

Who Will Solve the AI Productivity Puzzle?

di Robin Rivaton

Generative AI tools have been widely available for several years now, but productivity growth has little sustained improvement, because firms have yet to convert task-level time savings into measurable economic output. Is the most effective solution to overhaul existing companies—or to bypass them altogether?

PARIS—Last week, OpenAI and Anthropic both announced that they will launch private-equity-backed consulting businesses to push their models into the workflows of mid-market companies. With backing from Blackstone, Hellman & Friedman, and Goldman Sachs, Anthropic is reportedly raising \$1.5 billion for the venture, whereas OpenAI [has reportedly raised](#) \$4 billion from 19 investors—led by TPG, Brookfield, and Bain—at a \$10 billion valuation.

The combined \$5.5 billion is aimed at cracking the one problem the frontier labs cannot solve with better models: the AI productivity paradox. Although generative AI tools are now widely available, their benefits have yet to show up in the productivity statistics. Gallup finds that the share of American adults using AI tools at work at least a few times a year [grew](#) from under a quarter in early 2023 to almost a majority by 2026.

Yet US nonfarm business productivity growth has shown little sustained improvement. Though there was a brief acceleration in nonfarm business-sector productivity in 2024, it had [fallen back](#) toward its baseline by mid-2025. These trends echo the IT productivity paradox of the 1990s, which eventually resolved itself when firms reorganized around new IT capabilities. The open question is whether the same conditions are in place today, and if so, where.

What's the problem? In a recent [paper](#), I identify organizational plasticity as the key variable. Productivity growth depends on a firm's capacity to convert task-level gains into system-level throughput, and that capacity in turn depends on whether work is legible, processes are modular, task routing is substitutable, and managers have the authority to redesign workflows. Where those conditions are absent, AI saves individuals time without changing what the firm produces.

The standard economics literature has the same blind spot. The [canonical empirical studies](#) by David Autor, [Erik Brynjolfsson](#), Shakked Noy, Whitney Zhang, and their co-authors measure individual productivity on isolated tasks. A customer-service agent may resolve more tickets, a consultant may write better memos, and a programmer may ship usable code faster. The gains are real, but they do not necessarily make the firm move faster.

A [six-month randomized experiment](#) by Eleanor Dillon, Sonia Jaffe, Nicole Immorlica, and Christopher Stanton demonstrates the problem. They gave 7,137 knowledge workers across 66 large firms access to Microsoft 365 Copilot, which was integrated into email, documents, and meetings. Time spent on email dropped, but time in meetings did not, and task quantity and composition remained flat. Most of the hours saved went into fewer evenings working from home rather than additional output. Individual cogs achieved higher efficiency, but the machine ran at the same pace.

The AI labs' new vehicles are an admission that selling tokens does not close this gap. With forward-deployed engineers, multiyear statements of work, deployment templates copied across portfolio companies, and sector-specific workflow redesign, the consulting businesses they are launching will be priced and staffed accordingly. The labs have concluded that the value of the models they have built sits one layer up, in rewiring firms, rather than in the unit economics of an API.

China is moving in the same direction, albeit through a different channel. In March, the open-source AI agent OpenClaw escaped the developer community and achieved [mass-market](#) status among Chinese users. Within days, a platform war had erupted. Zhipu released AutoClaw, a one-click local installation bundled with dozens of pre-loaded skills. ByteDance shipped ArkClaw as cloud software. Tencent rolled out WorkBuddy for enterprise users and pushed QClaw and ClawBot into WeChat.

And Alibaba responded with its own orchestration tools. Because the open-source code was free, the model was no longer a defensible source of rent. Value migrated elsewhere—to default installations, the user interface, integration with messaging apps, and so forth.

Chinese industrial policy has also adapted. For two decades, local governments competed for factories, corporate headquarters, and supply-chain anchors by offering land grants, infrastructure, and tax incentives. Then, in early 2026, several districts and development zones began rolling out subsidy packages for OpenClaw adoption aimed not at firms but at individual operators. Hefei offered compute vouchers. Hangzhou's Xiaoshan district [went further](#). Shenzhen [positioned](#) itself as a hub for AI-driven one-person companies. Guangdong [followed](#) with a provincial framework for AI-enabled solo entrepreneurship, and Sichuan [pursued](#) similar measures.

This approach works because the physical infrastructure is already in place. A single operator with a laptop in Shanghai can plug into the supply chains of Dongguan or Suzhou. Subsidized compute (processing power), a pre-installed agent, integrated messaging, embedded payments, and direct access to manufacturers, logistics providers, and marketplaces are enough to turn an individual into a viable unit of production. AI agents transform the firm externally rather than redesigning it from within.

The American bet is on expertise. Low-plasticity firms must be made plastic from the inside, with paid teams of forward-deployed engineers working through one mid-market company at a time. But while this approach could unlock a great deal of value, it is slow and capital-intensive.

The Chinese strategy is more radical. It starts with the insight that the firm itself is the bottleneck. Instead of raising the plasticity of incumbents, China is lowering the minimum organizational scale required to produce. The result is a blossoming of free agents, default installation, super-app integration, public subsidies, and millions of micro-businesses iterating in parallel.

Neither approach is guaranteed to pay off. American consultants may find that organizational change resists scale, no matter how many engineers they embed. Chinese local governments may subsidize a wave of low-quality automation and

speculative one-person ventures that crash and burn. On both sides, the hard problem is the same: turning task-level time savings into measurable economic output. AI competition has moved past chips, data-center capacity, and frontier model performance. What matters now is the social architecture through which intelligence is deployed.