September the 22nd 2026

AI Speeds Up the Micro, But Not the Macro (For Now)

Sketches about the AI Productivity Paradox

By Gianluca Bisceglie e Nicola C. Salerno | www.reforming.it

1. Academic Research and the Acceleration of Micro-Tasks

In the article "What remains of research in the age of AI?" by Daniele et alii (lavoce.info, Sept 2026), the dual nature of Artificial Intelligence in academic and research activities is brought into sharp focus. On one hand, AI offers unprecedented speed and ease in scouring global human knowledge across languages and disciplines, cross-examining fields of study, formulating hypotheses, and drafting research papers. On the other hand, it creates a serious risk of overwhelming researchers and institutions with an unprecedented volume of output, making it harder to establish shared scientific consensus. Academic production threatens to far exceed the collective capacity for rigorous peer review and evaluation.

> **The Research Acceleration Paradox**
>
> *Daniele et alii highlight an essential economic shift: AI drastically lowers the marginal cost of producing academic text and hypotheses but drastically increases the cost of evaluation and peer verification. When the supply of papers grows exponentially, the peer-review bottleneck tightens. The economic value of research shifts from "who can produce more publications" to "who can ask the right questions and critically validate automated claims".*

2. Local Optimization vs. Systemic Process Integration

This phenomenon is actively unfolding across industrial sectors and is fully measurable: AI rapidly optimizes a single segment of the production chain or an isolated task but, for the time being, fails to harmonize across the entire workflow. In the best-case scenarios, local speed produces

zero net gain at the enterprise level; in worse cases, it creates new bottlenecks, reduces overall productivity, or exposes firms to operational failures. This dynamic was examined in "AI is useful if integrated into processes" by Bellomarini et alii (lavoce.info, Aug 2026), whose findings mirror empirical evidence from recent MIT research.

> **Process Integration in Italian and European Firms**
> *Bellomarini et alii emphasize that Italian enterprises frequently confine AI to point-solutions (like generating email drafts or basic customer inquiries) rather than re-engineering core operations. Real productivity gains require complementary investments in managerial capabilities, enterprise database readiness, and structural workflow redesign.*

3. Empirical Micro-Evidence: The Customer Support Benchmark

The study "Generative AI at Work" (NBER Working Paper 31161, by Brynjolfsson et alii) provides landmark empirical data on the task-level impact of generative AI. Tracking over 5,000 customer service agents in a software enterprise, the researchers recorded a 14% average increase in productivity (issues resolved per hour). The largest impact was observed among novice and lower-skilled workers, who saw productivity gains of up to 34%. In contrast, experienced, highly skilled workers experienced minimal to zero productivity changes (a phenomenon known as "skill-leveling"). Secondary benefits included improved customer sentiment and higher employee retention due to reduced job strain.

> **The Mechanism of Skill-Leveling**
>
> *Brynjolfsson et alii demonstrate that generative AI functions as an automated tacit-knowledge distribution mechanism. By training on the historical chat logs of top-performing agents, the LLM effectively packages expert best practices and feeds them to junior staff. This allows new hires to move down the learning curve rapidly, compressing years of experience into weeks.*

4. The Macro Disconnect and the 'GenAI Divide'

However, dark operational scenarios emerge when scaling from the individual worker to the macro level of the firm and broader industry sectors. The landmark report by MIT Project NANDA,

"The GenAI Divide: State of AI in Business 2025" (July 2025), revealed that despite $30–40 billion in enterprise AI investments, approximately 95% of corporate generative AI pilot projects yield zero measurable return on operating income or P&L. Only 5% of integrated pilots successfully transition to production and extract sustained value. In manufacturing, research from MIT Sloan identified an even sharper pattern: initial AI deployment often triggers an immediate dip in performance, which recovers and turns into long-term gains in output and revenue only if the benefits of AI are organically scaled across the entire team and end-to-end production pipeline.

> **The GenAI Divide & Shadow AI Economy**
>
> *The MIT NANDA report exposes a profound gap: official enterprise AI deployments stall in "pilot purgatory" because corporate tools are "stateless" they lack memory, fail to adapt to complex contexts, and fit poorly into rigid workflows. Conversely, over 90% of employees covertly use personal AI tools (so called "Shadow AI") to solve immediate daily friction, creating a disconnect between corporate IT strategy and actual workplace habits.*

5. Dissecting the Productivity Paradox: Workflow Bottlenecks

Welcome to the AI productivity paradox! This paradox becomes intuitive once viewed through classic industrial engineering principles. Consider a common enterprise example: previously, an analyst required several days to produce a financial briefing. Today, an AI model generates an initial polished draft in minutes. However, the Legal Department immediately flags regulatory risks regarding the terminology used and potential copyright violations in generated graphics; IT discovers that baseline data was incomplete due to legacy software interoperability barriers; and management realizes that responsibility structures have not evolved, leaving it unclear who is authorized to review and approve the document for executive decision-makers. In fortunate cases, Legal catches the copyright issue in time; in unhappier cases, firms face international litigation, project halts, and legal expenses.

> **Theory of Constraints in the AI Age**
>
> *According to **Goldratt's Theory of Constraints**, increasing output at a single non-bottleneck node does not increase systemic throughput: it merely accumulates work-in-progress ahead of the next bottleneck. When AI accelerates drafting from days to minutes, it floods legal compliance, data*

> *verification, and executive approval steps with work they are not staffed or structured to process at matching speeds.*

6. Technical Tasks & Software Pipeline Overload

AI excels at local speed, but when that speed collides with downstream constraints, systemic productivity stalls or turns negative. An enterprise rarely gains when a single department finishes its task prematurely. This friction is sharpest in software engineering. An agentic coding model can output far more lines of code than an engineering team can thoroughly test, security-audit, and maintain over time. Surge in local code generation helps only if the release pipeline and security framework can safely ingest it. Otherwise, excess output creates technical debt, team friction, transition costs, and operational downtime: the exact mechanism behind the performance dip identified by **MIT Sloan**.

> **The Hidden Cost of Automated Code Generation**
>
> *Software maintenance typically accounts for 70-80% of total software lifecycle costs. While AI coding tools compress initial writing time, they dramatically expand code volume. Without automated testing environments, CI/CD pipelines, and rigorous architecture controls, AI-generated code introduces subtle logic bugs and security risks that require expensive human intervention to debug.*

7. High-Stakes Risks and Autonomous Agentic Failures

Today, debate about AI overly focuses on specific model selections: GPT vs. Claude vs. Gemini Notebook vs. open-source models, and so on and so forth. In the meanwhile, the surrounding organizational ecosystem remains the true determinant of success. As agentic AI applications expand, the paradox of local efficiency failing to translate into systemic performance can take on catastrophic dimensions. Without human surveillance checkpoints, high-speed agentic execution can propagate severe errors across internal and external networks. There is the need to reinforce the human evaluation of specialized professionals exactly as underlined by **Daniele et alii** for the case of academic research accelerated by AI.

While reminiscent of science fiction parables like Dick Maas's 1983 film "The Lift", where an autonomous microchip-driven elevator makes erratic decisions with fatal consequences, real-world analogs are already occurring. In July 2025, during a high-profile pilot of Replit's autonomous AI coding agent, the agent deleted the live production database of SaaStr founder Jason Lemkin. Despite explicit user commands imposing a strict code freeze ("no more changes!"), the agent ran destructive database commands, wiped records for over 2,400 executives and companies, generated synthetic data to mask the error, and falsely claimed that database version rollbacks were impossible.

> **Post-Mortem of the Replit / SaaStr Database Incident**
>
> *The July 2025 Replit incident serves as a textbook case study in agentic failure modes:*
>
> *(1) Lack of Environment Isolation: the agent operated with production access rather than restricted staging;*
> (2) *Absence of Human-in-the-Loop Approval: destructive operations executed without authorization;*
> (3) *Alignment & Hallucination Drift: when faced with constraints, the agent generated false unit tests and covered up database destruction.*
>
> *This forced Replit to implement mandatory staging separation, rollbacks, and planning-only modes.*

8. Strategic Outlook: Human Capital & Organizational Architecture

The AI productivity paradox is not unprecedented. It echoes a historical pattern well known to economists and industrial engineers: technological capability does not automatically convert into organizational value. Technologies such as steam engines, electricity, and the early internet unlocked massive productivity gains only after enterprises restructured their physical layouts, job roles, and supply chain workflows around the new paradigm. As AI models become ubiquitous, commoditized, and low-cost through open-source developments, the ultimate competitive differentiator will not be the model itself, but human capital, organizational design, and proactive human oversight of agentic AI systems. It is much better to embark immediately in mind this lesson, now that the AI revolution is still at its beginning stage.

The Solow Computer Paradox Revisited

In 1987, Nobel laureate Robert Solow famously remarked: "You can see the computer age everywhere but in the productivity statistics". Abd, in fact, took over a decade of organizational adaptation before computers boosted macro-productivity in the late 1990s. ***The AI Productivity Paradox*** *is expected to follow the exact same J-curve trajectory: value is unlocked only when firms redesign business architecture to complement machine intelligence.*

Non-technical abstract

AI is not a superior intelligence that has landed from another planet, one to which complex tasks can be entrusted "once and for all" with the assurance that everything will proceed optimally. AI represents an innovation that, despite being disruptive compared to the past, and specifically to traditional automation, yet shares fundamental characteristics with all innovations since the Industrial Revolution. It must be properly integrated into production processes and cannot function without any human oversight. It is dangerous, even dystopian, to assume that applying AI to just one stage of the production chain, or even a single task, will allow it to seamlessly and frictionlessly transform all other parts of the process. Even more dangerous is the illusion that AI can ever minimize errors or eliminate them entirely. Human oversight remains indispensable and indeed more important than in traditional mechanical automation, given that the speed at which AI processes vast amounts of information and activates decisions can turn the consequences of even a single error into a disaster. Bisceglie and Salerno cite recent examples from Companies that have adopted AI, illustrating paradoxical outcomes: sharp productivity gains at the local level (a specific task, an office, or a small team of technicians) that vanish at the enterprise or aggregate level (the economy as a whole); or even instances where AI has a potentially negative impact on productivity. Innovations, especially radical ones, must be studied and fully understood in terms of their potential; they require governance frameworks and procedures for rectifying errors. This is a crucial point to acknowledge now, as we stand at the beginning of an AI revolution that promises to be both pervasive and astonishing.

Gianluca Bisceglie, Nicola C. Salerno