FEDERAL RESERVE BANK
PHILADELPHIA

MARCH 27, 2026

Monetary Policy and Productivity

2026 Macroeconomics and Monetary Policy Conference
San Francisco

Anna Paulson

President and Chief Executive Officer
Federal Reserve Bank of Philadelphia

The views expressed today are my own and not necessarily those of the Federal Reserve System or the Federal Open Market Committee (FOMC).

Remarks as prepared for delivery.

Monetary Policy and Productivity
2026 Macroeconomics and Monetary Policy Conference
March 27, 2026

Anna Paulson
President and CEO
Federal Reserve Bank of Philadelphia

Thank you for the invitation to join you today. It's great to be at the San Francisco Fed and to be part of this conversation.

I'm going to focus today on a question that is on my mind as a policymaker and that elicits strong and wide-ranging views. What might the rapid advance of artificial intelligence (AI) mean for the U.S. economy, and what does it mean for monetary policy?

Is AI more than hype? At the Children's Hospital of Philadelphia, AI systems are transforming pediatric care. AI tools help diagnose rare diseases and can predict patient decline hours before more traditional methods. They also allow for treatment protocols that are personalized in ways that would have been impossible just a few years ago.

But our relationship with technology is complicated. Some people embrace it, and others not so much. The Children's Hospital of Philadelphia is just a few miles away from the place where "hitchBOT" met an untimely end back in 2015. HitchBOT was a cute, hitchhiking robot designed to explore how people relate to technology. He had successfully hitchhiked across Canada and parts of Europe before being decapitated in downtown Philadelphia.

You don't have to look far to find passionate views on AI and the future of the economy. Some of my business contacts tell me AI will fundamentally transform their operations — changing not just how they produce, but what it is possible to produce. Others are more skeptical. They expect AI tools to produce incremental improvements rather than revolutionary changes.

These perspectives matter far beyond abstract economic debates. Right now — today — CEOs are making billion-dollar investment decisions, CFOs are developing hiring plans, and investors are setting valuations based on their assessment of AI's economic impact. And at the same time, young people are trying to factor in AI when they evaluate potential career paths. These decisions are being made based on a wide range of views about how AI might shape the future.

The questions I will explore today are: What does the range of alternative scenarios for AI and, hence, for productivity growth mean for monetary policy? Do policymakers and the private sector forecasters fully appreciate the range of possible outcomes?

Here's my roadmap. First, I'll lay out three plausible scenarios for productivity growth. Second, I'll discuss what each means for monetary policy, through the lens of the long-run federal funds rate. Then, I'll examine whether current projections and forecasts reflect the wide range of possibilities these scenarios imply. I'll end with some thoughts on how monetary policy should navigate the range of possibilities — especially since, unlike in the mid-1990s, we're navigating this with inflation above our 2-percent target. And, of course, these will be my views and not necessarily those of my Federal Open Market Committee (FOMC) colleagues.

Scenarios for Productivity Growth

Let's start with productivity. Nonfarm business productivity growth has been running at an average of 2.6 percent a year for the past couple of years. That is well above the 1.5 percent average we saw in the 2010s. This sizeable increase raises the question of whether we are seeing a structural shift in the economy's productive capacity — possibly with more to come given AI. But it's also possible that current gains won't last.

I'm going to lay out three productivity scenarios that are all plausible given what we know today.

Scenario One: Low — Scenario one is low-productivity growth. In this scenario, recent productivity gains mostly reflect post-pandemic adjustments — firms reallocating resources, cutting costs, and implementing efficiency measures in response to tight labor markets and high interest rates. Current productivity growth could be maintained for a while as these effects fade and firms transition to integrating AI. But eventually, all these adjustments run their course and productivity growth returns to the subdued 1- to 1.5-percent pace of the 2010s.

Scenario Two: Medium — Scenario two is the medium scenario. In this scenario, AI generates a significant productivity surge over the next three-to-five years or so as firms fully integrate new tools and automate tasks. After that though, structural productivity growth settles back to something like 2-and-a-quarter percent, which is close to the post-World War II average. The gains from AI are real and meaningful but front-loaded rather than permanent.

Scenario Three: High — Scenario three is the high scenario. In this scenario, AI is truly transformative. It lifts productivity growth above historical trends on a long-term basis. If this happens, the implications for living standards and potential output would be profound and lasting. In this scenario, long-run annual productivity growth would be noticeably higher than 2-and-a-quarter percent.

Which scenario is most likely? I don't know — and I'm not sure anyone does. AI certainly feels transformative. When I talk to businesses, I hear genuine enthusiasm. But I also hear about cautious implementation. The big productivity effects, if they come, seem likely to mostly lie ahead of us.

We are living through this in real time, and distinguishing between these scenarios with confidence will likely take years, and not months. The wide range of possibilities definitely plays a role in how I am thinking about monetary policy.

What Does This Mean for Monetary Policy?

Different scenarios about long-run productivity growth have implications for monetary policy through what they imply about the long-run federal funds rate —sometimes called the neutral or natural rate of interest, or, affectionately, R*.

Long-run R* is a little bit like a mythical beast — it's not visible to the human eye. When we are at full employment, inflation is at 2 percent and the economy is growing at potential — and no shocks are affecting either demand or supply, long-run R* is the short-term interest rate that is consistent with this unlikely alignment.

Achieving our dual mandate goals of price stability and maximum employment often means that, in real time, we want the federal funds rate to be either above, or below, its longer-run value. If growth is above potential and threatens to create inflation, the funds rate should be above R*. If a negative shock to demand threatens high unemployment, the funds rate should be below R* to help offset the shock. Long-run R* provides a reference point that helps in assessing whether monetary policy is restrictive, accommodative, or neutral.

There is a fundamental connection between R* and long-run productivity growth. When productivity growth is higher, there are more investment opportunities and a stronger demand for capital, and this puts upward pressure on interest rates. A world of sustainably higher productivity growth is, all else equal, a world with a higher R*. Other factors matter, too — demographics, global savings demand — but productivity growth is central.

So, let's consider the implications of the productivity growth rate scenarios I just discussed for long-run R*.

Under the low-productivity scenario, longer-run R* would be relatively low. For example, prior to the pandemic, the median FOMC participant had a projection for the longer-run federal funds rate of just 2.5 percent.

Under the medium-productivity scenario, long-run R* would be higher. Something like 3.1 percent, the median in the most recent Summary of Economic Projections (or SEP), might be reasonable.

Under the high-productivity scenario, long-run R* would be higher still. One plausible value would be the upper end of the SEP's central tendency for the long-run federal funds rate, which was 3.5 percent in March.

The three scenarios for productivity growth correspond to three different values of longer-run R* that differ by about one percentage point.

Do Market Participant and Policymaker Views Reflect the Range of Possible Outcomes?

The wide range of plausible outcomes for productivity growth should translate to a wide range of views for the long-run federal funds rate. I have already touched on the breadth of estimates for the long-run fed funds rate that show up in the Summary of Economic Projections that the FOMC produces four times a year. FOMC participants have boosted their estimate of the median long-run fed funds rate from about 2.5 percent prior to the pandemic to 3.1 percent in the most recent SEP. In addition, the dispersion in views among FOMC participants has doubled since before the pandemic.[1] This shift is consistent with the idea that, collectively, policymakers see a broader range of possibilities for productivity growth now compared to prior to the pandemic.

The private sector also produces estimates of the longer-run fed funds rate. For example, the Philadelphia Fed's Survey of Professional Forecasters asks survey participants for a forecast of the average return of the three-month Treasury Bill over the next 10 years, expressed as an annual percentage. This roughly corresponds to the longer-run federal funds rate.

In addition, the New York Fed's Survey of Market Expectations asks survey participants for their expectations of the longer-run fed funds rate. Like the SEP, the median estimates from both surveys have moved up in recent years and closely track the median in the SEP.

In contrast to FOMC participants, however, the dispersion in private sector views on long-run R* has either stayed the same or gotten smaller over the period beginning in 2012.[2] One possible interpretation of this contrast is that private sector experts see a much narrower range of possibilities for productivity growth compared to policymakers.

---

[1] The measure of dispersion in views for the FOMC comes from the SEP's central tendency for the longer-run federal funds rate. This is the range between the lowest and the highest projections for the longer-run federal funds rate after the three lowest and three highest values are dropped. This range has averaged 0.7 percentage points since 2023, double what it was prior to the pandemic. Interestingly, it has narrowed a bit very recently.

[2] Dispersion for the Survey of Professional Forecasters is measured as the range between the highest and lowest responses, after the top and bottom three have been dropped. This is the same measure that is used in the SEP. Dispersion in the Survey of Market Expectations is measured as the range between the high and the low values of the interquartile range.

Interestingly, professional forecasters are also asked directly about their estimates of productivity growth over the next 10 years. The range of responses to this question is much wider than the range of responses to the question about the longer-run federal funds rate.

Why This Disconnect Matters

What should we make of this potential disconnect? The broad range of views about R* on the FOMC seems to map reasonably well into plausible scenarios for AI and productivity growth. But we don't see the same pattern when we look at professional forecasters and market participants. Does this disconnect matter? In my view, it bears close watching because it could have implications for monetary policy transmission and how financial conditions respond to news.

There may be institutional reasons for market experts to produce similar forecasts for R*, but the apparent agreement among private-sector survey respondents is worth pausing on. It might suggest that there is a market consensus that productivity growth will stay near current levels over the longer term. This may well turn out to be right.

But the historical record on predicting productivity growth suggests caution. The path from invention to measured productivity impact is long, uneven, and poorly predicted in real time. We have repeatedly been wrong about both the timing and the magnitude of technology-driven gains. Given the wide range of possibilities that I see for productivity growth, I was surprised that there was so much apparent agreement among the market experts on long-run R*.

One thing that I took away from this investigation is that it is important for me to communicate clearly about the range of longer-term productivity growth scenarios that I see as plausible. My current estimate of longer-run R* is close to the SEP median of 3.1 percent. But, as I've emphasized, I could see that number being revised up, or possibly even down, as we get more data and learn more about how AI adoption and other forces are shaping longer-term productivity.

Implications for Monetary Policy

What does this all mean for monetary policy? In many ways, it is business as usual.

The goal of monetary policy is to deliver stable prices and maximum employment. This is true whether we end up in a low-, medium-, or high-productivity scenario, or in some other scenario altogether. And this is also true no matter what forecasters and the private sector believe about future productivity growth.

But how do we deliver stable prices and maximum employment when there are fundamentally different ways that the economy might evolve depending on how AI plays out? Here, risk management considerations intersect with our progress on getting inflation to 2 percent.

Suppose we see a surge in growth that could be productivity-driven — the kind of surge that, if it reflects genuine AI-enabled efficiency gains, need not lead to inflation. If that's the case, monetary policy does not need to adjust. But if it's not the case and policy stays on hold, inflationary pressures could build. And, of course, in real time, we won't know if the growth we are seeing is a reflection of productivity gains or overheating.

Among the things that I will be looking at to determine appropriate policy in a situation like that is: where we are, and where we have been, relative to our goals. Put simply, if inflation were at the 2-percent target, I would feel more comfortable being patient, keeping monetary policy on hold and waiting to see if a hypothetical growth surge puts upward pressure on inflation. But if inflation is above 2 percent and has been for some time, I would be more cautious. I would be inclined to weight the possibility of overheating more heavily in determining appropriate policy.

This brings me to the lessons that I take from the late 1990s. The economy experienced what appeared at first to be above-trend growth driven by a surge in information technology investment. Meeting after meeting, FOMC members anticipated they would need to raise rates, expecting inflationary pressures to emerge. But the inflation surge never came. The economy eventually delivered strong growth, falling unemployment, low inflation, and ultimately a lower federal funds rate.

Chairman Greenspan and his colleagues were navigating exactly the kind of risk management tradeoffs I've described. Was strong growth inflationary? Or were productivity gains boosting how fast the economy could grow without generating inflation? In this case, keeping the funds rate steady was the right call and their patience paid off.

But there's an important caveat: the FOMC had earned inflation-fighting credibility through the difficult Volcker years. The Greenspan Fed came into this period with inflation at 2 percent and with inflation expectations firmly anchored. That credibility gave the Committee room to be patient and to hold rates steady, even as growth ran above what conventional models suggested was sustainable. Markets and businesses trusted that if inflation did emerge, the Fed would act decisively.

We're in a different position today. Inflation has been above our 2-percent target for six years. While we've made significant progress, inflation is still 2.8 percent, not 2 percent. Long-term inflation expectations are consistent with 2 percent, but they may also be a little more fragile. And, of course, the conflict in the Middle East has created new risks to both inflation and growth.

The labor market is also very relevant for how I would evaluate a surge in growth that might, or might not, be driven by technological progress. If the unemployment rate is high, there is clearly more room to be patient. In a situation like today though, where we are closer to full employment, the calculus is trickier. But trends in labor market momentum and wage growth can provide valuable signals about the extent to which a

surge in growth is creating inflationary pressures. For example, I see no signs that the labor market is contributing to inflation today.

I have taken two lessons from the late 1990s. The first is that monetary policy credibility is valuable not just for fighting inflation but also for allowing economic growth to flourish. The second lesson I take from this episode is that whatever happens — the low-, medium-, high-productivity scenario, or something else altogether — we are better able to calibrate monetary policy when we start from a good place.

Conclusion

The question of whether AI will transform the economy — or turn out to be something more modest — is an important issue for everyone. The key priority for monetary policy, today, and in the wide range of scenarios we might see in the future, is to deliver on stable prices and maximum sustainable employment. Getting inflation to 2 percent and keeping the unemployment rate low creates benefits for individuals, businesses, and communities today, and creates credibility and flexibility that help us navigate the future.

Thank you.