

## AI Productivity Growth Won't Match the Computer Revolution

*di Carl Benedikt Frey*

*Those promoting AI would be lucky to see the technology's impact on output per hour match even the short-lived burst of the 1990s and 2000s. Productivity growth will underwhelm, not because the technology is weak, but because it creates a bottleneck that earlier digital tools largely avoided.*

OXFORD—“Something Big Is Happening,” wrote AI startup founder Matt Shumer in a recent [viral essay](#) that captured his industry's swelling confidence that the technology will power the next great productivity boom. So far, the economy has not played along. In fact, since slowing sharply in the 1970s, US productivity has experienced only one brief burst of growth: the computer age. Output per hour [surged](#) by roughly 3% per year in the late 1990s and early 2000s, and then it petered out.

Could AI be different? Optimists point to headline labor productivity, which [grew](#) at a 1.8% annualized rate in the fourth quarter of 2025. But a cleaner measure by the US Federal Reserve Bank of San Francisco, which strips out cyclical intensity (the effect of simply running people and machines harder), shows that labor productivity [grew](#) just 0.2% year on year. That is hardly suggestive of “something big.”

On the contrary, we would be fortunate to see the technology match even the short-lived computer revolution. Productivity growth will likely underwhelm, not because the technology is weak, but because it automates something fundamentally different from what the personal computer and the internet did. More to the point, AI creates a bottleneck that earlier digital tools largely avoided.

Consider what the computer revolution actually automated—faster calculation and access to knowledge. PCs, email, spreadsheets, and the web removed friction from the

process of finding, storing, and transmitting information. A researcher who needed a source no longer had to search in a library or wait for it to arrive by mail. The productivity gains were relatively straightforward because humans could simply substitute the faster method (Google) for the slower one (a library). Information found online was the same as what you would have found on a shelf.

Crucially, when computers did perform core work, they did it deterministically. A spreadsheet could propagate bad inputs, but it did not invent arithmetic. Search engines could surface irrelevant material, but they did not fabricate sources. The principal risk was human error, not persuasive invention.

AI automates something different: the production of cognitive outputs themselves—from writing to coding. It often performs these tasks quite well. But because it can also be confidently wrong in ways that look plausible, it creates a tension that those navigating the computer revolution never faced: if humans need to remain in the loop to verify AI outputs, they will still need the domain knowledge that AI is supposedly substituting for. Ensuring reliability still requires scarce expertise and time. Thus, some of the time saved in generation is partly—and sometimes entirely—offset by the time spent reconstructing the reasoning, testing the claims, and taking responsibility for the result.

A Manhattan bankruptcy court provided the latest illustration of this problem just this month. Sullivan & Cromwell—one of Wall Street’s most prestigious firms—filed an emergency motion riddled with [fabricated citations](#) and other AI-generated errors. The mistakes were caught not by the firm’s own review process but by opposing counsel. The episode was absurd, but also diagnostic. It showed what happens when a tool that produces fluent output meets a world that demands verifiable truth.

The deeper issue is not merely that AI can be wrong. It is that the cost of errors is changing. As systems become more agentic—as they act autonomously, rather than just generating text or code in response to discrete prompts—mistakes become more consequential. A chatbot that hallucinates a paragraph is annoying. An agent that changes code, moves money, files paperwork, deletes a database, or triggers actions across systems can create real damage at machine speed.

Call it the verification tax. In any setting where someone is accountable for an outcome—law, medicine, regulated finance, engineering, or public policy—an AI output is not a finished product. It is a draft that must be checked. The work does not disappear; it shifts from producing to supervising. Net productivity becomes time saved generating a draft minus time spent ensuring its trustworthiness.

Hence, in a large field [study](#) of customer support, a generative AI assistant increased productivity by about 14% on average, with much larger gains for novices and little benefit for the most experienced workers. Because the tasks were standardized, the outputs were easier to evaluate, and the tool could distribute best practices quickly.

But when the context is more complicated, and correctness is harder to observe, the verification burden can overwhelm the benefit. A randomized [trial](#) of experienced open-source developers working on their own repositories found that access to frontier AI tools made them about 19% slower—largely because their time went into prompting, waiting, reviewing, and correcting.

These results imply that AI's payoff depends on task structure. Where errors are cheap and outputs are easy to test, AI can accelerate work. Where mistakes are costly, and correctness is hard to observe, the bottleneck shifts from “doing the work” to “certifying” it. The machine can produce endless output, but the organization cannot absorb endless verification. As the economists Christian Catalini, Xiang Hui, and Jane Wu [argue](#), when AI pushes the cost of execution toward zero, the binding constraint becomes human verification bandwidth—our limited capacity to validate outcomes and underwrite responsibility.

This framing also clarifies a longer-run risk. If firms respond to AI by hiring fewer junior lawyers and analysts, training less, and assuming the machine will handle the first draft, they erode the very expertise needed to check the machine's output. The organization will look leaner until the hidden error surfaces in public.

What, then, would it take for AI to deliver broad productivity gains, rather than a lot of activity and a pile of unpriced risk? The answer is verification infrastructure. For example, a federal judge in Texas now requires lawyers to certify that any AI-drafted language has been verified using traditional legal research.

A similar shift is needed across white-collar work. If companies want AI agents to change code, move money, and file paperwork, they will need provenance for claims, audit trails, and clear standards of due diligence. Such institutional change does not happen at the speed of model releases. Until regulations, compliance departments, professional norms, insurance, and courts catch up, AI's potential will remain limited.