BROOKINGS

SEPTEMBER 2026

ARTIFICIAL INTELLIGENCE AND THE ECONOMY

SUMMARY OF REMARKS

Benjamin F. Jones

Benjamin F. Jones is the Gordon and Llura Gund Family Professor of Entrepreneurship and Professor of Strategy at the Kellogg School of Management of Northwestern University and is affiliated with the NBER. These remarks were prepared for the Brookings–University of Utah Conference on “The Forces Shaping America’s Fiscal Future” in June 2026.

DISCLOSURES

The Brookings Institution is committed to quality, independence, and impact. We are supported by a diverse array of funders. In line with our values and policies, each Brookings publication represents the sole views of its author(s).

Much of the discussion about artificial intelligence (AI) swings between two poles. In one, AI automates most of what people do, growth accelerates without limit, and labor is left behind. In the other, AI is a useful tool that will nonetheless leave the broad contours of the economy much as they are. My remarks offer a simple framework that helps discipline these views and, in several places, pushes back on both of them. The framework is meant less as a forecast than as a way of organizing the questions we should be asking, and the empirical objects we should be measuring, before we take a strong stand on where AI leads.

Disciplining our approach to AI is also central to assessing America's fiscal future. Whether AI raises growth, and how the gains are divided between capital and labor and within labor itself, bears directly on the fiscal questions the other papers take up: the trajectory of revenues, the pressure on spending programs, and the distribution of income that any fiscal response must contend with. My aim is to help give those discussions an economic anchor.

I. A Task-Based Framework

The starting point is to think of any outcome we care about as the product of many underlying tasks. The outcome could be gross domestic product. It could equally be the output of research and development, where the "outcome" is a productivity gain, a new drug, or an increase in longevity. In either case, producing the outcome requires that a large set of distinct tasks get done, and each task can be performed either by a person or by a machine. By "machine" I mean AI, but also microscopes, tractors, and computers—any capital input that can take over work that labor once did.

Three features of this task structure govern what AI will do to the outcome. First is the share of tasks that machines can perform. Second is the productivity of machines at the tasks they perform. This includes productivity for both newly automated tasks and improvements at tasks machines already do. Third, the strength of bottlenecks, meaning how much the tasks that remain hard to automate constrain the whole. In the model, these map to three objects: the fraction of tasks automated, which I will call β; the average productivity of machines at their tasks, Z; and a substitution parameter, ρ, that determines how binding the bottlenecks are.

To make the role of bottlenecks precise, one can write the outcome, Y, as a generalized mean of the individual task outputs. Let y(j) denote the output of task j, with tasks indexed on the unit interval. The outcome is then:

$$Y = \left(\int_0^1 y(j)^\rho \, dj \right)^{\frac{1}{\rho}}$$

A single parameter, ρ, governs how the tasks combine. It corresponds to an elasticity of substitution across tasks through the relation σ = 1 / (1 − ρ). Four values help fix the intuition. At ρ = 1 the outcome is the arithmetic mean of task outputs. Tasks are perfect substitutes, there are no bottlenecks, and the best-performed tasks dominate, so excelling

at some tasks is enough to carry the whole. As ρ → 0 the outcome becomes the geometric mean.[1] At ρ = −1 the outcome is the harmonic mean, where the low-productivity tasks dominate the average and getting better at the tasks one already does well adds little. As ρ → −∞ the outcome collapses to the minimum of the task outputs, so the result is only as good as the single worst task.[2] Bottlenecks grow stronger as ρ falls, and the minimum function is their extreme.

My reading of the evidence is that the empirically relevant case is roughly ρ = −1, close to the harmonic mean, in both ordinary production and research. In that case the outcome takes the form:

$$Y = \left(\int_0^1 y(j)^{-1} \, dj \right)^{-1}$$

For example, take the economy as a whole. Most estimates of capital-labor substitutability imply an elasticity of substitution well below one, consistent with a ρ near −1. For research, studies spanning every field of science and invention point to a similar strength of complementarity across tasks. That the world sits near the strong-bottleneck end of the range rather than the substitutable end, does a great deal of work: It is what makes the results that follow, on growth, on the labor share, and on the limits to AI, come out as they do.

The value of this framework is that it forces three concrete questions for any setting, from cancer therapeutics to software to the economy as a whole: What share of tasks can AI take over? How good is AI at those tasks? And how strong are the bottlenecks? One cannot say much about AI's consequences without taking a stand on all three.

II. Automation Raises Output by Leveraging Scarce Labor

The first result is the most intuitive. Output per person rises as a greater share of tasks is automated, whether in the production of goods and services or in the production of ideas. The mechanism is worth stating plainly, because it also anticipates the results on labor that follow. When machines take over more tasks, labor concentrates on a smaller set of remaining tasks, and more output can be produced for each of those tasks. Automation, in this sense, leverages scarce labor. It lets a fixed number of workers stand behind a much larger volume of production.

III. Two Dimensions of Progress, and a Surprise About the Labor Share

The popular fear is that automation lowers labor's share of income, and in part it does. As machines perform a larger share of tasks (that is, β goes up) more of what is produced is

[1] To economists, this is the familiar Cobb-Douglas case.
[2] This is the Leontief or fixed-proportion production function in economics.

paid out to machines, and labor's share falls. If β were the whole story, rising automation would steadily erode labor's position.

But β is not the whole story. Improving the productivity of machines, raising Z, works in the opposite direction. When machines get better at the tasks they already do, the price of what they produce falls. With an elasticity of substitution below 1 (consistent with ρ = −1), prices fall faster than the quantity rises, so the expenditure share flowing to machines shrinks. In short, better machines are paid less as a share of the whole, and labor's share rises. The capital share can be written compactly as $\frac{\beta}{Z}$. Automation pushes it up, but machine productivity pushes it down.

This is the piece missing from automation models that track β but underweight Z. It can help explain a long-standing puzzle, namely why the labor share was roughly stable for so long even as machines advanced. Automation and machine improvement were pushing the labor share in opposite directions and roughly offsetting. It also carries an optimistic implication that cuts against the standard worry. AI productivity gains, even when fully embodied in machines, can raise labor's share of income rather than lower it. Doing automated tasks better, rather than simply automating more of them, is central to sustaining labor's share.

IV. Bottlenecks Are More Powerful than Infinities

If there is one idea I would want to highlight, it is that bottlenecks are extraordinarily powerful, more powerful than intuition suggests. Consider this: Over the past 50 years, computing power per dollar has risen by a factor of roughly 100,000,000,000,000,000. Total factor productivity over the same period has risen by a factor of about 3 (Nordhaus 2001, Jones and Liu 2024). An almost unimaginable improvement in an important input produced a modest improvement in the aggregate. The gap is the work of bottlenecks.

A simple example makes the point clear. Suppose every task produces output of 1. Then imagine that half of all tasks become infinitely productive while the other half are unchanged. If the tasks combine as a harmonic mean, which is close to what the evidence suggests, output does not explode. It merely doubles. Making half of everything infinitely good doubles the outcome, because the tasks you did not improve now dominate the average. Amazing gains at some things coexist with little movement in the aggregate, precisely because the aggregate is governed by what did not improve.

This is Baumol's cost disease (Baumol 1967), seen through the task lens. The activities we cannot easily automate, including much of healthcare, housing, and many in-person services, become expensive relative to everything else. The things we automate become cheap, but these other tasks remain hard, and therefore expensive. They become a larger share of the economy and a larger share of what constrains its growth. The same logic applies to research: A tool that solves step one of a 10-step discovery process does not deliver the cure if the other steps remain slow.

Bottlenecks therefore cut against the most dramatic techno-optimism. Powerful AI, meaning many highly capable systems operating on a limited set of tasks, does not by itself produce transformative growth. What matters far more than radical improvement in a narrow set of tasks is broadening the share of tasks that AI can perform at all. The marginal return to intelligence, to borrow a phrase (Amodei 2024, Jones 2026), is strongly limited when that intelligence operates on only a minority of tasks.

There is a mirror image worth naming, because it connects the two halves of the argument. Bottlenecks make growth harder, but they are also what protect labor. The hard-to-automate tasks are exactly where labor earns its rents. Baumol's cost disease is not only a drag on measured productivity; it is the mechanism by which labor holds its share as machines get better at everything else.

V. Skill Bias and Inequality Within Labor

The framework so far concerns the split between capital and labor. But the dominant driver of measured inequality over the past several decades has been inequality within labor, the gap between higher-skill and lower-skill workers. The natural additional question is then whether AI complements or substitutes for higher-skill labor. Two forces in the framework suggest that automation has historically been biased against lower-skill work, and a third suggests AI may not follow that pattern.

The first force is technological opportunity. A machine displaces labor at a task when it can produce the output more cheaply. This is more likely where machine productivity is high relative to human productivity, as with narrow, repetitive, or highly standardized activities. Machines often win on input prices as well. Automation thus tends to arrive first historically where labor is engaged in more routine, lower-skill tasks.

The second force is market size. The bottleneck logic implies that the largest markets, measured by expenditure, are precisely the tasks where humans are relatively unproductive and costs are high. From the perspective of directed technical change, those large and expensive markets are the most attractive to automate. One provocative way to put it is that AI may be an endogenous response to Baumol's cost disease. That is, innovation is drawn toward the very tasks that have become expensive because they resisted earlier automation.

Together these forces suggest AI could continue the long pattern of skill bias, replacing lower-skill work first. But there is a reason AI may be different. Earlier waves of automation acted on physical and routine tasks. AI is cognitive. It attacks complex analytical work that higher-skill workers do. Further, even as robotics improves, the material cost of physical machines will not necessarily be cheap, so that human physical labor may still be in substantial demand. Ultimately, the tasks most insulated from AI may turn out to be manual, physical, and face-to-face. As the discussion at the conference put it memorably: One can automate a coder but not a caulker, and the last task to be automated in a restaurant may be the person who brings your plate. Whether AI complements or substitutes for higher-

skill labor is, in my view, genuinely unresolved, and it is among the most important open questions for the distributional and fiscal analysis that follows.

VI. AI and the Acceleration of Ideas

AI matters for the economy not only through the ordinary production of goods and services but through the production of ideas. Research and development is itself a set of tasks, and the same three features govern how much AI can accelerate discovery.

Here a distinctive feature of human researchers becomes relevant. As the frontier of knowledge expands, any individual can master only a shrinking share of it. This is the burden of knowledge mechanism, where researchers grow narrower over time, not because people have become less capable but because there is so much more to know. Recent evidence on the "pivot penalty" (Hill et al. 2025) suggests that a given researcher is increasingly confined to a narrow area where that researcher has deep expertise. Human search is local; it clusters around what a given researcher already knows.

AI enters this picture with distinctive research capabilities. It can draw on the full breadth of recorded knowledge rather than the narrow slice any one person commands, which lets it widen the search and point toward more promising directions. In the creative-search framing of explore versus exploit, AI can appear more creative than human experts precisely because it is not confined to a local neighborhood of expertise. AI also, because it operates in fast and low-cost compute environments, has the potential to complete many more research tasks per hour or per dollar. Tools like AlphaFold in structural biology and GNoME in materials discovery are early evidence that AI can broaden search at low cost, massively scaling the avenues of discovery.[3]

But the bottleneck logic applies here with equal force. AI may transform one step of a research pipeline while leaving the others as slow as before, and it is the slow steps that ultimately set the pace. For example, as impressive as AlphaFold is, there has been little effect on drug discovery so far, and nor has GNoME obviously accelerated the generation of new materials. The challenge in both domains is that many other steps, including difficult experimental steps, are rate-limiting for achieving breakthroughs.

Evidence on the strength of bottlenecks points in a consistent direction. In the aggregate economy, most estimates of capital-labor substitutability imply strong bottlenecks. In research, work examining every field of science and invention points to a similar place. In both the production of goods and the production of ideas, bottlenecks appear strong. This is a critical fact for calibrating expectations.

[3] AlphaFold determines the 3D shape of proteins, which is essential to understanding their function. Its creators shared the 2024 Nobel Prize in Chemistry. GNoME suggests novel materials that could be useful in applications from batteries to computer chips. It has produced an approximate 10-fold increase in the number of known, potentially stable materials.

VII. What Would Break the Bottleneck, and Why It Matters Here

The bottleneck result relies on an important qualification, and the more dramatic scenarios emerge when that qualification no longer holds. The reason bottlenecks cap AI's aggregate effect is that some tasks remain in human hands. If AI came to perform essentially all tasks, in ordinary production, in research, or in both, the constraint would dissolve and very large growth accelerations would become possible. The emergence of more autonomous, agentic AI, systems that pursue broad objectives and coordinate with one another with little human mediation, is the development most likely to push the share of automated tasks toward that limit, and it would represent a different regime than the one the framework describes.

That qualification also frames the scarier distributional outcomes. The framework is a general-equilibrium one in which labor always has bottlenecked tasks to do, so it does not by itself generate mass unemployment; markets clear through wages and labor reallocation shifts workers to other labor tasks when certain tasks are automated. More difficult scenarios emerge if the speed of adjustment is slow compared with the pace of technological change. For labor, displacement could outrun labor's ability to reallocate, which would be especially hard for workers whose entire job is automated but have few outside options in the near term. A related question, which bears directly on fiscal policy, is who captures the returns to capital if the capital share does rise. That depends not only on the share itself but on how concentrated ownership and market power turn out to be. To the extent that AI market value becomes concentrated among a small set of capital owners, it will be for political economy mechanisms, as opposed to market mechanisms, to confront capital-labor inequality. Recent dynamics may be somewhat reassuring, especially as token prices are falling, open-source models sit close to the frontier, and many competing firms are developing data-center capacity. How that competitive picture evolves is an open and consequential question.

VIII. Summary

The implications of AI for the economy, whether for goods and services or for ideas, depend on three things: the share of tasks AI performs, how good AI is at those tasks, and the strength of bottlenecks. Bottlenecks are powerful, more powerful than infinities, and the evidence in both ordinary production and research points to strong bottlenecks. This tempers the most transformative forecasts. At the same time, to the extent that AI both automates a substantial share of tasks and performs them very well, it could deliver an unusually large gain in income per person, and doing those tasks very well is exactly what sustains a high labor share. The harder questions, inequality within labor, the concentration of capital returns, and the conditions under which rapid displacement outruns adjustment, are where the framework hands off to the fiscal and distributional analysis that is the heart of this conference.

References

Amodei, Dario (2024). "Machines of Loving Grace: How AI Could Transform the World for the Better." Personal Essay, October 2024.

Baumol, William J. (1967). "Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis." American Economic Review, 57(3), 415–426.

Hill, Ryan, Yian Yin, Carolyn Stein, Xizhao Wang, Dashun Wang, and Benjamin F. Jones (2025). "The Pivot Penalty in Research." Nature, 642, 999–1006.

Jones, Benjamin F., and Xiaojie Liu (2024). "A Framework for Economic Growth with Capital-Embodied Technical Change." American Economic Review, 114(5), 1448–1487.

Jumper, John, Richard Evans, Alexander Pritzel, et al. (2021). "Highly Accurate Protein Structure Prediction with AlphaFold." Nature, 596, 583–589.

Merchant, Amil, Simon Batzner, Samuel S. Schoenholz, et al. (2023). "Scaling Deep Learning for Materials Discovery." Nature, 624, 80–85.

Nordhaus, William D. (2007). "Two Centuries of Productivity Growth in Computing." Journal of Economic History, 67(1), 128–159.

BROOKINGS

1775 Massachusetts Ave NW
Washington, DC 20036
(202) 797-6000
www.brookings.edu