AI & Society Institute

PSL ENS Dauphine UNIVERSITÉ PARIS

Does AI Really Boost Productivity?

The Paradox, the Evidence, and What Policy-Makers Should Do About It

Éric Hazan, Ardabelle Capital, HEC, Sciences Po

1794

Eric Hazan is the Head of Value Creation and Transformation, as well as Strategic Partnerships at Ardabelle. He is a graduate of HEC, and Lectures on Corporate Strategy and Digital Transformation at HEC and Science Po Paris. He is also the author of numerous articles and reports on tech, AI and the economy, and author of the books « Révolution par les territoires » with Frédéric Salat Baroux, and "Faut-il encore décider?" with Olivier Sibony.

The **AI and Society Institute** fosters the responsible development and use of AI by studying its interactions with society. It produces analyses grounded in scientific methods to inform policy in France and Europe. It is led by **Hugo Mercier** (Scientific Director) and **Constance de Leusse** (Executive Director).

Published in May 2026.

 Suggested citation: Hazan, E., *Does AI Really Boost Productivity? The Paradox, the Evidence, and What Policy-Makers Should Do About It.* AI and Society Institute, 13 May 2023.

Table of content

Key messages

1. Task-level experiments consistently show that AI boosts individual productivity by 14–40%, particularly for less-experienced workers. Yet aggregate productivity statistics and CEO surveys show only limited and uneven effects to date, a pattern reminiscent of the Solow paradox of the 1980s, though early growth-accounting evidence suggests the macro signal may be beginning to emerge.

2. The gap between micro and macro is explained by a well-documented pattern: general-purpose technologies require years of complementary investment, in skills, processes, organisational redesign, and data infrastructure, before their aggregate effects materialise.

3. In Europe, AI adoption has more than doubled in two years (8.7% to 20% of firms; Eurostat 20.0%, OECD 20.2%), but remains sharply unequal: 52% of large firms versus 17% of small firms, and a tenfold gap between leading and lagging member states.

4. AI's impact depends less on the technology than on organisational choices: Harvard-BCG experiments show that identical AI tools produce three radically different outcomes, upskilling, newskilling, or no skilling, depending on how workers integrate them. The default market trajectory favours automation over augmentation.

5. To convert AI's demonstrated micro-level potential into measurable aggregate productivity growth, Europe needs coordinated action on five fronts: workforce skills, SME adoption infrastructure, organisational transformation, the direction of AI development itself, and resource reallocation.

The Paradox: AI Everywhere Except in the Statistics

In February 2026, financial analysts at CitriniResearch published a thought experiment that sent shockwaves through investment circles: a fictional dispatch from June 2028 in which AI had worked *too well*. In their scenario, productivity surged, margins expanded, stock markets hit record highs, and then the economy collapsed. White-collar workers, displaced by machines that did their jobs better and cheaper, stopped spending. Mortgages defaulted. A "negative feedback loop with no natural brake" produced an unemployment rate of 10.2% and a 38% market crash. The most unsettling aspect of the exercise was not its conclusions but its mechanism: every individual corporate decision was rational; the collective result was catastrophic.[1]

The scenario is fiction. But the paradox it dramatises is real. In 1987, the economist Robert Solow observed: *"You can see the computer age everywhere but in the productivity statistics."* Nearly four decades later, the same paradox has returned. Generative AI is transforming daily workflows across knowledge work: by December 2025, 35.9% of American workers reported using generative AI tools. Yet the macroeconomic signature remains almost invisible. The question this brief seeks to answer is not whether the Citrini scenario will materialise, but what determines whether AI's demonstrated micro-level productivity gains translate into shared prosperity or concentrated disruption.[2][3]

A landmark NBER study published in February 2026, surveying approximately 6,000 CEOs, CFOs, and senior executives from firms in the United States, the United Kingdom, Germany, and Australia, found that more than 80% of managers reported no change in productivity, measured as sales volume per employee, over the preceding three years, despite AI adoption among surveyed firms rising from 61% to 71% between early 2025 and early 2026. PwC's 2024 Global CEO Survey reached a similar conclusion: only 30% of CEOs reported any revenue increase from AI, while 22% said their costs had actually risen.[4][5]

> "AI is everywhere except in the incoming macroeconomic data." — *Torsten Slok, Chief Economist, Apollo Global Management, 2025. Researchers at the Federal Reserve Bank of St. Louis estimate that US labour productivity has run only about 1.9 percentage points above its 2015–2019 pre-pandemic trend in the roughly two and a half years since ChatGPT was released — an excess of just 0.7pp per year.*[6][7]

The University of Virginia economist Anton Korinek, in a March 2026 exchange with David Autor and Natasha Sarin, draws attention to a telling asymmetry in the investment data. Alphabet, Meta, Microsoft, and Amazon have collectively spent more than $300 billion on AI infrastructure in the past year alone, more than triple their spending just a few years ago. Yet the companies at the centre of the AI revolution employ remarkably few people: OpenAI has roughly 4,000 employees for a $500 billion valuation (7–8 employees per billion dollars of market capitalisation), Anthropic 2,300 at $350 billion. By comparison, Walmart has roughly 2,000 to 3,000 employees per billion of market capitalisation, depending on the valuation date, and Ford comparable figures. The investment is enormous; the employment it creates is minimal. Korinek argues that the employment effects we are waiting to see may simply be "lagging indicators of a transformation that's already locked in by the capital being deployed."[8]

The range of expert forecasts is itself revealing. At the cautious end, the MIT economist Daron Acemoglu estimates that AI will increase total factor productivity by approximately 0.53% over the next ten years, a figure far below transformative claims. At the optimistic end, Goldman Sachs projects that generative AI could raise global GDP by 7% (nearly $7 trillion) and lift productivity growth by 1.5 percentage points over a decade. McKinsey Global Institute estimates annual economic value of $2.6–4.4 trillion from generative AI alone.[9][10][11]

Philippe Aghion occupies an instructive middle ground. In a widely cited 2024 paper with Simon Bunel, and a subsequent 2025 essay with Xavier Jaravel, Aghion estimates that AI should increase aggregate productivity growth by 0.8 to 1.3 percentage points per year over the coming decade, comparable to the productivity acceleration delivered by electricity (1.3 points) or information technology (0.8 points) in their respective eras. Crucially, Aghion's framework emphasizes that this gain is conditional on institutional adaptation: countries that reorganise work, invest in skills, and reform labour markets will capture the dividend; those that do not will see AI reinforce existing inequalities rather than resolve them.[12][13]

Yet in the short term, even optimists concede the gap is real. Goldman Sachs' own chief economist, Jan Hatzius, acknowledged in early 2026 that AI investment had contributed "basically zero" to US GDP growth in 2025, despite hundreds of billions in corporate expenditure. The Federal Reserve Bank of San Francisco, in a February 2026 Economic Letter, frames the situation as an "AI moment", rich with long-run potential but with measurable aggregate effects still largely in the future. Growth-accounting work by Bontadini, Corrado, Haskel, and Jona-Lasinio does, however, detect an early signal: software and AI-related capital already account for approximately 50% of the 2% average US labour productivity growth observed between 2017 and 2024, and for half of the most recent productivity acceleration — suggesting that the macro effects may be beginning to materialise, albeit unevenly.[14][15][16]

What Experiments Actually Show: The Micro-Level Evidence

The apparent contradiction between subdued macro-level statistics and the experimental evidence reviewed below is not an artefact of measurement error. At the task and individual level, the evidence of AI-driven productivity gains is robust, consistent, and in some cases, remarkable.

The most rigorous study to date is Brynjolfsson, Li, and Raymond's *"Generative AI at Work"* (2025). Studying 5,172 customer-support agents at a Fortune 500 software firm, they found that access to an AI assistant increased the number of issues resolved per hour by 14% on average. The effect was sharply heterogeneous: novice and low-skilled workers improved by 35%, while the most experienced agents saw marginal or negligible gains. The study also found improvements in customer sentiment, employee retention, and evidence of accelerated learning — the AI appeared to disseminate the tacit knowledge of top performers to less experienced colleagues.[17]

A second experiment involved 758 consultants at Boston Consulting Group, conducted by researchers from Harvard Business School, MIT Sloan, and Wharton. Participants using GPT-4 completed 12.2% more tasks, 25.1% faster, and produced 40% higher-quality output for tasks within AI's capability frontier. However, the study introduced a critical concept: the *"jagged technological frontier"*. For tasks outside this frontier, consultants using AI were 19% less likely to produce correct solutions than those working without it — suggesting that AI can actively degrade performance when users rely on it for tasks it cannot reliably handle.[18]

Noy and Zhang (2023) found that ChatGPT reduced the time required for professional writing tasks by 40% while increasing output quality by 18%. A consistent finding across studies is that AI compresses the skill distribution: it lifts the floor far more than the ceiling, making less-skilled workers substantially more productive while offering modest benefits to experts.[19]

Taken together, the experimental literature presents a remarkably consistent picture: task-level productivity gains of 14–40%, concentrated among less-experienced workers, with significant quality improvements for tasks within AI's competence but potential quality degradation beyond it. The Berkeley *California Management Review* identified seven recurring myths in the public discourse on AI and productivity; chief among them is the assumption that task-level gains automatically translate into organisational or economy-wide productivity improvements. The evidence to date suggests that this translation is neither automatic nor immediate.[20]

Why the Gap? General-Purpose Technologies and the J-Curve

The pattern we observe with AI is not new. It is, in fact, the defining signature of general-purpose technologies — innovations so fundamental that they reshape production across virtually all sectors, but only after extended periods of complementary investment and institutional adaptation. Bresnahan and Trajtenberg formalised this concept in 1995; the canonical examples are the steam engine, electricity, and the microprocessor.[21]

Brynjolfsson, Rock, and Syverson have modelled this dynamic as a "productivity J-curve". When a major new technology arrives, firms must invest heavily in intangible capital — new business processes, worker training, data infrastructure, organisational redesign — before the technology's productive potential is fully realised. During this transition period, measured productivity may actually decline, because the investment expenditure is counted as a cost but the complementary capital being built is not captured in conventional output statistics.[22]

The historical precedent is instructive. The economic historian Paul David showed that the electric dynamo, introduced in the 1880s, did not produce measurable aggregate productivity gains until the 1920s — a lag of nearly four decades. The reason was not that electricity did not work; it was that factories had to be physically redesigned, from the multi-storey shaft-driven architecture of steam power to the single-storey flexible layout that electric motors made possible. Similarly, the personal computer diffused through offices from the late 1970s, but the productivity surge associated with IT, an increase from 1.4% to 2.7% annual productivity growth, did not arrive until 1995–2005.[23][24]

> *The Workday 2026 study illustrates the transition cost concretely: 37–40% of time ostensibly saved by AI is currently consumed by reviewing, correcting, and verifying AI-generated output. AI does not yet automate tasks end-to-end; it displaces effort from production to supervision. Until* organisations *redesign workflows to account for this — distinguishing tasks where AI output can be trusted from those where it cannot — the net productivity gain will remain well below the gross task-level improvement.*[25]

Carl Benedikt Frey, in his 2025 book *How Progress Ends* and in subsequent commentary in early 2026, argues that the J-curve metaphor, while useful, is incomplete. The deeper bottleneck is not merely time and complementary investment, it is *institutions and incentives*. Frey notes that the personal computer and the internet "gave us the world's stored knowledge in our pockets", yet the resulting productivity surge lasted barely a decade and was largely confined to the United States. The reason, he argues, is that the institutions surrounding innovation, publish-or-perish academic incentives, regulatory bottlenecks, declining business dynamism, prevented the sustained diffusion of gains. The lesson for AI is clear: without reform of the institutional environment, competition policy, labour market regulation, skills systems, even a powerful general-purpose technology can produce a productivity surge that peaks and fades rather than compounding over time.[26][27]

Frey and Osborne's original 2013 working paper, which estimated that approximately 47% of US jobs were at "high risk" of automation, catalysed the modern debate on AI and labour. A decade later, Frey acknowledges that the timeline was too aggressive — but insists that the structural forces remain operative. The critical variable, he argues, is not the pace of AI capability but the structure of the economy receiving it. An economy dominated by large incumbents optimising existing processes will experience AI very differently from one in which new entrants are creating entirely new sectors and occupations.[28]

Beyond the J-Curve: Why Firms Must Change How They Work

The J-curve framework explains why aggregate productivity gains are delayed. But it does not, by itself, explain what firms must actually do to accelerate the transition. A growing body of research now provides a far more granular picture. The central finding is that AI's productivity impact depends less on the technology than on the organisational choices surrounding its deployment.

In February 2026, Daron Acemoglu, David Autor, and Simon Johnson published a paper for the Hamilton Project at Brookings: *"Building Pro-Worker Artificial Intelligence"*. Building on the analytical framework of their 2023 book *Power and Progress*, the three MIT economists argue that AI technologies fall into five economically distinct categories, ranging from labour-augmenting tools that make workers more effective at existing tasks to fully autonomous systems that eliminate human involvement entirely. Their core thesis is that the market, left to its own devices, systematically under-invests in the pro-worker end of this spectrum: the business models of the major AI vendors — built on automation, scale, and data capture — are structurally misaligned with worker augmentation.[29][30]

Acemoglu identifies three barriers to pro-worker AI. First, business model misalignment: "The business models of tech companies are not really aligned with that pro-worker dimension." Second, a one-size-fits-all architecture: general-purpose models are not designed for the highly contextual needs of professionals — electricians, teachers, nurses — whose expertise depends on domain-specific judgment and real-world context. Third, the gravitational pull of artificial general intelligence, which Acemoglu characterises as "the enemy of pro-worker AI" because it shifts attention from systems designed to enhance human capability toward systems designed to replace it. Without deliberate intervention, automation — not augmentation — becomes the default path.[31][32]

The experimental evidence supports this analysis with striking specificity. A field study of 244 management consultants at Boston Consulting Group, conducted by researchers from Harvard, MIT Sloan, and Wharton, identified three distinct modes of human-AI collaboration emerging from use of the same tools: *Centaurs*, who strategically delegate specific tasks to AI while retaining authority over the overall workflow; *Cyborgs*, who fuse their work with AI in a continuous, iterative loop; and *Self-Automators*, who effectively abdicate cognitive authority by handing entire workflows to the model. The skilling outcomes diverge dramatically: Centaurs develop deeper domain expertise; Cyborgs acquire new AI-related capabilities; Self-Automators experience no skill development at all. Professionals using the same technology, in the same organisation, arrived at radically different productivity and learning outcomes depending on how they chose to integrate it.[33]

The Burning Glass Institute's analysis of US job postings spanning three years before and after ChatGPT's release reaches a complementary conclusion. Burning Glass classifies individual skills (rather than jobs) as exposed to automation or augmentation, and tracks the resulting shifts in employer demand. Skills exposed to automation are 16% more likely to see demand decline; skills exposed to augmentation are 7% more likely to see demand increase. The critical finding is that automation- and augmentation-exposed skills cluster within the same occupations: "The jobs most exposed to automation are not a separate population from the jobs most exposed to augmentation. They are, in large part, the same jobs." The binary framing — will AI replace workers or help them? — is empirically false. Both forces operate simultaneously within the same roles, and the balance between them is determined by organisational choices, not by the technology alone.[34]

> *A Harvard Business Review analysis adds a further dimension: AI does not simply reduce work — it intensifies it. As AI handles routine tasks, the cognitive burden shifts to complex judgment, interpretation, and decision-making. Decision velocity increases while reflection time decreases. McKinsey's task-level analysis of 800 occupations finds that more than 70% of employer-sought skills remain relevant in both automatable and non-automatable work — but the demand for the ability to use and manage AI tools has grown sevenfold in two years, faster than any other skill, including the ability to design AI systems themselves.*[35][36]

The implications for European policy are direct. The World Economic Forum's January 2026 scenario analysis identifies four plausible futures for AI and work by 2030, differentiated not by the pace of technological progress but by the quality of institutional response: a "Co-Pilot Economy" in which early investment in training, mobility, and governance creates absorptive capacity; a "Supercharged Progress" scenario of rapid AI breakthroughs with new occupations emerging at scale; an "Age of Displacement" where AI advancement outpaces workforce adaptation; and a "Stalled Progress" in which patchy skills investment leaves gains concentrated among those already ahead. Notably, 54% of business executives expect AI to displace existing jobs, but only 12% expect it to lead to higher wages — a gap that itself reflects the absence of deliberate pro-worker design.[37]

Data from Anthropic's survey of 80,508 individuals across 159 countries and 70 languages confirms that the gap between AI's promise and its perceived impact is global: while 32% of respondents cite productivity gains as AI's primary benefit, 22.3% identify economic disruption as their top concern — a ratio that suggests widespread awareness of AI's potential paired with deep uncertainty about its distributional consequences. The scale of the survey — the largest known multi-country consultation on AI's societal implications — makes clear that the organisational and policy choices discussed in this brief are not abstract: they correspond to anxieties and expectations held by millions of workers worldwide.[38]

The research paper by Bontadini, Corrado, Haskel, and Jona-Lasinio provides the growth-accounting complement to these findings. Analysing AI as both a general-purpose technology and an "innovation in the method of innovation", they estimate that software and AI already account for 50% of the 2% average US labour productivity growth observed between 2017 and 2024, and for 50% of the 1.2 percentage-point acceleration in the most recent period. Their central estimate is that AI will add up to 1 percentage point of annual labour productivity growth in the United States — but only approximately 0.3 points in Europe, reflecting the continent's smaller ICT sector, lower capital deepening, and slower organisational adaptation.[39]

Frey's historical analysis adds a structural dimension to these findings. He distinguishes between process innovation — in which large firms use AI to optimise existing operations (automation) — and product innovation — in which new entrants create entirely new sectors. "If all we'd done since 1800 was automation, we would have productive agriculture. Cheap textiles. But that would be about it." The implication is that an economy in which AI is deployed primarily for cost reduction within incumbent firms will miss the larger prize: the creation of new industries, new occupations, and new sources of demand. Market concentration — which has increased steadily across OECD economies — is therefore not merely a competition problem; it is a productivity problem, because it shifts the balance of AI deployment away from the kind of innovation that historically drove sustained growth.[40][41]

The MIT economist David Autor frames the underlying risk with a historical analogy that illuminates the stakes. In the artisanal era, most goods were handmade by skilled workers whose expertise was revered. The Industrial Revolution decimated the value of that expertise: it took five decades before working-class living standards began to rise. Autor warns that AI threatens a comparable "commodification of human expertise" — not necessarily destroying jobs, but hollowing out their economic value. As he puts it: "We should be worried not about the number of jobs but about the commodification of human expertise, since expertise is what gives labour its economic value." Korinek goes further: if artificial general intelligence succeeds, "labour itself becomes optional for the economy" — and the reassuring historical patterns in which new technologies create more jobs than they destroy would no longer apply, because machines would learn the new jobs faster and do them more cheaply.[42]

Aghion, Bergeaud and colleagues have shown that productivity depends not only on invention and adoption, but on the economy's capacity to reallocate capital, labour, and market share toward more productive firms. Their research demonstrates an inverted-U relationship between credit access and productivity growth: too little credit starves innovative firms of capital, but too much flows to incumbents and sustains unproductive ones. If AI is layered onto stagnant market structures, weak diffusion channels, or distorted credit allocation, task-level gains may coexist with disappointing aggregate outcomes.[43]

The European Picture: Rapid Adoption, Deep Inequalities

Europe's AI adoption trajectory has accelerated sharply. Eurostat data show that the share of EU enterprises with ten or more employees using AI technologies rose from 8.7% in 2023 to 13.5% in 2024 and 20.0% in 2025 — a 130% increase in two years. The OECD, using a slightly broader sample, reports a comparable figure of 20.2% for 2025. The ICT sector leads at 57.3%, followed by professional and scientific services at 36.8%.[44][45]

But aggregate figures conceal profound disparities. By firm size: 52% of large enterprises used AI in 2025, against 30% of medium-sized firms and just 17.4% of small firms — a gap of nearly 35 percentage points between the largest and smallest firms. By geography: Denmark leads Europe at 42.0%, followed by Finland (37.8%) and Sweden (35.0%). At the other end, Romania stands at 5.2%, Poland at 8.4%, and Bulgaria at 8.5% — a tenfold gap between the leading and lagging member states.[46][47]

The IMF estimates that the medium-term productivity gains from AI for Europe as a whole will be modest: approximately 1.1% cumulatively over five years. However, a separate EIB study, using US peer adoption rates as an instrument, finds that AI adoption increases EU labour productivity by 4% on average, driven primarily by capital deepening rather than job displacement — suggesting that if adoption rates could be raised to US levels, the aggregate gains would be substantially larger.[48][49]

> *The central finding is that Europe's AI productivity challenge is not principally one of technology or even of investment. It is one of diffusion. The benefits are real but sharply concentrated — among large firms, high-skill workers, advanced member states, and ICT-intensive sectors. The policy question is how to broaden diffusion without diluting the quality of adoption.*[50]

Frey, drawing on IMF estimates, adds a striking figure to the European picture: barriers to trade in services inside the EU amount to approximately 110% — the equivalent, as he puts it, of "Trump Liberation Day tariffs self-imposed on services inside the European Union." Given that AI's productivity potential is concentrated in knowledge-intensive services — precisely the sectors most affected by these internal barriers — the completion of the Digital Single Market is not merely a trade policy objective; it is a precondition for capturing AI's productivity dividend at scale.[51]

Policy Recommendations

The following eight recommendations draw on the experimental, macroeconomic, and organisational evidence set out above. They are designed to address the five structural bottlenecks — skills, SME diffusion, organisational transformation, the direction of AI development itself, and resource reallocation — that explain the gap between AI's demonstrated micro-level potential and its absent macro-level impact.

1. A European AI Skills Guarantee. The OECD has identified workforce skills as the decisive factor in whether economies capture AI's productivity potential. The experimental evidence consistently shows that less-skilled workers benefit most from AI assistance — but only if they receive adequate training in how to use it effectively. Europe should establish a binding AI Skills Guarantee: every employee in a firm with more than 50 staff should have access to at least 40 hours of certified AI training by 2028, co-funded by the European Social Fund and national training levies. Priority should be given to sectors with the highest latent AI productivity potential: professional services, healthcare, public administration, and manufacturing. Two operating models are instructive. Singapore's SkillsFuture, launched in 2015, has lifted the annual training participation rate of the resident workforce from roughly 35% in 2015 to around 50% by 2021, with about 500,000 individuals — close to a fifth of the workforce — using subsidised programmes each year; cumulatively, more than half of the resident workforce has activated its SkillsFuture Credit since the scheme's introduction.[52] France's Compte Personnel de Formation, which in 2015 converted a collective bargaining right into individual, portable training accounts, allocates €500 per year, capped at €5,000 — and €800 per year, capped at €8,000, for low-skilled workers — and has, since its launch, attracted low-skilled workers in proportions that exceed their share of the workforce.[53] Annual training participation among that group rose markedly in the five years following the scheme's launch. Recent moves to introduce a €100 co-payment in 2024, however, illustrate the design tension between scale and quality control that any European scheme will need to address.[54] Both programmes suggest that uptake depends critically on portability and on employer co-investment[55] — design features that the Council of the European Union formalised in its 2022 Recommendation on Individual Learning Accounts,[56] and that any future EU Skills Guarantee should incorporate from the outset. The case for accelerating this agenda is reinforced by recent field-experimental evidence that the productivity gains of generative AI accrue overwhelmingly where deployment is paired with worker training and workflow redesign.[57][58]

A recent proposal by former US Commerce Secretary Gina Raimondo offers a concrete template. Raimondo calls for a "grand bargain" between government, industry, and workers: public investment in modular, stackable credentials — short-duration, industry-recognised certificates that can be completed while working — paired with employer tax credits for on-the-job AI training. She cites the CHIPS Act partnership between TSMC and Maricopa Community College in Arizona as proof of concept: a program that trained thousands of semiconductor technicians in months, not years, through a combination of classroom instruction and paid apprenticeship. A European adaptation of this model, linking AI training vouchers to the existing European Qualifications Framework, could accelerate workforce readiness without requiring the creation of new institutional infrastructure.[59]

2. An SME AI Adoption Facility. The adoption gap between large and small firms is the single largest obstacle to translating task-level productivity gains into aggregate growth. The 35-percentage-point gap between large and small enterprises is not primarily a technology gap — it is a gap in organisational capacity, data readiness, and integration expertise. A European SME AI Adoption Facility, modelled on the G7 SME AI Adoption Blueprint and co-financed by the EIB and national development banks, should provide three services: subsidised access to pre-configured, sector-specific AI tools; on-site implementation support through a network of accredited AI advisors; and shared data infrastructure enabling SMEs to pool anonymised data for AI training without surrendering competitive advantage. Existing programs show what such a facility can achieve. Germany's Mittelstand-Digital initiative operates through regional and sector-specific SME digital hubs, combining advisory services, training, demonstrators and hands-on support; the BMWK's 2025 Ramboll evaluation of the associated "Digital Jetzt" investment grant found that, between programme launch in September 2020 and December 2022, €134.3 millions of grants triggered roughly €447 million in additional revenue at the supported firms. In the United States, the Manufacturing Extension Partnership reaches more than 36,000 manufacturers a year through a federal–state network of 51 centres and hundreds of service locations; on NIST's FY2023 reported figures, it generated $24.60 in new sales growth for every federal dollar invested, although these impact figures should be read with the usual caution attached to survey-based program evaluations. Japan's Kohsetsushi technology centres provide an even older precedent: locally embedded public technology centres offering SMEs technical consultation, training, testing facilities, applied R&D and access to equipment.

The lesson is not that subsidies alone work. It is that successful diffusion systems pair subsidised tools with trusted, on-site human expertise. This matters especially for AI, because adoption is not a one-off software purchase: as the OECD notes, SMEs need complementary investments in skills, data, compute, organisational change and finance before productivity gains materialise, often with a J-curve dynamic. A European AI adoption facility should therefore be sized for comparable reach, for example, at least one accredited adoption hub per 10,000 SMEs, and funded on a multi-year basis, so that firms have time to absorb the complementary investments that effective AI deployment requires.[60][61]

3. Organisational Transformation Vouchers. The J-curve evidence demonstrates that complementary intangible investment — process redesign, management practices, data architecture — is as important as the technology itself. Yet current EU innovation support programmes focus overwhelmingly on technology acquisition, not on the organisational changes required to deploy it productively. A new instrument, AI Transformation Vouchers of €50,000–200,000 per firm, should fund external consulting on workflow redesign, change management, and AI integration — replicating at scale the kind of controlled deployment that produced measurable gains in the Brynjolfsson and Dell'Acqua experiments. The Netherlands' innovation voucher scheme, pioneered in 2004 with €7,500 vouchers allocated by lottery, provided some of the only randomised-controlled-trial evidence in innovation policy: Cornet, Vroomen and van der Steeg (CPB, 2006), exploiting the lottery allocation, found significant short-run effects on SME–research interaction.[62] Subsequent register-based evaluations have suggested that these short-run effects translated, at least in part, into longer-term performance gains.[63] The voucher should be conditional on the firm reporting anonymised productivity metrics to a public AI Observatory for three years, and a meaningful share, at least 10% of allocations, should be assigned by lottery to enable robust impact evaluation from the outset, replicating the original Dutch design that made causal inference possible. Innovation voucher schemes targeting consulting and process redesign, rather than equipment purchase, have shown positive productivity effects in evaluations of the Dutch SME programme and of the UK schemes piloted in the early 2010s, with the strongest gains where vouchers funded external expertise rather than internal hardware. The lesson for AI adoption is straightforward: the binding constraint at firm level is rarely the cost of the technology itself; it is the cost of the complementary organisational change — the management practices, the workflow redesign and the intangible capital that make a general-purpose technology productive. Field-experimental work on knowledge workers using state-of-the-art AI tools confirms that gains are large but uneven, and concentrated where tasks are well aligned with the technology — making the case for subsidising the firm-level capacity to identify and reorganise those tasks rather than the technology itself.[64][65]

4. A European AI Productivity Observatory. One of the most striking findings of the current literature is the sheer paucity of reliable data on AI's actual impact at the firm level. The NBER CEO survey, the EIB working paper, and the Eurostat adoption data all capture different dimensions of the phenomenon, but no single institution tracks AI's productivity effects systematically, longitudinally, and at granular sectoral resolution. Europe should create an AI Productivity Observatory, housed within the Joint Research Centre or Eurostat, with a mandate to conduct annual firm-level surveys, publish standardised productivity indicators, and provide the evidence base for policy adjustment. Without rigorous measurement, policy will continue to oscillate between hype and disillusion.[66][67]

5. Pro-Worker AI Incentives. The Acemoglu-Autor-Johnson framework demonstrates that market forces alone will not produce pro-worker AI: the dominant business models of the major AI vendors are structurally aligned with automation and data capture, not with worker augmentation. Europe should create a Pro-Worker AI Label and Tax Incentive, modelled on existing eco-certification schemes, that provides a 25% R&D tax credit for firms developing AI systems demonstrably designed to augment, rather than replace, human capabilities. Eligibility criteria, defined by the AI Productivity Observatory, should include evidence of domain-specific design, human-in-the-loop architecture, and measurable skill-development outcomes. Complementary procurement preferences should require that EU-funded AI deployments in healthcare, education, and public administration meet pro-worker design standards.[68][69]

6. Convergence Funding for Lagging Member States. The tenfold gap in AI adoption rates across EU member states, from Denmark's 42% to Romania's 5%, risks creating a permanent two-speed digital economy within the Single Market. A dedicated AI Convergence Fund, drawn from cohesion policy and the Recovery and Resilience Facility, should finance AI infrastructure, training, and advisory services in the twelve member states with adoption rates below 15%. The fund should target industries where AI adoption could most rapidly narrow productivity gaps: agriculture, manufacturing, and public services.[70][71]

7. Transition Insurance: Universal Basic Capital and Wage Insurance. David Autor has proposed two complementary instruments — building on his work with Neil Thompson on the erosion of expertise — to insure against the risk that AI permanently undermines labour scarcity. Universal Basic Capital would grant every citizen a meaningful ownership stake in productive assets — in effect, a stock-market portfolio at birth — creating permanent stakeholders in the automated economy and providing income through capital returns rather than through ongoing political redistribution. Wage Insurance would subsidise part of the wage gap — say, 50% for a transitional period — for workers displaced into lower-paying roles, keeping them in the labour force rather than relying on benefits. Wage insurance has historical precedent on both sides of the Atlantic: the US Trade Adjustment Assistance program — created in 1962 and extended in 2002 with a dedicated wage-insurance component (Alternative TAA / Reemployment TAA) — has provided wage supplements to older displaced workers, and randomised evaluations of similar schemes have found that wage insurance increases re-employment speed and reduces long-term earnings losses for older displaced workers. European short-time work schemes — most prominently Germany's Kurzarbeit during the 2008 and 2020 crises — demonstrate that publicly co-funded wage support can preserve employment relationships at scale during structural transitions. A European adaptation, funded through a modest levy on AI inference compute, would provide a structural safety net calibrated to the distinctive nature of AI-driven displacement: not cyclical job loss, but permanent repricing of human expertise.[72][73]

8. Build a European Diffusion and Reallocation Agenda. Long-run productivity research by Bergeaud, Cette, and Lecat shows that the diffusion of general-purpose technologies across firms and sectors has historically mattered as much as the initial invention. Aghion, Bergeaud *et al.* further demonstrate that aggregate productivity gains from innovation depend critically on the economy's capacity to reallocate resources — capital, labour, and market share — toward more productive firms, and that distorted credit markets can block this process. Europe should launch a Diffusion and Reallocation Agenda that combines three instruments: first, a Technology Diffusion Service modelled on the US Manufacturing Extension Partnership, providing hands-on AI integration support to mid-cap and small firms; second, reallocation-friendly reforms to insolvency frameworks and labour mobility that allow capital and workers to move more quickly from stagnant to dynamic firms; and third, a review of credit allocation mechanisms — including public development banks and guarantee schemes — to ensure they do not inadvertently sustain low-productivity incumbents at the expense of innovative entrants.[74][75]

Conclusion: From Paradox to Policy

The AI productivity paradox is real, but it is not mysterious. Every major general-purpose technology in modern economic history — electricity, the internal combustion engine, the personal computer — exhibited the same pattern: transformative micro-level effects, followed by years or decades of macro-level invisibility, followed eventually by a measurable surge in aggregate productivity. The question for policy-makers is not whether AI will boost productivity — the experimental evidence leaves little doubt that it can — but whether Europe will create the conditions for the transition to happen quickly enough, and broadly enough, to matter.[76][77]

The current evidence suggests that, left to market forces, the benefits of AI will concentrate among large firms, high-skilled workers, and advanced economies — widening, rather than narrowing, the productivity disparities that already characterise the European landscape. Acemoglu, Autor, and Johnson have shown that the default trajectory of AI development favours automation over augmentation; the Harvard-BCG field experiments demonstrate that identical tools produce radically different outcomes depending on how workers and organisations choose to use them. The eight recommendations set out in this brief — a skills guarantee, an SME adoption facility, organisational transformation support, pro-worker AI incentives, a productivity observatory, convergence funding, transition insurance, and a diffusion and reallocation agenda — are designed to ensure that AI's demonstrated micro-level potential translates into shared, measurable, economy-wide productivity growth.[78][79][80][81]

> *The paradox will resolve itself. The policy question is whether it resolves in a way that lifts all firms, all workers, and all regions — or only those that were already ahead. Europe's industrial history shows that it knows how to diffuse transformative technologies broadly when it chooses to invest in the complementary institutions. The window for that choice is now.*[82]

Annex. Key Data on AI Adoption and Productivity

Indicator	Finding	Source
Task-level productivity gain (customer service)	+14% avg; +35% novice	Brynjolfsson et al. (2025)
Task-level quality gain (consulting)	+40% quality; +25% speed	Dell'Acqua et al. (2023)
Time saved on writing tasks	−40% time; +18% quality	Noy & Zhang (2023)
CEOs reporting no productivity change	89% of ~6,000 executives	NBER CEO Survey (2026)
US workers using generative AI (Dec 2025)	35.9%	Hartley et al. (2026)
EU enterprise AI adoption (2025)	20.0% (Eurostat) / 20.2% (OECD)	Eurostat; OECD
EU AI adoption change (2023–2025)	8.7% → 20.0% (+130%)	Eurostat; OECD
Large firms vs small firms (EU, 2025)	52.0% vs 17.4%	OECD (2026)
Denmark vs Romania (2025)	42.0% vs 5.2%	Eurostat (2025)
EU productivity gain from AI (5-year)	~1.1% cumulative (IMF)	IMF WP/25/067
EU productivity gain (at US adoption rates)	+4% labour productivity	EIB WP 2026/02
Acemoglu TFP estimate (10-year)	+0.53%	NBER WP 32487 (2024)
McKinsey gen. AI value estimate	$2.6–4.4 trillion / year	MGI (2023)
Time consumed verifying AI output	37–40% of time saved	Workday (2026)
Aghion aggregate productivity gain	+0.8–1.3pp / year	Aghion & Bunel (2024)
Frey & Osborne automation risk (2013)	~47% of US jobs at high risk	Frey & Osborne (2017)
EU internal services trade barriers	~110% (tariff equivalent)	IMF / Frey (Davos 2026)
Global AI perception (81k survey)	32% cite productivity; 22% cite disruption	Anthropic (2024)
AI labs: employees per $bn valuation	7–8 (OpenAI) vs 2,200 (Walmart)	Korinek / NYT (2026)
Citrini 2028 scenario: S&P drawdown	−38% from Oct 2026 highs	CitriniResearch (2026)

Sources: Brynjolfsson et al. (2025); Dell'Acqua et al. (2023); Noy & Zhang (2023); NBER (2026); OECD (2025–2026); Eurostat (2025); IMF (2025); EIB (2026); McKinsey MGI (2023); Workday (2026).

Notes

[1] James van Geelen and Alap Shah, “The 2028 Global Intelligence Crisis: A Thought Exercise in Financial History, from the Future”, CitriniResearch Macro Memo, 22 February 2026. https://www.citriniresearch.com/p/2028gic
[2] Robert Solow, "We'd better watch out", New York Times Book Review, 12 July 1987.
[3] Jonathan S. Hartley, Filip Jolevski, Vitor Melo, and Brendan Moore, "AI, Productivity, and Labor Markets: A Review of the Empirical Evidence", NBER/ICLE, 2026. 35.9% of US workers used generative AI by December 2025. https://laweconcenter.org/resources/ai-productivity-and-labor-markets-a-review-of-the-empirical-evidence/
[4] Yotzov, Barrero, Bloom, Bunn, Davis, Foster, Jalca, Meyer, Mizen, Navarrete, Smietanka, Thwaites and Wang, “Firm Data on AI”, NBER Working Paper No. 34836, February 2026. https://www.nber.org/papers/w34836 — see also press coverage in Fortune, “Thousands of CEOs just admitted AI had no impact on employment or productivity”, 17 February 2026.
[5] PwC, 29th Global CEO Survey, January 2026. Only 30% of CEOs reported any revenue increase from AI in the preceding 12 months; 22% reported increased costs. https://www.pwc.com/gx/en/issues/c-suite-insights/ceo-survey.html
[6] Torsten Slok, Apollo Chief Economist: "AI is everywhere except in the incoming macroeconomic data", Apollo Academy, 2025. https://www.apolloacademy.com/waiting-for-the-ai-j-curve/
[7] Federal Reserve Bank of St. Louis, On the Economy blog / FRED data, observed cumulative excess productivity growth of approximately 1.9% since the introduction of ChatGPT in late 2022. Alexander Bick, Adam Blandin, and David Deming, “The State of Generative AI Adoption in 2025”, Federal Reserve Bank of St. Louis, On the Economy, November 2025. https://www.stlouisfed.org/on-the-economy/2025/nov/state-generative-ai-adoption-2025. “Cumulative excess productivity growth” denotes the total gap, in percentage points, between observed US nonfarm labour productivity and the level implied by extrapolating the 2015–2019 pre-pandemic trend, measured from Q4 2022 (the public release of ChatGPT) through Q2 2025 — not an annualised rate.
[8] David Autor, Anton Korinek, and Natasha Sarin, “A.I. and Jobs: What Do Economists Actually Know?”, The New York Times (Opinion), February 2026. Moderated by David Leonhardt.
[9] Daron Acemoglu, "The Simple Macroeconomics of AI", NBER Working Paper No. 32487, April 2024. Estimates AI will increase total factor productivity by 0.53% over ten years.
[10] Goldman Sachs, "Generative AI Could Raise Global GDP by 7%", March 2023; Goldman Sachs Research, "The Potentially Large Effects of AI on Economic Growth", Briggs and Kodnani, 2023. https://www.goldmansachs.com/insights/articles/generative-ai-could-raise-global-gdp-by-7-percent
[11] McKinsey Global Institute, The Economic Potential of Generative AI: The Next Productivity Frontier, June 2023. https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier
[12] Philippe Aghion and Simon Bunel, "AI and Growth: Where Do We Stand?", Federal Reserve Bank of San Francisco Conference Paper, June 2024. https://www.frbsf.org/wp-content/uploads/AI-and-Growth-Aghion-Bunel.pdf
[13] Philippe Aghion, Simon Bunel, and Xavier Jaravel, "What AI Means for Growth and Jobs", Project Syndicate, 14 January 2025. https://www.project-syndicate.org/commentary/ai-will-boost-productivity-growth-without-harming-jobs-employment-by-philippe-aghion-et-al-2025-01
[14] Goldman Sachs Research, AI investment contribution to US GDP growth, 2025–2026 series. Comments by Jan Hatzius, Goldman Sachs Chief Economist, that AI investment contributed “basically zero” to US GDP growth in 2025; as also reported in Fortune, “Goldman finds no relationship between AI and productivity”, 3 March 2026.
[15] Federal Reserve Bank of San Francisco, "The AI Moment? Possibilities, Productivity, and Policy", Economic Letter 2026-02, February 2026. https://www.frbsf.org/research-and-insights/publications/economic-letter/2026/02/ai-moment-possibilities-productivity-policy/
[16] Filippo Bontadini, Carol Corrado, Jonathan Haskel, and Cecilia Jona-Lasinio, "AI as an Innovation in the Method of Innovation: Implications for Productivity Growth in the US and Europe", October 2025. Software/AI contributed 50% of the 2% average US labour productivity growth (2017–2024); Europe likely +0.3pp vs US +1pp.
[17] Erik Brynjolfsson, Danielle Li, and Lindsey Raymond, "Generative AI at Work", The Quarterly Journal of Economics, Vol. 140, No. 2, May 2025, pp. 889–942. https://academic.oup.com/qje/article/140/2/889/7990658
[18] Fabrizio Dell'Acqua et al., "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality", Harvard Business School Working Paper 24-013, 2023. https://www.hbs.edu/faculty/Pages/item.aspx?num=64700
[19] Shakked Noy and Whitney Zhang, "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence", Science, Vol. 381, No. 6654, 2023, pp. 187–192. https://www.science.org/doi/10.1126/science.adh2586
[20] California Management Review, "Seven Myths about AI and Productivity: What the Evidence Really Says", Berkeley, October 2025. https://cmr.berkeley.edu/2025/10/seven-myths-about-ai-and-productivity-what-the-evidence-really-says/

[21] Timothy Bresnahan and Manuel Trajtenberg, "General Purpose Technologies: Engines of Growth?", Journal of Econometrics, Vol. 65, No. 1, 1995, pp. 83–108.
[22] Erik Brynjolfsson, Daniel Rock, and Chad Syverson, "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies", American Economic Journal: Macroeconomics, Vol. 13, No. 1, 2021, pp. 333–372.
[23] Paul David, "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox", American Economic Review, Vol. 80, No. 2, 1990, pp. 355–361.
[24] US Bureau of Labor Statistics: non-farm labour productivity growth averaged 2.7% per year 1995–2005, compared with 1.4% per year 1973–1995.
[25] Workday, "Beyond Productivity: Measuring the Real Value of AI", January 2026. 37% of time saved by AI is consumed by reviewing, correcting, and verifying AI-generated output.
https://newsroom.workday.com/2026-01-14-New-Workday-Research-Companies-Are-Leaving-AI-Gains-on-the-Table
[26] Carl Benedikt Frey, How Progress Ends: Technology, Innovation, and the Fate of Nations, Princeton University Press, September 2025. Winner of the 2026 PROSE Award in Economics; shortlisted for the 2025 FT/Schroders Business Book of the Year.
https://press.princeton.edu/books/hardcover/9780691233079/how-progress-ends
[27] Carl Benedikt Frey, interview at WEF Annual Meeting, Davos, January 2026 (Radio Davos podcast).
https://www.weforum.org/podcasts/radio-davos/
[28] Carl Benedikt Frey and Michael Osborne, "The Future of Employment: How Susceptible Are Jobs to Computerisation?", Technological Forecasting and Social Change, Vol. 114, 2017, pp. 254–280. Original 2013 working paper estimated ~47% of US jobs at high risk of automation.
[29] Daron Acemoglu, David Autor, and Simon Johnson, "Building Pro-Worker Artificial Intelligence", Hamilton Project at Brookings, February 2026. https://www.hamiltonproject.org/publication/paper/building-pro-worker-ai/
[30] Daron Acemoglu and Simon Johnson, Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity, PublicAffairs, 2023. https://shapingwork.mit.edu/power-and-progress/
[31] Daron Acemoglu, “Pro-worker AI doesn’t just happen. Companies need to act”, MIT Sloan Ideas Made to Matter, 2025; and Daron Acemoglu, “The World Needs a Pro-Human AI Agenda”, Project Syndicate, November 2024. See also Acemoglu, Autor and Johnson, “Building Pro-Worker Artificial Intelligence”, NBER Working Paper No. 34854, February 2026.
[32] MIT Sloan, "Pro-Worker AI Doesn’t Just Happen. Companies Need to Act", 2026.
https://mitsloan.mit.edu/ideas-made-to-matter/pro-worker-ai-doesnt-just-happen-companies-need-to-act
[33] Steven Randazzo, Hila Lifshitz, Katherine C. Kellogg, Fabrizio Dell’Acqua, Ethan Mollick, François Candelon, and Karim R. Lakhani, "Cyborgs, Centaurs and Self-Automators: The Three Modes of Human-GenAI Knowledge Work", Harvard Business School Working Paper 26-036, 2025.
[34] Matt Sigelman, Benjamin Francis, Shrinidhi Rao, and Gwynn Guilford, "Beyond the Binary: How Automation and Augmentation Are Combining to Reshape Work", Burning Glass Institute, January 2026.
[35] Aruna Ranganathan and Xingqi Maggie Ye, "AI Doesn’t Reduce Work — It Intensifies It", Harvard Business Review, February 2026.
https://hbr.org/2026/02/ai-doesnt-reduce-work-it-intensifies-it
[36]Lareina Yee and Anu Madgavkar (McKinsey Global Institute), "How Workers Will Adapt in the AI Era", TIME, 17 December 2025.
[37] World Economic Forum, "Four Futures for Jobs in the New Economy: AI and Talent in 2030", January 2026.
[38] Anthropic, "What 81,000 People Want from AI", December 2024. 80,508 interviews across 159 countries in 70 languages. 32% cite productivity gains; 22.3% cite economic disruption as top concern. https://www.anthropic.com/81k-interviews
[39] Bontadini et al., op. cit. (note 16).
[40] Frey, How Progress Ends, op. cit. (note 26).
[41] Frey, Davos 2026 remarks, op. cit. (note 27).
[42] Autor, Korinek & Sarin, op. cit. (note 8).
[43] Philippe Aghion, Antonin Bergeaud, Gilbert Cette, Rémy Lecat, and Hélène Maghin, “The Inverted-U Relationship Between Credit Access and Productivity Growth”, Economica, Vol. 86, No. 341, January 2019, pp. 1–31. Also Banque de France Working Paper No. 617, 2018.
[44] Eurostat, "20% of EU enterprises use AI technologies", News Release, 11 December 2025.
https://ec.europa.eu/eurostat/web/products-eurostat-news/w/ddn-20251211-2
[45] OECD, "AI use by individuals surges across the OECD as adoption by firms continues to expand", January 2026. 20.2% of firms used AI in 2025, up from 8.7% in 2023. 52% of large firms vs 17.4% of small firms.
https://www.oecd.org/en/about/news/announcements/2026/01/ai-use-by-individuals-surges-across-the-oecd-as-adoption-by-firms-continues-to-expand.html
[46] Eurostat, op. cit. (note 44).
[47] OECD, op. cit. (note 45).
[48] IMF Working Paper WP/25/067, "Artificial Intelligence and Productivity in Europe", April 2025.
https://www.imf.org/en/publications/wp/issues/2025/04/04/ai-and-productivity-in-europe-565924

[49] EIB Working Paper 2026/02, "AI Adoption, Productivity and Employment: Evidence from European Firms". AI adoption increases labour productivity by 4% on average in the EU, with no evidence of reduced employment in the short run.
https://www.eib.org/en/publications/20250383-economics-working-paper-2026-02
[50] IMF Blog, "How Europe Can Capture the AI Growth Dividend", 20 November 2025.
https://www.imf.org/en/blogs/articles/2025/11/20/how-europe-can-capture-the-ai-growth-dividend
[51] Carl Benedikt Frey, remarks at the World Economic Forum, Davos 2026, citing IMF estimates that intra-EU barriers to trade in services amount to approximately 110% — described by Frey as the equivalent of "Trump Liberation Day tariffs self-imposed on services inside the European Union." Primary IMF source to be added.
[52] SkillsFuture Singapore Agency (SSG), "SkillsFuture Year-in-Review 2023", published 2024, consolidating training participation, learners and SkillsFuture Credit activation data; and Ministry of Manpower / SSG manpower research, "Training participation rate of residents in the labour force" series. https://www.skillsfuture.gov.sg/aboutskillsfuture
[53] DARES, Le compte personnel de formation, annual series (DARES Analyses / DARES Résultats), Ministère du Travail. Source for CPF entitlement amounts, the over-representation of low-skilled workers among users, and the post-2015 evolution of training participation among that group.
[54] Cour des Comptes, La formation professionnelle des salariés, Rapport public thématique, 2023; Décret n° 2024-394 du 29 avril 2024 relatif à la participation du titulaire d'un compte personnel de formation au financement d'une formation (legal basis for the €100 co-payment).
[55] OECD, Individual Learning Accounts: Design is Key for Success, Policy Brief on the Future of Work, OECD Publishing, Paris, 2019. Comparative synthesis of individual learning accounts (SkillsFuture, CPF and others) and of the conditions for effectiveness, including portability and employer co-investment.
[56] Council of the European Union, Council Recommendation of 16 June 2022 on individual learning accounts, OJ C 243, 27.6.2022, p. 26.
[57] OECD, "Making AI Work: Why Investing in Skills Matters", January 2026.
https://www.oecd.org/en/blogs/2026/01/making-ai-work-why-investing-in-skills-matters.html
[58] Brynjolfsson, Li & Raymond, op. cit. (note 17).
[59] Gina Raimondo, "America Cannot Withstand the Economic Shock That's Coming", The New York Times (Opinion), 6 March 2026. Proposes a "grand bargain" between government, industry, and workers on AI-era skills, citing the CHIPS Act/TSMC/Maricopa Community College partnership as proof of concept.
[60] OECD, AI Adoption by Small and Medium-Sized Enterprises, December 2025.
https://www.oecd.org/en/publications/ai-adoption-by-small-and-medium-sized-enterprises_426399c1-en.html
[61] G7, "SME AI Adoption Blueprint", Industry, Digital and Technology Ministerial Statement, 2025.
https://www.g7.utoronto.ca/ict/2025-sme-ai-adoption-blueprint.html
[62] Maarten Cornet, Björn Vroomen, and Marc van der Steeg, "Do innovation vouchers help SMEs to cross the bridge towards science?", CPB Discussion Paper No. 58, CPB Netherlands Bureau for Economic Policy Analysis, 2006.
[63] See Albert Bravo-Biosca, "Experimental innovation policy", in Innovation Growth Lab / Nesta working paper series, 2019, and follow-on register-based analyses linking Dutch voucher recipients to firm-level performance data at Statistics Netherlands.
[64] Brynjolfsson, Rock & Syverson, op. cit. (note 22).
[65] Dell'Acqua et al., op. cit. (note 18).
[66] National Bureau of Economic Research, CEO Survey on AI and Productivity, February 2026 — as reported in: Fortune, "Thousands of CEOs just admitted AI had no impact on employment or productivity", 17 February 2026.
https://fortune.com/2026/02/17/ai-productivity-paradox-ceo-study-robert-solow-information-technology-age/
[67] IMF WP/25/067, op. cit. (note 48).
[68] Acemoglu, Autor & Johnson, op. cit. (note 29).
[69]Ajay K. Agrawal, John McHale, and Alexander Oettl, "Enhancing Worker Productivity Without Automating Tasks: A Different Approach to AI and the Task-Based Model", NBER Working Paper No. 34781, 2026. Argues augmentation pathways offer fuller accounting of AI's economic effects than standard task-substitution frameworks. https://www.nber.org/papers/w34781
[70] Eurostat, op. cit. (note 44).
[71] Mario Draghi, The Future of European Competitiveness, European Commission, September 2024.
https://commission.europa.eu/topics/competitiveness/draghi-report_en
[72] *Diagnosis*: David Autor and Neil Thompson, "Expertise", NBER Working Paper No. 33941, June 2025. *Policy proposals*: David Autor, contribution on Universal Basic Capital and Wage Insurance in *The Digitalist Papers*, Vol. 2, Stanford Digital Economy Lab, 2025.
[73] Autor, Korinek & Sarin, op. cit. (note 8).
[74]Antonin Bergeaud, Gilbert Cette, and Rémy Lecat, "Productivity Trends in Advanced Countries between 1890 and 2012", Review of Income and Wealth, Vol. 62, No. 3, September 2016, pp. 420–444.
[75] Philippe Aghion, Antonin Bergeaud, Timo Boppart, Peter J. Klenow, and Huiyu Li, "Missing Growth from Creative Destruction", American Economic Review, Vol. 109, No. 8, August 2019, pp. 2795–2822.
[76] Brynjolfsson, Rock & Syverson, op. cit. (note 22).
[77] Paul David, op. cit. (note 23).

[78] McKinsey Global Institute, Accelerating Europe: Competitiveness for a New Era, 2024.
https://www.mckinsey.com/mgi/our-research/accelerating-europe-competitiveness-for-a-new-era
[79] Éric Hazan, "IA : il faut accélérer la politique industrielle", Les Echos, 2024
(https://www.lesechos.fr/idees-debats/cercle/opinion-ia-il-faut-accelerer-la-politique-industrielle-2080326); Éric Hazan, "L'Europe face à l'âge de l'intelligence", Les Echos, 2024
(https://www.lesechos.fr/idees-debats/cercle/opinion-leurope-face-a-lage-de-lintelligence-une-strategie-technologique-pour-eviter-le-declassement-2147934).
[80] Acemoglu, Autor & Johnson, op. cit. (note 29).
[81] Randazzo et al., op. cit. (note 33).
[82] WEF, "AI Paradoxes: Why AI's Future Isn't Straightforward", December 2025.
https://www.weforum.org/stories/2025/12/ai-paradoxes-in-2026/