AI] technology is not replacing a single job but acting as a 'general labor substitute for humans.'" OpenAI CEO Sam Altman once proclaimed that "a lot of jobs will go away" because of AI. Both have since tried to walk these predictions back.

[5] That prices fall when costs do is not an assumption that relies on highly competitive markets. Except in rather peculiar theoretical situations, companies with market power — even complete monopolies — want to charge lower prices when their production costs fall.

[6] See, e.g., Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond. 2025. "Generative AI at Work." Quarterly Journal of Economics 140(2): 889–942; Humlum, Anders and Emilie Vestergaard. 2026. "Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI." NBER Working Paper No. 33777; 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. 2026. "Firm Data on AI." NBER Working Paper No. 34836; Bonney, Kathryn, Cory L. Breaux, Emin Dinlersoz, Lucia S. Foster, John C. Haltiwanger, and Aditya A. Pande. 2026. "The Microstructure of AI Diffusion: Evidence from Firms, Business Functions, and Worker Tasks." NBER Working Paper No. 35141; Baslandze, Salome, Zachary Edwards, John R. Graham, Ty McClure, Michael Sparks, Brent Meyer, Sonya Ravindranath Waddell, and Daniel Weitz. Forthcoming. "Artificial Intelligence, Productivity, and the Workforce: Evidence from Corporate Executives." Journal of Finance: Insights and Perspectives.

[7] For a historical example of the importance of complementary intangibles and obtaining the full productivity benefits of a new general purpose technology (the electric motor), see David, Paul A. 1990. "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox." American Economic Review Papers and Proceedings 80(2): 355–61.