

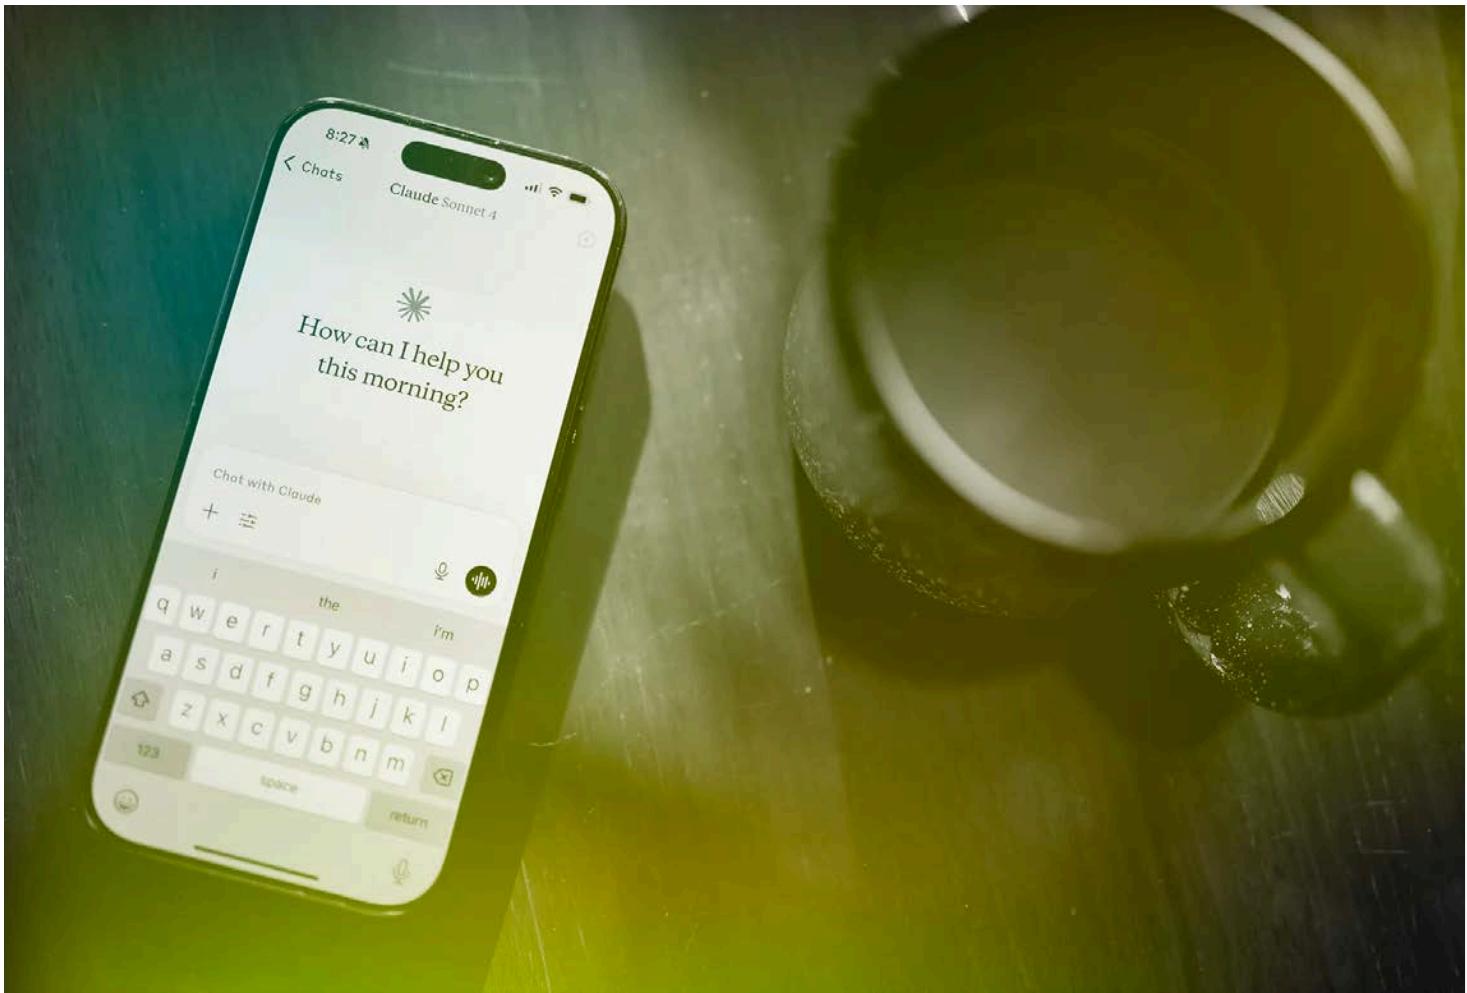
THE FINANCIAL PAGE

THE A.I.-PROFITS DROUGHT AND THE LESSONS OF HISTORY

Like the steam engine, electricity, and computers, generative artificial intelligence could take longer than expected to transform the economy.

By John Cassidy

August 25, 2025



Source photograph by Gabby Jones / Bloomberg / Getty

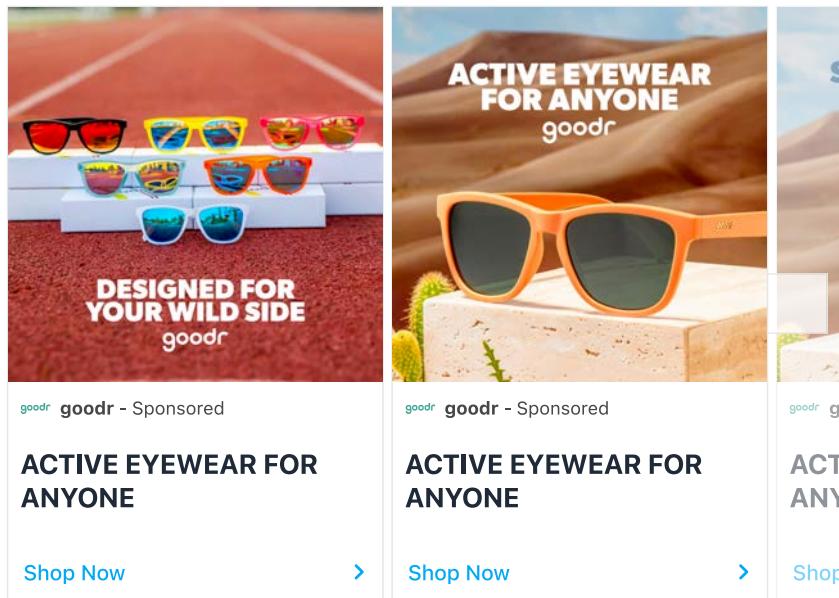
 [Save this story](#)

In a 1987 article in the *Times Book Review*, Robert Solow, a Nobel-winning economist at M.I.T., commented, “You can see the computer age everywhere but in the productivity statistics.” Despite massive increases in computing power and the rising popularity of personal computers, government figures showed that over-all output per worker, a key determinant of wages and living standards, had stagnated for more than a decade. The “productivity paradox,” as it came to be known, persisted into the nineteen-nineties and beyond, generating a huge and inconclusive body of literature. Some economists blamed mismanagement of the new technology; others argued that computers paled in economic importance compared to older inventions such as the steam engine and electricity; still others blamed measurement errors in the data and argued that once these were corrected the paradox disappeared.

Nearly forty years after Solow’s article, and almost three years since OpenAI released its ChatGPT chatbot, we may be facing a new economic paradox, this one involving generative artificial intelligence. According to a recent survey carried out by economists at Stanford, Clemson, and the World Bank, in June and July of this year, almost half of all workers—45.6 per cent, to be precise—were using A.I. tools. And yet, a new study, from a team of researchers associated with M.I.T.’s Media Lab, reports, “Despite \$30 - \$40 billion in enterprise

investment into GenAI, this report uncovers a surprising result in that 95% of organizations are getting zero return.”

The study’s authors examined more than three hundred public A.I. initiatives and announcements, and interviewed more than fifty company executives. They defined a successful A.I. investment as one that had been deployed beyond the pilot phase and had generated some measurable financial return or marked gain in productivity after six months. “Just 5% integrated AI pilots are extracting millions in value, while the vast majority remain stuck with no measurable P&L”—profit-and-loss—“impact,” they wrote.



The survey interviews elicited a range of responses, some of which were highly skeptical. “The hype on LinkedIn says everything has changed, but in our operations, nothing fundamental has shifted,” the chief operating officer at a midsize manufacturing firm told researchers. “We’re processing some contracts faster, but that’s all that has changed.” Another respondent commented, “We’ve seen dozens of demos this year. Maybe one or two are genuinely useful. The rest are wrappers or science projects.”

To be sure, the report points out that some firms have made successful A.I. investments. For example, it highlights efficiencies created by customized tools aimed at back-office operations, noting, “These early results suggest that learning-capable systems, when targeted at specific processes, can deliver real value, even without major organizational restructure.” The survey also cites some firms reporting “improved customer retention and sales conversion through automated outreach and intelligent follow-up systems,” which suggests that A.I. systems could be useful for marketing.

But the idea that many companies are struggling to achieve substantial returns jibed with another recent survey, by Akkodis, a multinational consulting firm. After contacting more than two thousand business executives, the firm found that the percentage of C.E.O.s who are “very confident” in their firm’s A.I.-implementation strategies has fallen from eighty-two per cent in 2024 to forty-nine per cent this year. Confidence had also fallen among corporate chief technology officers, although not by as much. These

developments “may reflect disappointing outcomes from previous attempts at digital or AI initiatives, delays or failures in implementation as well as concerns around scalability,” the Akkodis survey said.

Last week, media accounts of the M.I.T. Media Lab study coincided with a fall in highly valued stocks associated with A.I., including Nvidia, Meta, and Palantir. Correlation isn’t causation, of course, and recent comments by Sam Altman, the chief executive of OpenAI, may have played a bigger role in the sell-off, which was surely inevitable at some point, given recent price increases. At a dinner with reporters, Altman said valuations were “insane” and used the term “bubble” three times in fifteen seconds, CNBC reported.

Still, the M.I.T. study garnered a lot of attention, and after the initial raft of news stories about the research, a report emerged that the Media Lab, which has ties to many technology companies, was quietly restricting access to it. Messages that I left with the organization’s communications office and two of the report’s authors went unreturned.

Although the report is more nuanced than some news coverage made out, it certainly raises questions about the grand economic narrative that has underpinned the tech boom since November, 2022, when OpenAI released ChatGPT. The short version of this narrative is that the economy-wide diffusion of generative A.I. would be bad for workers, particularly knowledge workers, but great for companies, and their shareholders, because it would generate a big leap in productivity and, by extension, profits.

One possible reason this doesn’t seem to have happened yet recalls the suggestion that management failures were constraining the productivity benefits of computers in the nineteen-eighties and early nineties. The Media Lab study found that some of the most successful A.I. investments were made by startups that use highly customized tools in narrow areas of workflow processes. On the other side of the “GenAI Divide,” the study pointed to less successful startups that were “either building generic tools or trying to develop capabilities internally.” More generally, the report said the divisions between success and failure “does not seem to be driven by model quality or regulation, but seems to be determined by approach.”

Conceivably, the novelty and complexity of generative A.I. may be holding some companies back. A recent study, by the consultancy firm Gartner, found that fewer than half of C.E.O.s are confident that their chief information officers are “AI-savvy.” But there is another possible explanation for the disappointing record highlighted in the Media Lab report: for many established businesses, generative A.I., at least in its current incarnation, simply isn’t all it’s been cracked up to be. “It’s excellent for brainstorming and first drafts, but it doesn’t retain knowledge of client preferences or learn from previous edits,” one respondent to the Media Lab survey said. “It repeats the same mistakes and requires extensive context input for each session. For high-stakes work, I need a system that accumulates knowledge and improves over time.”

Of course, there are plenty of people who find A.I. useful, and there is academic evidence to back this up: in 2023, two economists at M.I.T. found that exposure to ChatGPT enabled participants in a randomized trial to complete “professional writing tasks” more quickly and improved the quality of their writing. The same year, other research teams identified productivity-enhancing outcomes for computer programmers who used Github’s Copilot, and for customer-support agents who were given access to proprietary A.I. tools. The Media Lab researchers found that many workers are using their personal tools, such as GPT or Claude, at their jobs;

the report refers to this phenomenon as the “shadow AI economy,” and comments that “it often delivers better ROI” than employer initiatives. But the question remains, and it’s one that senior corporate executives will surely be asking more frequently: Why haven’t more firms seen these types of benefits feeding through to the bottom line?

Part of the problem may be that generative A.I., remarkable as it is, has limited application in many parts of the economy. Taken together, leisure and hospitality, retail, construction, real estate, and the care sector—child-minding and looking after people who are old or infirm—employ about fifty million Americans, but they don’t look like immediate candidates for an A.I. transformation.

Another important thing to note is that adoption of A.I. throughout the economy could well be a lengthy process. In Silicon Valley, people like to move fast and break things. But economic history tells us that even the most transformative technologies, which economists refer to as general-purpose technologies, can’t be exploited to maximum effect until infrastructure, skills, and products that can complement them are developed. And this can be a long process. The Scottish inventor James Watt invented his cylindrical steam engine in 1769. Thirty years later, most cotton factories in Great Britain were still powered by water wheels, partly because it was difficult to transport coal for use in steam engines. That didn’t change until the development of steam-powered railways in the early nineteenth century. Electricity also spread slowly and didn’t immediately lead to an economy-wide spurt in productivity growth. As Solow noted, the development of computers followed the same pattern. (From 1996 to 2003, economy-wide productivity growth finally increased, which many economists attributed to the delayed effect of information technology. Subsequently, however, it fell back.)

In some cases, economists argue, new technologies may even reduce productivity growth because they are so disruptive and difficult to incorporate into existing ways of doing things. Only later do productivity gains show through—a pattern known as the “J curve.” Earlier this year, four economists from a variety of institutions published a paper in which they argued that American manufacturing may now be in the down portion of the A.I. J curve. After examining firm-level data on A.I. adoption that they collected in collaboration with the Census Bureau, the economists said they had found evidence that “Short-term performance losses precede long term gains.” In an article about this research which was published by M.I.T.’s Sloan School of Management, one of its authors, Kristina McElheran, a professor at the University of Toronto, wrote, “AI isn’t plug and play. It requires systemic change, and that process introduces friction, particularly for established firms.”

If this thesis is taken at face value, it’s ultimately an optimistic one for businesses—although not necessarily for workers whose skills can be replicated by A.I. (As some entry-level programmers are already finding, the latter have ample reason to be alarmed.) On the technology J curve, once the “frictions” are overcome, productivity takes off. But since the journey along the curve can be lengthy, it’s hard to predict which firms will emerge as the winners and losers. During the commercialization of the internet, many of the ultimate winners didn’t emerge until after the dot-com bubble had burst, in 2000. (Google was founded in 1998, but it

didn't go public until 2004. Facebook wasn't created until 2004, Airbnb until 2008.) There is no guarantee that history will repeat itself. But investors who are still riding the A.I. boom might be wise to cash in some more of their chips. ♦