International
Center for
Law & Economics

AI, Productivity, and Labor Markets: A Review of the Empirical Evidence

Eric Fruits & Kristian Stout

ICLE Issue Brief 2026-02-05

Executive Summary

Generative artificial intelligence (AI) has diffused with unusual speed since late 2022. By late 2024, nearly 40% of U.S. adults ages 18–64 reported using AI tools, a pace that exceeds comparable stages for personal computers and the internet. That rapid uptake has sharpened two policy questions: whether AI will generate measurable gains in output and productivity at the aggregate level, and whether the adjustment process will produce labor-market disruption large enough to justify new regulatory intervention.

The empirical literature points to a clear pattern. Controlled field experiments and randomized trials document large productivity gains at the task and firm level, often alongside quality improvements. Across writing, customer support, software development, accounting, law, and translation, studies report 15% to more than 50% reductions in task-completion time, meaningful quality gains, and disproportionately large benefits for less-experienced workers, producing "skill compression" within occupations.

At the same time, aggregate labor-market indicators through 2024–2025 show limited disruption, despite rapid adoption. Most datasets find little evidence of economywide job loss or wage decline. Where effects appear, they are concentrated in entry-level segments of highly exposed occupations, while senior employment remains largely stable. Adjustment to date has occurred through task reallocation and within-firm productivity gains, rather than mass displacement.

Macroeconomic projections remain uncertain. Credible estimates range from modest productivity gains to large output increases, depending on assumptions about task coverage, diffusion speed, organizational redesign, and complementary investment. Disagreement in the macro literature reflects divergent assumptions and measurement challenges, including the well-documented productivity J-curve, rather than conflicting data.

The evidence also informs competition policy. AI lowers minimum viable scale in many downstream markets, facilitating entry by startups and small firms, even as upstream model development may remain concentrated. Open-source components and contracting already provide broad access to AI inputs, weakening claims that mandatory access or data-sharing rules are necessary to preserve competition.

Overall, the literature offers limited support for regulatory approaches premised on widespread worker displacement, durable monopoly power, or exclusionary "data moats." A more defensible approach emphasizes targeted enforcement of existing law, reduced regulatory fragmentation, and investment in complements such as skills, governance, and infrastructure.

The issue brief concludes with an extensive annotated bibliography summarizing the empirical and theoretical studies underlying these findings.

I. Macroeconomic Effects of AI: Evidence, Assumptions, and Uncertainty

The macroeconomic literature on AI remains driven largely by models, rather than realized output data. That imbalance is partly mechanical. Large-scale diffusion began only recently, and firm-level reorganization and supply-chain adjustment typically lag initial adoption.

Two additional issues shape how existing macro evidence should be interpreted: measurement and complements.

Diane Coyle and John Lourenze S. Poquiz (2025) argue that standard national accounts may understate AI's economic contribution when gains appear as quality improvements, time savings, or outputs that markets price poorly. When workers produce higher-quality text, code, or designs in less time, welfare gains at the firm level can exceed what revenue-based aggregates capture. Measurement limits do not imply the absence of output gains, but they weaken inferences drawn from a short window of aggregate data.

Erik Brynjolfsson, Daniel Rock, and Chad Syverson's (2021) "productivity J-curve" highlights a second interpretive challenge. General-purpose technologies can raise long-run productivity even as measured productivity initially stagnates or falls.[1] Early adoption requires firms to invest in organizational redesign, worker training, and process change before efficiency gains appear in output statistics. The J-curve implies a transition period in which complementary investments determine whether productivity gains ultimately materialize.

Observed macroeconomic flatness therefore supports competing interpretations. One view holds that AI delivers modest aggregate gains even after diffusion, either because the affected task share of GDP is limited or because current capabilities remain poorly matched to high-value work. Another view holds that the economy remains in an adjustment phase, with organizational and measurement frictions delaying visible gains. The literature does not yet resolve which interpretation dominates, but it provides a structured framework for evaluating for future research.

A. Assumption-Driven Divergence in Macroeconomic Forecasts of AI

The sharpest disagreements in the macroeconomic literature turn on three margins: task share, adoption, and complementary investment.

Daron Acemoglu (2025) develops a task-based macro model that distinguishes "easy-to-learn" tasks with objective verification from "hard-to-learn" tasks requiring context-sensitive judgment. He

[1] When a general-purpose technology—such as electricity, computers, or artificial intelligence—emerges, firms must invest in intangible assets, including process redesign, worker training, and new software and business models. These investments absorb time and capital without generating immediate, measurable output, often producing an initial slowdown in recorded productivity. As complementary investments mature and integration deepens, productivity growth accelerates, forming the upward stem of the J-curve.

applies Hulten's Theorem to bound aggregate gains by the GDP share of tasks that AI can perform.[2] Under his parameterization, AI raises total factor productivity by less than 0.66% over 10 years. The estimate is deliberately conservative and mechanism-driven: when AI excels only on a subset of tasks and struggles where verification is costly or ambiguous, aggregate gains remain limited even with broad diffusion.

Alex Arnon and Kent Smetters (2025), writing for the Penn Wharton Budget Model, offer a middle-ground estimate grounded in exposure classifications, adoption assumptions, and micro-level evidence. They conclude that roughly 10% of current U.S. GDP could be affected in the short run, rising to 15% over two decades under partial adoption. Their calculations assume average labor-cost savings of about 25% from current tools, potentially increasing to 40% as systems improve. The results hinge on adoption speed and on whether measured cost reductions translate into higher output rather than rent reallocation.

More optimistic projections assume faster diffusion and stronger downstream effects. Joseph Briggs and Devesh Kodnani (2023) estimate a 7% increase in annual global GDP over 10 years—roughly $7 trillion—based on broad occupational task exposure. Erkan Erdem and Dileep Birur (2025), using a dynamic computable general equilibrium framework, estimate that rapid adoption could raise U.S. GDP by about $2.48 trillion by 2030.[3] The Congressional Research Service synthesizes these projections and notes that estimates are typically positive but highly sensitive to diffusion timelines that may unfold over decades, as with earlier general-purpose technologies (Lida Weinstock and Paul Tierno, 2025).

Taken together, these forecasts differ less over data than over assumptions. Aggregate effects depend on three factors: diffusion across sectors, whether productivity gains expand output rather than reshuffle rents, and whether organizational redesign occurs at scale. The wide range of estimates reflects divergence along these margins rather than measurement error alone.

Isabel Aldasoro, Sebastian Doerr, Leonardo Gambacorta, and Daniel Rees (2024) underscore this point by modeling AI adoption as a positive productivity shock. They show that aggregate output, consumption, and investment rise in both the short and long run, while inflation dynamics depend on expectations. When households and firms anticipate productivity gains, demand can increase early and push inflation upward. When gains arrive unanticipated, supply may initially outpace demand, producing disinflation.

[2] Hulten's Theorem holds that the aggregate economic effect of a productivity change in a given industry equals that industry's share of national GDP, measured by total sales. The result implies that an activity's economywide importance depends on its relative size, not on the complexity of its upstream or downstream linkages.

[3] A computable general-equilibrium (CGE) model simulates how shocks—such as new technologies, taxes, or trade policies—propagate across an economy. Using observed economic data, the model estimates how producers and consumers adjust until the system converges to a new equilibrium.

For policy, the implication is straightforward: short-run stabilization outcomes during AI adoption can diverge sharply from long-run productivity effects. Regulators should not infer long-run welfare consequences from near-term inflation responses alone.

B. Plausible AI-Driven Geographic Divergence, with Uncertain Causal Attribution

The Council of Economic Advisers (2026) argues that AI could generate an international growth divergence analogous to earlier industrial transitions, with the United States positioned to benefit from advantages in investment, compute capacity, and innovation ecosystems. The report points to stronger growth among so-called “Pax Silica” economies and the United States’ large share of global compute resources. It reflects an official policy position favoring rapid deployment and diffusion.

The empirical case for attributing recent growth differentials to AI remains limited. Short observation windows and confounding macroeconomic forces complicate efforts to isolate AI’s contribution from broader cyclical, fiscal, and geopolitical factors.

Shahid Yusuf’s (2025) United Nations Development Programme analysis similarly highlights the risk of divergence but shifts the focus to distributional effects. It emphasizes uneven gains within countries and across regions, particularly in Asia-Pacific economies facing long-run productivity slowdowns.

II. Labor-Market Effects of AI: Evidence of Adjustment Rather Than Displacement

Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock (2024) provide a foundational mapping between occupational tasks and large language model capabilities. They estimate that roughly 80% of U.S. workers have at least 10% of their tasks exposed to LLM assistance, and about 19% have 50% or more of tasks exposed. They stress that “exposure” measures technical feasibility for task assistance, not a prediction of job destruction. That distinction matters, because task reallocation and mass job elimination carry fundamentally different welfare implications.

Xianguo Huang (2025) extends the exposure framework across countries and demographic groups. He finds higher exposure in higher-income economies and disproportionate exposure among women and mid-education workers in some contexts. Consistent with the task-based view, the analysis characterizes current AI systems as primarily augmenting human labor, rather than automating it outright.

A. Limited Aggregate Labor-Market Disruption Through 2024–2025

Several studies report null or modest aggregate labor-market effects despite rapid AI adoption. The Budget Lab at Yale finds no clear relationship between AI exposure and unemployment through August 2025 (Martha Gimbel *et al.*, 2025). Anders Humlum and Emilie Vestergaard (2025) link survey-reported ChatGPT use to Danish administrative records across 11 exposed occupations and find essentially zero effects on earnings or hours through 2024.

U.S.-based evidence points in the same direction. Jonathan S. Hartley, Filip Jolevski, Vitor Melo, and Brendan Moore (2026) report that 35.9% of U.S. workers used generative AI by December 2025 and find small positive wage effects, with no statistically significant declines in job openings or employment in exposed occupations. Bharat Chandar's (2025) Current Population Survey analysis similarly finds no aggregate employment decline, while documenting heterogeneity across education levels and occupations.

Taken together, these studies find no evidence of immediate economywide labor displacement through 2024–2025. The results instead point to early adjustment through task reallocation, quality improvement, and within-firm productivity gains, rather than broad-based layoffs.

B. Entry-Level Effects in Some AI-Exposed Occupations

While aggregate indicators remain stable, several studies identify concentrated entry-level effects in highly exposed segments. Erik Brynjolfsson, Bharat Chandar, and Ruyu Chen (2025), using ADP payroll data, estimate that workers ages 22–25 in highly exposed occupations experienced employment declines of roughly 16% relative to trend following ChatGPT's release, while senior employment remained stable. Bouke Klein Teeselink (2025), examining the United Kingdom, finds that exposed firms reduced employment and hiring concentrated among junior and entry-level roles and posted lower advertised salaries for exposed occupations.

These findings point to a plausible adjustment channel. AI can automate discrete tasks that have traditionally served as entry-level work, reducing marginal demand for junior labor without displacing more senior workers.

The evidence suggests structural adjustment in career ladders and human-capital formation, rather than broad job loss. If entry-level pathways narrow, labor markets may respond by shifting screening and training mechanisms or reallocating junior workers toward tasks that remain complementary to AI. The literature has not yet resolved whether these effects persist or fade as firms adapt.

C. Complementarity and Skill Compression in Many AI-Exposed Settings

Andrew C. Johnston and Christos A. Makridis (2025) find that higher-exposure sectors experienced wage and employment gains, particularly among younger and more educated workers, while roles characterized by direct substitution saw declines. The pattern points to task-level complementarity, rather than economywide labor replacement.

Evidence from controlled experiments shows that realized productivity gains also depend on workers' ability to judge when AI assistance improves outcomes. Andrew Caplin *et al.* (2024), studying an age-classification task, find that AI raises performance across ability levels but reduces performance dispersion most when users are well calibrated about their own skills. Low-ability participants who accurately recognized their limitations achieved the largest gains by relying selectively on AI, while overconfident or underconfident users underperformed relative to calibrated peers. A counterfactual exercise suggests that universal calibration would nearly double AI's inequality-reducing effects.

Across settings, micro-level evidence points to skill compression. AI tools disproportionately boost output among lower-performing workers, narrowing performance gaps within job categories. This pattern recurs across productivity studies and carries distributional implications: AI can reduce inequality within occupations, even as it reshapes demand for certain entry-level roles.

III. Within-Firm Productivity Gains and Their Limits

The micro-level evidence provides the strongest empirical support in the current literature. Shakked Noy and Whitney Zhang (2023) conduct a randomized experiment with 453 professionals and find that ChatGPT use reduces task-completion time by roughly 40% and increases output quality by about 18%, with larger gains concentrated among initially lower-performing workers. Erik Brynjolfsson, Danielle Li, and Lindsey Raymond (2025) examine a Fortune 500 customer-support deployment and find a 15% average productivity increase, measured as issues resolved per hour, and a 36% increase for workers in the bottom skill quintile. Customer satisfaction remained stable, while attrition among newer workers declined sharply. Together, these results show that AI can transmit best practices and raise performance at the lower end of the skill distribution.

Field evidence from professional services points in the same direction. Jung Ho Choi and Chloe L. Xie (2025) analyze accounting work at a technology firm and survey 277 accountants, finding that AI adoption correlates with an 18% increase in weekly client support and a reallocation of roughly 9% of work hours from routine data entry to higher-value tasks, such as client communication. AI use reduced monthly book-closing timelines by 7.5 days and increased ledger detail by 12%, indicating quality improvements alongside time savings. The authors also document complementarity between professional expertise and AI confidence scores: experienced accountants used model outputs to target review effort, rather than replace judgment. A framed experiment reveals occasional overreliance on inaccurate suggestions, underscoring the importance of verification protocols.

Evidence from legal practice reinforces these findings. Daniel Schwarcz *et al.* (2025) conduct a randomized controlled trial evaluating AI tools used by law students on complex legal tasks. Both a retrieval-augmented legal system and a general-purpose reasoning model substantially improved document quality across clarity, organization, and analytical depth. Students using AI completed assignments 50% to 130% faster than control groups. The study also identifies functional differentiation across tools: retrieval-augmented systems reduced citation errors, while reasoning models improved substantive analysis. These results suggest that limitations observed in earlier model generations are not fixed and that appropriate tool selection and workflow design can mitigate risks without restricting deployment.

Software development studies show similarly large effects. Sida Peng *et al.* (2023) report that GitHub Copilot users complete coding tasks 55.8% faster in controlled settings, with larger gains among less experienced developers. Kevin Zheyuan Cui *et al.* (2025), studying nearly 5,000 developers across three large field experiments, find a 26.08% increase in weekly task completion, driven by higher adoption and disproportionately larger gains for junior developers. These findings undermine claims that AI primarily benefits top performers and instead show AI reducing frictions for early-career workers.

Experimental evidence from translation extends the pattern. Ali Merali (2024) conducts a randomized trial with 300 professional translators and links increased training compute to economic outcomes. A tenfold increase in compute reduced completion time by 12.3%, improved quality by 0.18 standard deviations, and raised earnings per minute by 16.1%. Lower-skilled translators experienced gains roughly four times larger than their higher-skilled counterparts.

The magnitude and replication of these effects across contexts suggest that productivity gains arise consistently in specific task categories—including writing, customer support, software development, and translation—while remaining sensitive to task structure and workflow integration.

A. The Jagged Technological Frontier and Deployment Constraints

Fabrizio Dell'Acqua *et al.* (2023) describe a "jagged technological frontier," in which AI exhibits uneven capabilities across tasks that appear similar in difficulty. Within this frontier, AI can substantially improve productivity and quality for some complex tasks, while producing errors on others that seem straightforward. The unevenness requires knowledge workers to exercise judgment in deciding when AI assistance is appropriate and when human oversight remains essential.

In a Boston Consulting Group field experiment using GPT-4, Dell'Acqua *et al.* (2023) find that AI improved performance on tasks within its capability boundary but reduced performance on tasks just beyond it. The decline stemmed from overreliance on plausible but incorrect model outputs. The results show that productivity effects depend critically on task selection and verification systems, not on model capability alone.

This literature identifies an internal governance challenge for firms. Verification protocols, worker training, and clear task assignment can mitigate boundary failures more effectively than broad constraints on model development. A blanket regulatory approach that raises adoption costs across all use cases would fail to target the source of error and risk suppressing productivity gains where AI performs reliably.

B. Why Intra-Firm Productivity Gains Do Not Automatically Translate to Macro Growth

The micro-level evidence indicates substantial productivity potential, but its translation to macroeconomic outcomes depends on complementary investments. The J-curve framework predicts that firms incur adjustment costs—through reorganization, training, and process redesign—before realizing large productivity gains (Erik Brynjolfsson, Daniel Rock, and Chad Syverson, 2021).

Measurement frictions can further delay visibility in aggregate data. Diane Coyle and John Lourenze S. Poquiz (2025) show that standard GDP statistics often miss quality improvements and time savings, weakening inferences drawn from short-run macro indicators.

Diffusion also remains uneven. Some occupations exhibit large productivity effects in controlled experiments, while others show limited exposure or face high verification costs (Tyna Eloundou *et al.*, 2024; Daron Acemoglu, 2025). Together, these constraints explain why macroeconomic effects may lag or appear smaller than micro-level gains would suggest.

IV. Entrepreneurship and Business Formation in the AI Economy

Junhui Jeff Cai *et al.* (2025) analyze administrative business-formation data using a difference-in-differences design around ChatGPT's release. They find increased entry by first-time and resource-constrained founders and show that post-ChatGPT firms tend to have fewer shareholders and smaller founding teams. The mechanism is direct: AI substitutes for managerial, operational, and technical tasks that previously required additional hires or cofounders. By lowering the minimum viable team size, AI reduces entry costs, increases the number of entrants, and strengthens downstream competition. These effects arise independently of concentration at the foundation-model layer.

Survey evidence from the Organisation for Economic Co-operation and Development (2025b) reinforces this pattern. About 31% of small and medium-sized enterprises across seven countries had adopted generative AI by 2024, and 83% of adopters reported no change in staffing levels. Firms cited reduced workloads and improved performance and many described AI adoption as a response to labor shortages, rather than a labor-replacement strategy. Non-adopters most often cited lack of suitability (57%), legal or data-privacy concerns (54%), and insufficient internal skills (50%).

These findings carry two policy-relevant implications. First, uncertainty about compliance and legal risk remains a material barrier to adoption, suggesting that clearer rules and regulatory harmonization could accelerate diffusion. Second, workforce capability constrains uptake, pointing toward training and institutional learning, rather than command-and-control regulation aimed at model developers.

Capital-market evidence highlights a complementary margin. Iuri Struta (2024) reports more than $20 billion invested in generative AI startups through Q3 2024. Abu Bakkar Siddik, Yong Li, and Anna Min Du (2024), studying 556 generative AI startups, find that investor influence—measured by the number of investors, lead investors, and funding rounds—strongly predicts total funding, while measures of "technological influence," such as IT spending and patenting, do not. Access to capital networks and distribution channels remains central to entry and scaling, underscoring the importance of contract freedom and partnership formation for competitive AI markets.

V. Dynamic Competition and Market Structure in AI Markets

The Organisation for Economic Co-operation and Development's (2025a) analysis of competitive dynamics in downstream markets concludes that AI can lower entry barriers and the minimum efficient scale by automating functions that previously required large, specialized teams. This finding aligns with the entrepreneurship evidence. The result is a plausible decoupling: upstream model development may exhibit concentration driven by fixed costs in compute and training, while downstream markets experience increased entry, experimentation, and rivalry.

OECD competition roundtables also document mixed effects from vertical integration between model providers and application firms (OECD, 2025a). Some forms of integration may foreclose rivals, while others improve coordination, reduce transaction costs, and accelerate diffusion. Absent evidence of durable foreclosure, vertical contracting is generally presumptively efficiency-enhancing. Broad structural presumptions risk deterring procompetitive arrangements that facilitate deployment and innovation.

Cross-country evidence underscores the role of human capital, infrastructure, and trade openness. Alessandra Bonfiglioli *et al.* (2025) find that countries with larger STEM graduate pipelines, higher internet penetration, and greater export volumes hold comparative advantages in AI-intensive industries, while restrictive digital trade policies correlate with weaker export performance. Alex Haag's (2025) Federal Reserve analysis identifies U.S. advantages in infrastructure, compute capacity, and investment conditions, alongside constraints facing China in advanced semiconductors and persistent gaps in cloud scale and private investment across Europe. These findings suggest that national competitiveness turns on enabling inputs and that heavy restrictions on cross-border data flows or digital trade can impose meaningful costs.

Evidence on model openness further complicates simple competition narratives. Thibault Schrepel and Jason Potts (2025) evaluate 11 foundation models using an 18-variable index focused on licensing and governance. They show that openness operates along a spectrum, rather than a binary divide, with many models clustered in the middle and only modest score differences between systems often labeled "open" or "closed." For competition policy, the implication is narrow but important: arguments that hinge on categorical labels should be tested against enforceable license terms and governance structures.

Open-source AI adoption reinforces this point. Anna Hermansen and Cailean Osborne (2025) report that 89% of organizations using AI incorporate open-source components, with higher adoption rates among small and mid-sized firms due to lower deployment costs and greater flexibility. Survey and case evidence indicate that open-source systems often reduce business-unit costs by more than 50%, by enabling inter-organizational collaboration and faster development cycles. In specialized domains such as health care, open-source models perform comparably to proprietary alternatives, while offering superior integration flexibility in manufacturing and edge-computing environments. These findings suggest that access to AI inputs frequently emerges through market mechanisms, rather than regulation, and that restrictions on contracting or licensing risk disrupting existing diffusion pathways.

Policy commentary from the International Center for Law & Economics (ICLE) echoes these concerns. In comments responding to U.S. Justice Department (DOJ) proposals, Geoffrey A. Manne *et al.* (2024) argue that mandatory data-sharing and access requirements can reduce competition by weakening investment incentives and undermining startups' ability to form partnerships that provide compute and distribution. They further contend that data advantages are not necessarily exclusionary because performance gains from data can saturate, shifting competition toward algorithmic and product innovation.

Giorgio Castiglia's (2025) case for dynamic competition policy reinforces this perspective. In fast-moving technology markets, innovation—not static market share—defines rivalry. In AI markets, where performance improvements can rapidly reallocate competitive advantage, competition enforcement that treats current positions as durable risks misdiagnosing the competitive process.

VI. Policy Implications: Strategic Forbearance and Complementary Investment

ICLE's comments to the U.S. Office of Science and Technology Policy advance a case for "strategic forbearance," urging regulators to rely on existing technology-neutral law while agencies modernize rules built around assumptions of human operators or static systems (Eric Fruits, Ben Sperry, and Kristian Stout, 2025). The argument rests on uncertainty and rapid technological change. When model capabilities, best practices, and deployment structures evolve quickly, prescriptive *ex ante* rules risk imposing high compliance costs, while failing to target actual harm channels. Kristian Stout's later comments similarly warn against burdensome frameworks and emphasize the difficulty of defining "AI" in legally stable terms as technology evolves (Stout, 2025a). Overbroad definitions can impose compliance costs across low-risk uses, creating entry barriers unrelated to the harms regulators seek to address.

Stout (2025b) further argues that federal preemption of conflicting state AI regulations would reduce market fragmentation and compliance costs, facilitating interstate deployment and lowering fixed costs for smaller firms. From a competition perspective, fragmentation functions as an entry barrier. Legal overhead scales with the number of jurisdictions, rather than with output, disproportionately burdening startups and small and medium-sized enterprises.

The broader literature supports applying existing legal frameworks to AI rather than constructing sector-wide regulatory regimes. Many concerns raised in policy debates map cleanly onto established doctrines. Consumer deception falls within consumer protection and unfair practices law. Employment discrimination remains governed by civil rights and labor statutes. Product defects and safety risks are addressed through product-liability rules and sector-specific safety regulation. This approach aligns with the economic principle of targeting: regulation should address specific externalities or market failures, rather than impose general constraints that burden benign uses. It also avoids raising diffusion costs in contexts where the literature finds productivity gains and improvements in work quality (Shakked Noy and Whitney Zhang, 2023; Erik Brynjolfsson, Danielle Li, and Lindsey Raymond, 2025).

Where genuine regulatory gaps exist, Fruits, Sperry, and Stout (2025) recommend pilot programs, waivers, and conditional approvals that allow agencies to learn about risks and benefits before imposing permanent requirements. This approach preserves flexibility while generating evidence.

Multiple strands of evidence underscore the importance of complements. Micro-level experiments show that workers gain substantially when trained to use AI tools effectively, with especially large benefits for less-experienced workers (Noy and Zhang, 2023; Brynjolfsson, Li, and Raymond, 2025; Kevin Zheyuan Cui *et al.*, 2025; Jung Ho Choi and Chloe L. Xie, 2025; Daniel Schwarcz *et al.*, 2025). Evidence on the jagged technological frontier shows that poor task assignment and overreliance can reduce performance, strengthening the case for internal governance and verification, rather than technology-wide restrictions (Fabrizio Dell'Acqua *et al.*, 2023). Comparative-advantage studies highlight the role of STEM supply and infrastructure and link restrictive digital-trade policies to weaker performance in AI-intensive industries (Alessandra Bonfiglioli *et al.*, 2025).

Policy that raises adoption costs through compliance burdens risks delaying the complement-building phase described by the productivity J-curve (Erik Brynjolfsson, Daniel Rock, and Chad Syverson, 2021). Taken together, the evidence points toward an enabling posture: reduce fragmentation, clarify rules, and invest in the complements that allow productivity gains to materialize.

VII. Conclusion

The empirical literature supports several conclusions with relatively high confidence. First, controlled workplace studies consistently show large productivity gains from AI in task categories such as writing, customer support, software development, and translation, often improving quality as well as speed (Shakked Noy and Whitney Zhang, 2023; Erik Brynjolfsson, Danielle Li, and Lindsey Raymond, 2025; Sida Peng *et al.*, 2023; Kevin Zheyuan Cui *et al.*, 2025; Ali Merali, 2024). These gains appear repeatedly across firms, occupations, and experimental designs and are strongest among initially lower-performing workers, producing skill compression, rather than elite-only benefits.

Second, aggregate labor-market effects through 2024–2025 remain limited in most datasets. Studies using administrative records and large surveys find little evidence of economywide job loss or wage decline despite rapid adoption (Martha Gimbel *et al.*, 2025; Anders Humlum and Emilie Vestergaard, 2025; Jonathan S. Hartley *et al.*, 2026). At the same time, several datasets identify pressure in entry-level segments of highly exposed occupations, particularly among younger workers and new hires (Erik Brynjolfsson, Bharat Chandar, and Ruyu Chen, 2025; Bouke Klein Teeselink, 2025). The emerging pattern is adjustment at the margin—through task reallocation and changes in career ladders—rather than broad displacement.

Third, macroeconomic effects remain uncertain. Credible estimates range from modest productivity gains to large output increases, depending on assumptions about task share, diffusion speed, and complementary investment (Daron Acemoglu, 2025; Alex Arnon and Kent Smetters, 2025; Joseph Briggs and Devesh Kodnani, 2023; Erkan Erdem and Dileep Birur, 2025). Measurement challenges and the productivity J-curve further complicate interpretation, as organizational redesign and intangible investment can delay visible gains in aggregate data (Erik Brynjolfsson, Daniel Rock, and Chad Syverson, 2021; Diane Coyle and John Lourenze S. Poquiz, 2025). Divergence across forecasts reflects disagreement over these margins, not simple data error.

Several cross-cutting themes emerge. AI's economic impact depends less on raw model capability than on deployment context, governance, and complements. The "jagged technological frontier" shows that productivity gains hinge on task selection and verification, not blanket automation (Fabrizio Dell'Acqua *et al.*, 2023). Diffusion appears strongest where firms invest in training, workflow redesign, and internal controls. At the market level, AI lowers minimum viable scale and facilitates entry downstream, even as upstream model development may remain concentrated due to fixed costs. Competition, in this setting, is dynamic and innovation-driven, rather than static and share-based.

For policy, these findings point toward restraint coupled with focus. After two years of rapid adoption, the most defensible posture remains strategic forbearance: targeted enforcement of existing law, combined with efforts to reduce regulatory fragmentation and support complementary investment in skills, governance, and infrastructure (Eric Fruits, Ben Sperry, and Kristian Stout, 2025; Kristian Stout, 2025b). Consumer protection, civil rights, and product-liability doctrines already address many identified risks. Where uncertainty persists, pilot programs, waivers, and conditional approvals offer a way to learn without locking in premature mandates.

Proposals for forced data sharing, mandatory access, or structural restrictions on partnerships require stronger evidence of persistent market failure and clearer proof that intervention improves welfare net of dynamic costs (Geoffrey A. Manne *et al.*, 2024). The current literature provides less support for those premises than the speed of AI adoption might suggest. The central policy challenge is not to slow diffusion, but to ensure that institutions, skills, and governance evolve quickly enough for productivity gains to materialize broadly and sustainably.

Annotated Bibliography

- **Acemoglu, Daron,** *The Simple Macroeconomics of AI*, 40 Econ. Pol'y 13 (2025).

Develops a task-based macroeconomic model to estimate AI's aggregate productivity effects. The model distinguishes easy-to-learn tasks, which have objective success metrics, from hard-to-learn tasks, which require contextual judgment. Finds that generative AI will raise total factor productivity by less than 0.66% over a 10-year horizon. Current systems perform well on easy tasks but deliver diminishing returns on hard tasks, which account for a larger share of overall economic activity. Applies Hulten's Theorem to show that aggregate productivity gains are bounded by the GDP share of tasks affected by AI. As a result, even substantial improvements at the task level translate into modest economywide effects. Offers the most conservative credible estimate in the current literature and pinpoints specific technical constraints that limit AI's macroeconomic impact. The analysis is essential for assessing the plausible range of AI-driven growth forecasts.

- **Aldasoro, Iñaki, Sebastian Doerr, Leonardo Gambacorta & Daniel Rees**, *The Impact of Artificial Intelligence on Output and Inflation*, BIS Working Paper No. 1179, Bank for Int'l Settlements (Apr. 2024), https://www.bis.org/publ/work1179.pdf.

Employs a multi-sector macroeconomic model showing that AI adoption functions as a positive productivity shock that raises aggregate output, consumption, and investment in both the short and long run. Simulations show that inflation effects hinge on expectations. When households and firms anticipate productivity gains, immediate demand growth produces inflation. When gains arrive unanticipated, supply initially outpaces demand, resulting in disinflation. At the sectoral level, the authors find little correlation between an industry's initial AI exposure and its long-term output growth. The results also indicate that AI adoption in consumption-good sectors generates substantially larger aggregate output gains than adoption in investment-good sectors, reflecting amplification through sectoral production linkages.

- **Arnon, Alex & Kent Smetters**, *The Projected Impact of Generative AI on Future Productivity Growth*, Penn Wharton Budget Model (Sept. 8, 2025), https://budgetmodel.wharton.upenn.edu/issues/2025/9/8/projected-impact-of-generative-ai-on-future-productivity-growth.

Presents a detailed report combining task-exposure data, adoption assumptions, and empirical productivity studies to estimate AI's macroeconomic effects. The analysis concludes that roughly 10% of current U.S. GDP could be affected in the short run, rising to 15% over two decades under partial-adoption scenarios. Assumes average 25% labor-cost savings from currently available AI tools, with potential gains reaching 40% as the technology matures. Builds a sector-level model that incorporates workforce composition and task-automation potential to translate firm-level efficiencies into aggregate outcomes. Synthesizes evidence from multiple micro-level productivity studies, including Brynjolfsson, Noy & Zhang, and Peng *et al.*, to project economywide effects. Positions its estimates between conservative and highly optimistic forecasts and provides a policy-relevant baseline for evaluating long-term budget and growth implications associated with AI adoption.

- **Bick, Alexander, Adam Blandin & David J. Deming**, *The Rapid Adoption of Generative AI*, NBER Working Paper No. 32966 (2025), https://www.nber.org/papers/w32966.

Documents that generative AI adoption reached nearly 40% of the U.S. population ages 18–64 by late 2024, with diffusion occurring faster than for personal computers or the internet. Analyzes demographic adoption patterns and finds parallels to early PC uptake. Uses Real-Time Population Survey (RPS) data to measure adoption frequency and intensity across users. Demonstrates rapid technology uptake and estimates potential productivity gains based on reported time savings, rather than on market-deregulation effects.

- **Bonfiglioli, Alessandra, Rosario Crinò, Mattia Filomena & Gino Gancia**, *Comparative Advantage in AI-Intensive Industries: Evidence from U.S. Imports*, CESifo Working Paper No. 11642 (2025), https://ssrn.com/abstract=5116412.

Investigates the determinants of global competitiveness in AI-intensive industries using a dataset of U.S. imports from 68 countries across 79 industries from 1999 to 2019. Constructs a novel AI-intensity index based on occupations that require machine-learning and data-analysis skills to measure industry exposure. Finds that countries with larger supplies of STEM graduates, broader internet penetration, and higher overall export volumes exhibit a strong comparative advantage in AI-intensive sectors. Links these structural factors to stronger export performance in high-technology industries. Shows that restrictive digital-trade regulations—particularly those affecting infrastructure and cross-border data flows—correlate with significantly lower export performance in AI-intensive fields. The results remain robust across multiple controls and instrumental-variable analyses using historical scientific data. Highlights the central role of human capital, digital infrastructure, and open regulatory environments in shaping comparative advantage in the digital economy.

- **Briggs, Joseph & Devesh Kodnani**, *The Potentially Large Effects of Artificial Intelligence on Economic Growth*, GOLDMAN SACHS ECON. RSCH. (Mar. 26, 2023), https://www.gspublishing.com/content/research/en/reports/2023/03/27/d64e052b-0f6e-45d7-967b-d7be35fabd16.pdf.

Estimates that generative AI could raise annual global GDP by 7%—nearly $7 trillion—over a 10-year period. Combines data on the task content of more than 900 occupations with adoption-rate assumptions to forecast productivity gains. Projects that roughly 300 million full-time jobs could face automation exposure, while noting that historical patterns show worker displacement often offset by new job creation. Frames the results as a baseline scenario with “potentially large macroeconomic effects” contingent on adoption timelines.

- **Brynjolfsson, Erik, Bharat Chandar & Ruyu Chen**, *Canaries in the Coal Mine? Six Facts About the Recent Employment Effects of Artificial Intelligence*, STAN. DIGIT. ECON. LAB (Nov. 2025), https://digitaleconomy.stanford.edu/app/uploads/2025/11/CanariesintheCoalMine_Nov25.pdf.

Analyzes ADP payroll data covering millions of workers to identify AI's employment effects. Uses Poisson-regression event-study estimations to control for firm-level shocks and isolate the relationship between AI exposure and hiring patterns following ChatGPT's release. Finds that workers ages 22–25 in highly AI-exposed occupations experienced a 16% employment decline relative to trend, while senior-level employment remained stable. Documents that employment effects concentrate in occupations where AI automates, rather than augments, labor. Provides early large-scale evidence of AI's differential workforce impacts, identifying entry-level workers as disproportionately affected.

- **Brynjolfsson, Erik, Danielle Li & Lindsey Raymond**, *Generative AI at Work*, 140 Q.J. ECON. 889 (2025).

Analyzes data from 5,172 customer-support agents at a Fortune 500 enterprise to evaluate deployment of a GPT-based conversational assistant. Uses a staggered-rollout design to identify causal effects, supplemented by a pilot randomized controlled trial. Finds a 15% average productivity increase, measured as issues resolved per hour, with a 36% gain among agents in the bottom skill quintile. The system diffused best practices from top performers, effectively delivering real-time coaching at scale. Reports stable customer-satisfaction scores alongside faster resolution times. Attrition declined by roughly 10 percentage points (40%) among newer agents with AI access. Demonstrates that generative AI can raise the productivity floor while improving work experience, challenging predictions that automation inevitably degrades working conditions.

- **Brynjolfsson, Erik, Daniel Rock & Chad Syverson**, *The Productivity J-Curve: How Intangibles Complement General Purpose Technologies*, 13 AM. ECON. J.: MACROECONOMICS 333 (2021).

Explains why general-purpose technologies often fail to raise measured productivity in the near term despite clear technical capabilities. Firms must undertake complementary investments in "organizational restructuring," worker training, and process redesign before realizing efficiency gains. Treats these intangible investments as short-run costs that temporarily depress measured productivity. Historical evidence from electrification and information technology shows that productivity gains emerge only after substantial lags, in some cases taking "a generation" or longer. Frames this dynamic as the upward trajectory of a productivity J-curve. The framework suggests the economy may remain in the investment phase of the curve with respect to AI, meaning the absence of immediate productivity surges does not undermine more optimistic long-term projections.

- **Cai, Junhui Jeff, Xian Gu, Liugang Sheng, Mengjia Xia, Linda Zhao & Wu Zhu**, *AI as "Co-founder": GenAI for Entrepreneurship*, ARXIV PREPRINT ARXIV:2512.06506 (2025), https://arxiv.org/pdf/2512.06506.

Uses administrative firm-formation data and a difference-in-differences design centered on ChatGPT's November 2022 release to assess effects on business creation. Finds that generative AI facilitated market entry by first-time founders and resource-constrained entrepreneurs, particularly in industries downstream of AI capabilities. Reports that firms formed after ChatGPT's introduction had fewer shareholders and smaller founding teams, consistent with AI substituting for managerial, technical, and operational labor at the startup stage. Interprets these patterns as evidence that AI supplies domain knowledge and functional capabilities that previously required larger teams. Shows that generative AI can lower barriers to entrepreneurship beyond automating existing tasks by enabling new business models viable at smaller scale. Highlights implications for market structure and business dynamism.

- **Caplin, Andrew, David J. Deming, Shangwen Li, Daniel J. Martin, Philip Marx, Ben Weidmann & Kadachi Jiada Ye**, *The ABC's of Who Benefits from Working with AI: Ability, Beliefs, and Calibration*, NBER Working Paper No. 33021 (Oct. 2024), http://www.nber.org/papers/w33021.

Presents a controlled experiment using an age-classification task to examine how individual ability and belief calibration—defined as the accuracy of self-assessment—shape AI-driven productivity gains. Finds that AI assistance improves performance across users but produces the largest effects when individuals are well calibrated in assessing their own skills. Shows that low-ability participants who accurately recognize their limitations realize the greatest performance gains because they rely more effectively on AI recommendations. In contrast, miscalibration—whether overconfidence or underconfidence—produces suboptimal decision-making and constrains the technology's benefits. Counterfactual analysis indicates that if all users held perfectly calibrated beliefs about their abilities, AI's equalizing effects on performance would nearly double. Highlights calibration training as a key workforce complement to AI deployment.

- **Castiglia, Giorgio**, *Rethinking Antitrust: The Case for Dynamic Competition Policy*, INFO. TECH. & INNOVATION FOUND. (Oct. 2025), https://www2.itif.org/2025-dynamic-competition.pdf.

Argues that innovation-based competition should take priority in technology markets where creative destruction can rapidly displace incumbent firms. Although not AI-specific, the framework maps directly onto foundation-model markets, where algorithmic advances can quickly reconfigure competitive positions. Critiques antitrust approaches that prioritize static market-share metrics over dynamic innovation incentives. Contends that such frameworks risk misidentifying competitive harm in fast-moving, technology-driven sectors. Provides a theoretical basis for skepticism toward structural interventions that could preserve existing market configurations at the expense of future innovation. Offers a policy-relevant lens for evaluating competition proposals in AI markets.

- **Chandar, Bharat**, *Tracking Employment Changes in AI-Exposed Jobs*, STAN. DIGIT. ECON. LAB (Aug. 1, 2025), https://ssrn.com/abstract=5384519.

Analyzes U.S. Current Population Survey (CPS) data from late 2022 through early 2025 to assess generative AI's labor-market effects. Finds no aggregate decline in employment or earnings among occupations with the highest AI exposure, countering broad displacement claims. Disaggregated results show substantial heterogeneity. High-exposure, college-degree roles—such as software development—experienced employment growth, while lower-education roles, including customer service, recorded employment declines. Identifies a divergence between online job-postings data and realized employment outcomes. While postings signaled a downturn in tech hiring, observed employment levels remained stable, suggesting postings may serve as an unreliable proxy for labor-market conditions.

- **Choi, Jung Ho & Chloe L. Xie**, *Human + AI in Accounting: Early Evidence from the Field*, STAN. UNIV. & MASS. INST. OF TECH. (Sept. 2025), https://ssrn.com/abstract=5240924.

Analyzes field data from a technology firm alongside survey responses from 277 accountants to assess workplace AI adoption. Finds that AI use correlates with an 18% increase in weekly client support and a reallocation of roughly 9% of work hours from routine data entry to higher-value tasks, including business communication. Documents measurable quality improvements in financial reporting. AI adoption reduced monthly book-closing timelines by 7.5 days and increased ledger-account detail by 12%. Identifies a "Human + AI" complementarity in which experienced professionals use AI confidence scores to target review efforts. A framed field experiment also finds that accountants sometimes over-rely on inaccurate AI outputs, underscoring the continuing importance of professional expertise as a control mechanism.

- **Council of Econ. Advisers**, *Artificial Intelligence and the Great Divergence*, EXEC. OFF. OF THE PRESIDENT OF THE U.S. (2026), https://www.whitehouse.gov/wp-content/uploads/2026/01/Artificial-Intelligence-and-the-Great-Divergence-5.pdf.

Posits that AI could produce a second "Great Divergence" analogous to the Industrial Revolution, with technologically leading nations accelerating ahead of peers. Reports that so-called "Pax Silica" economies—the United States and key allies—averaged 2.5% real GDP growth through Q3 2025, compared with 1.1% across G7 nations. Documents U.S. leadership in core AI inputs, including investment flows, compute capacity—estimated at 74% of the global total—and development of large-scale AI systems. Links this position to national innovation capacity and geopolitical advantage. Frames federal policy around accelerating deployment through infrastructure expansion and deregulation. Projects that exponential capability gains—measured by performance metrics doubling within months—will translate into sustained productivity growth. Reflects the Council of Economic Advisers' position favoring rapid diffusion and regulatory restraint to preserve U.S. technological leadership.

- **Coyle, Diane & John Lourenze S. Poquiz**, *Making AI Count: The Next Measurement Frontier*, NBER Working Paper No. 34330 (Oct. 2025), http://www.nber.org/papers/w34330.

Argues that conventional GDP statistics understate AI's economic impact by failing to capture quality improvements and intangible-capital formation. When AI enables workers to produce higher-quality outputs in less time, input- and revenue-based measures miss a substantial share of resulting welfare gains. Examines both conceptual and practical challenges in measuring AI-driven productivity improvements. Proposes extensions to national accounting frameworks designed to better capture AI-related value creation. Explains that the absence of a pronounced productivity surge in official statistics may reflect measurement constraints rather than a lack of real economic gains. Provides an interpretive framework for assessing macroeconomic data in periods of rapid technological change.

- **Cui, Kevin Zheyuan, Mert Demirer, Sonia Jaffe, Leon Musolff, Sida Peng & Tobias Salz**, *The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers*, Working Paper (Aug. 2025), https://ssrn.com/abstract=4945566.

Investigates generative AI's productivity effects on high-skilled labor through three large-scale field experiments involving nearly 5,000 software developers at Microsoft, Accenture, and a Fortune 100 firm. Evaluates deployment of an AI-based coding assistant across real-world production environments. Finds that developers with AI access increased weekly task completion by 26.08%, alongside measurable gains in code updates and compilations. Reports heterogeneous effects across experience levels, with junior and less-experienced developers showing higher adoption rates and larger productivity improvements than senior peers. Indicates that generative AI can materially increase output in high-skilled occupations without degrading work quality, with particularly strong effects among early-career professionals.

- **Dell'Acqua, Fabrizio, Edward McFowland III, Ethan Mollick, Hila Lifshitz-Assaf, Katherine C. Kellogg, Saran Rajendran, Lisa Krayer, François Candelon & Karim R. Lakhani**, *Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality*, Harv. Bus. Sch. Working Paper No. 24-013 (2023), https://www.hbs.edu/ris/Publication%20Files/24-013_d9b45b68-9e74-42d6-a1c6-c72fb70c7282.pdf.

Reports results from a field experiment with Boston Consulting Group consultants evaluating GPT-4's effects on productivity and work quality. Finds substantial performance gains on tasks within the system's capability range, alongside performance declines on tasks just beyond that frontier. Attributes these declines to overreliance on AI-generated suggestions for problems the system could not reliably solve. Introduces the "jagged frontier" concept to describe uneven AI capability boundaries, where superficially similar tasks can differ sharply in automation potential. Demonstrates the importance of human oversight and the risks of automation complacency. Concludes that effective deployment requires workers to exercise judgment about when to trust, verify, or override AI outputs, highlighting organizational and managerial challenges in integrating these systems.

- **Eloundou, Tyna, Sam Manning, Pamela Mishkin & Daniel Rock**, *GPTs Are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models*, 384 SCIENCE 1306 (2024).

Develops a foundational task-exposure classification for large language models. Uses expert human raters alongside GPT-4 to evaluate which occupational tasks AI systems could perform. Finds that roughly 80% of U.S. workers have at least 10% of their tasks exposed to AI capabilities. Exposure skews toward higher-income occupations, particularly those involving programming and writing, while roles centered on science and critical-thinking tasks show negative correlations. Estimates that about 19% of workers could face exposure across 50% or more of their tasks. Emphasizes that "exposure" reflects technical feasibility for augmenting labor efficiency, not a prediction of automation or displacement. Establishes an empirical baseline for assessing AI's economic and policy implications.

- **Erdem, Erkan & Dileep Birur,** *Generative AI and Economic Growth: A New Approach to Measuring Its Potential Economic Impact*, KPMG (Nov. 2025), https://kpmg.com/kpmg-us/content/dam/kpmg/pdf/2025/gen-ai-economic-growth.pdf.

Employs a dynamic computable general-equilibrium (CGE) model to forecast the economic effects of generative AI adoption across the U.S. and global economies from 2024 through 2050. Models three scenarios—rapid adoption, slow adoption, and a no-generative-AI baseline—to assess how integration rates and workforce-upskilling strategies would shape GDP, wages, and employment across labor categories. Projects that generative AI could generate trillions of dollars in incremental economic value. Under rapid-adoption assumptions, U.S. GDP could increase by approximately $2.48 trillion by 2030. Finds that realizing these gains depends on sustained investment in labor productivity and strategic workforce development to manage transitional labor-market adjustment costs.

- **Fruits, Eric, Ben Sperry & Kristian Stout**, *Comments of the International Center for Law & Economics: Office of Science and Technology Policy RFI, Regulatory Reform on Artificial Intelligence*, INT'L CTR. FOR L. & ECON. (Oct. 27, 2025), https://laweconcenter.org/wp-content/uploads/2025/10/OSTP-AI-Comments-2025.pdf.

Advocates a policy framework of "strategic forbearance" for artificial intelligence governance. Argues that because AI functions as a rapidly evolving general-purpose technology, policymakers should prioritize existing technology-neutral statutes—including those addressing fraud, discrimination, and product safety—rather than adopt premature, prescriptive mandates that risk constraining innovation. Identifies structural barriers to AI deployment. Highlights "regulatory mismatch," in which legacy federal rules assume human operators or static devices, and points to a fragmented landscape of conflicting state-level laws that impose elevated compliance costs. Recommends that federal agencies deploy administrative tools—such as waivers, pilot programs, and conditional approvals—to generate empirical evidence and modernize outdated regulatory frameworks, while preserving innovation incentives.

- **JOSHUA GANS**, THE MICROECONOMICS OF ARTIFICIAL INTELLIGENCE (MIT Press, 2025).

Presents a book-length analytical framework that conceptualizes AI as a technology that reduces prediction costs. Develops formal models showing that, as prediction becomes cheaper, the value of complementary human judgment rises—particularly in setting objectives, preferences, and values that guide how predictions are used. Analyzes AI's effects on pricing, organizational design, and market structure across a range of applications. Applies a microeconomic lens to policy domains including competition law, privacy, and misinformation. Argues that AI improves decision-making efficiency while introducing risks tied to identifiable market failures, including bias and externalities. Provides a theoretical foundation for regulatory approaches that target these discrete failures, rather than imposing broad constraints on capability development.

- **Gimbel, Martha, Molly Kinder, Joshua Kendall & Maddie Lee**, *Evaluating the Impact of AI on the Labor Market: Current State of Affairs*, BUDGET LAB AT YALE (Oct. 1, 2025), https://budgetlab.yale.edu/research/evaluating-impact-ai-labor-market-current-state-affairs.

Assesses AI's labor-market effects using industry- and occupation-level data through August 2025. Examines whether high-exposure sectors exhibit shifts in occupational composition or elevated unemployment rates. Finds no clear correlation between AI exposure and aggregate unemployment. The null results point to broad labor-market resilience during the current adoption phase. Complements studies that identify localized or demographic-specific impacts by showing overall employment stability. Indicates that effects observed in narrower populations have not yet translated into economywide disruption and may take longer to materialize.

- **Haag, Alex**, *The State of AI Competition in Advanced Economies*, FEDS NOTES (Oct. 6, 2025), https://www.federalreserve.gov/econres/notes/feds-notes/the-state-of-ai-competition-in-advanced-economies-20251006.html.

Presents a Federal Reserve comparative analysis of AI competition across advanced economies, focusing on investment flows, research output, and enabling infrastructure. Finds that the United States retains structural advantages in compute capacity, digital infrastructure, and investment conditions. China has expanded research output and accelerated adoption but faces constraints in advanced semiconductor manufacturing and segments of the software ecosystem. European economies lag in private investment and hyperscale cloud capacity. Concludes that existing initial conditions favor continued U.S. leadership, even as international competition intensifies. Supports policy approaches that reinforce foundational enablers, including chip supply, STEM education, and digital infrastructure.

- **Hartley, Jonathan S., Filip Jolevski, Vitor Melo & Brendan Moore**, *The Labor Market Effects of Generative Artificial Intelligence*, Working Paper (Jan. 2026), https://ssrn.com/abstract=5136877.

Presents a comprehensive survey of U.S. workers assessing generative AI adoption and labor-market outcomes. Finds that 35.9% of workers reported using generative AI tools by December 2025, with adoption concentrated among younger, college-educated, and higher-earning employees. Identifies small but statistically significant positive wage effects alongside no measurable change in job openings or aggregate employment across AI-exposed occupations, despite rapid uptake. Links survey responses—aggregated at the occupation level—with administrative data to measure relationships among AI exposure, earnings, and labor demand. Provides one of the most current snapshots of AI diffusion and near-term labor-market effects. Indicates that predictions of immediate, widespread displacement have not materialized and situates current adjustment dynamics within an early adoption phase.

- **Hawk, Ryan, Jeroen van Hoof, Nicki Wakefield & Allen Webb**, *The Leader's Guide to Value in Motion*, PwC (Apr. 29, 2025), https://www.pwc.com/leaders-guide-value-in-motion.

Uses general-equilibrium modeling to quantify the economic effects of artificial intelligence and climate change on global markets through 2035. Projects that AI-driven productivity gains could raise global GDP by nearly 15%, offsetting an estimated 7% contraction tied to physical climate risks and roughly 3% in losses from assets stranded by decarbonization. Maps current industrial sectors to future growth-domain value pools, estimating that these reconfigured markets could generate more than $132 trillion by 2035. Outlines three forward-looking growth scenarios, ranging from 24.9% to 37.2% cumulative expansion, showing how long-term performance will depend on technological adoption, institutional trust, and geopolitical stability.

- **Hermansen, Anna, & Cailean Osborne**, *The Economic and Workforce Impacts of Open Source AI: Insights from Industry, Academia, and Open Source Research Publications*, Linux Found. (May 2025), https://www.linuxfoundation.org/hubfs/Research%20Reports/lfr_marketimpacts25_052725a.pdf.

Analyzes the economic and workforce effects of open-source artificial intelligence (OSAI), identifying it as a key driver of industry innovation and cost efficiency. Reports that 89% of organizations deploying AI incorporate open-source components, with smaller firms adopting OSAI at higher rates due to accessibility and lower deployment costs. Finds that OSAI adoption can reduce business-unit costs by more than 50%, driven by interorganizational collaboration and accelerated development cycles for high-quality models. Shows that open-source models in specialized sectors, including health care, perform comparably to proprietary systems, while offering greater flexibility for operational integration in manufacturing and edge-computing environments.

- **Huang, Xianguo**, *Labor Market Exposure to AI: From GenAI to Future AGI*, Working Paper WP/25-12, ASEAN+3 Macroeconomic Rsch. Off. (Nov. 2025), https://amro-asia.org/wp-content/uploads/2025/11/GenAI_Labour_Huang2025_20251107.pdf.

Finds that higher-income economies face greater labor-market exposure to generative AI due to higher concentrations of cognitive and digital-infrastructure-dependent roles. Concludes that current systems primarily augment human productivity, rather than fully automate tasks, while noting that this balance could shift toward displacement with the emergence of artificial general intelligence. Presents a case study of Brunei showing moderate aggregate exposure but uneven sectoral effects. Finance, insurance, and administrative services display the highest transformation risk. Identifies demographic disparities in exposure. Women and workers with midlevel educational attainment face disproportionately higher exposure relative to other workforce groups.

- **Humlum, Anders, & Emilie Vestergaard**, *Large Language Models, Small Labor Market Effects*, NBER Working Paper No. 33777 (Sept. 2025), https://www.nber.org/papers/w33777.

Links survey data on ChatGPT adoption to administrative earnings and employment records in Denmark across 11 AI-exposed occupations. Finds essentially zero aggregate effects on earnings or hours worked through 2024 despite widespread, worker-reported adoption. Leverages Denmark's comprehensive administrative registers to track individual-level labor outcomes with high precision. The null results contrast with large productivity gains observed in controlled experiments, aligning with a productivity J-curve dynamic in which organizational reorganization and task restructuring delay measurable economic effects. Shows that task-level automation does not translate mechanically into labor-market disruption. Helps explain the gap between micro-level productivity findings and macro-level employment outcomes.

- **Jabarian, Brian, & Luca Henkel**, *Voice AI in Firms: A Natural Field Experiment on Automated Job Interviews*, Working Paper (2025), https://ssrn.com/abstract=5395709.

Reports results from a field experiment that replaced human recruiters with an AI voice agent for initial job interviews at a large recruiting firm. Finds that AI-led interviews increased job offers by 12%, job-start rates by roughly 18%, and 30-day retention by about 18% relative to human-conducted interviews. Shows that applicants offered a choice between AI and human interviewers predominantly selected the AI option. Attributes performance gains to reduced screening bottlenecks and improved candidate-position matching. Demonstrates that AI deployment can enhance labor-market efficiency and employment outcomes beyond cost reduction. Counters claims that AI adoption uniformly degrades hiring processes or worker experience.

- **Johnston, Andrew C., & Christos A. Makridis,** *The Labor Market Effects of Generative AI: A Difference-in-Differences Analysis of AI Exposure*, Working Paper (Oct. 25, 2025), https://ssrn.com/abstract=5375017.

Conducts a difference-in-differences analysis across sectors with varying levels of AI exposure to estimate labor-market effects. Uses the workplace rollout of AI tools beginning in 2021, combined with cross-sector variation in occupational exposure, to identify causal impacts. Finds that higher-exposure sectors recorded significant wage and employment gains, particularly among younger and more-educated workers. Interprets these patterns as evidence of labor-AI complementarity. Also identifies employment declines in sectors where AI directly substitutes for human labor. Highlights the heterogeneity of AI's workforce effects, emphasizing that outcomes depend on whether the technology complements or replaces human inputs.

- **Manne, Geoffrey A., Dirk Auer, Kristian Stout, Lazar Radic & Mario A. Zuñiga**, *Comments of the International Center for Law & Economics: DOJ Invitation to Comment on Promoting Competition in Artificial Intelligence*, INT'L CTR. FOR L. & ECON. (July 15, 2024), https://laweconcenter.org/wp-content/uploads/2024/07/Comments-of-the-International-Center-for-Law-DOJ-AI-RFI.pdf.

Presents a comprehensive critique of mandatory data-sharing and access mandates for AI models. Argues that incumbent data advantages do not constitute durable exclusionary barriers because model performance improves with data scale only up to a threshold, after which algorithmic innovation drives marginal gains. Warns that restrictions on partnerships between large technology firms and AI startups could reduce competition by limiting startup access to compute resources, technical infrastructure, and distribution channels. Frames such collaborations as mechanisms that facilitate entry and scaling, rather than entrench incumbent dominance.

- **Merali, Ali**, *Scaling Laws for Economic Productivity: Experimental Evidence in LLM-Assisted Translation*, ARXIV PREPRINT ARXIV:2409.02391 (Dec. 10, 2024), https://arxiv.org/pdf/2409.02391.

Presents experimental evidence from a randomized controlled trial involving 300 professional translators to measure the economic effects of large language model adoption across varying training-compute levels. Establishes quantifiable "scaling laws" linking model capability to productivity outcomes. Finds that a tenfold increase in training compute reduced task-completion time by 12.3% and improved translation quality by 0.18 standard deviations. These gains produced a 16.1% increase in earnings per minute for participants. Identifies substantial heterogeneity in outcomes. Lower-skilled translators realized productivity gains roughly four times larger than those of higher-skilled peers, indicating that continued model scaling could raise aggregate productivity while narrowing skill-based wage differentials.

- **Noy, Shakked, & Whitney Zhang**, *Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence*, 381 SCIENCE 187 (2023).

Reports results from a randomized experiment involving 453 professionals performing writing tasks. The treatment group used ChatGPT-3.5, while the control group completed assignments without AI assistance. Finds a 40% reduction in task-completion time and an 18% improvement in output quality, as rated by independent evaluators. Quality gains concentrated among workers in the bottom half of the initial skill distribution, producing a measurable reduction in performance inequality. Participants also reported higher job satisfaction and self-efficacy when using AI tools. Provides controlled evidence that generative AI can materially increase knowledge-worker productivity while improving, rather than degrading, job quality. Demonstrates a skill-compression effect in which AI enables lower-skilled workers to approach expert-level performance.

- **Organisation for Econ. Co-operation & Dev. (OECD)**, *Artificial Intelligence and Competitive Dynamics in Downstream Markets*, OECD Roundtables on Competition Pol'y Papers No. 331 (2025), https://www.oecd.org/en/publications/artificial-intelligence-and-competitive-dynamics-in-downstream-markets_ccf0624a-en.html.

Analyzes AI's effects on entry barriers and minimum-efficient scale across industries. Finds that AI enables smaller firms to automate functions—such as marketing analytics, customer service, and data processing—that previously required large, specialized teams. Examines vertical integration between model providers and application developers, identifying mixed competitive effects. Integration can foreclose rivals in some contexts, while improving coordination efficiency in others. Provides an international comparative perspective on AI-driven market structure. Clarifies how AI reshapes business organization and competitive dynamics beyond the foundation-model layer.

- **Organisation for Econ. Co-operation & Dev. (OECD)**, *How Are SMEs Using Generative AI?* (2025), https://www.oecd.org/en/publications/generative-ai-and-the-sme-workforce_2d08b99d-en/full-report.html.

Presents findings from a 2024 cross-country survey of more than 5,000 small- and medium-sized enterprises examining generative AI adoption and workforce effects. Reports that roughly 31% of SMEs across the seven surveyed countries have adopted generative AI, with users citing improved employee performance and reduced workloads. Finds limited evidence of automation-driven job loss: 83% of adopting firms reported no change in overall staffing levels, instead using AI to address labor shortages and offset skill gaps. Identifies persistent barriers to broader diffusion. Nonadopters most frequently cited lack of business-use suitability (57%), legal and data-privacy concerns (54%), and insufficient internal workforce skills (50%). Suggests that while generative AI delivers productivity gains and can level competitive conditions, effective deployment currently favors firms with higher-skilled labor, potentially widening the SME digital divide.

- **Peng, Sida, Eirini Kalliamvakou, Peter Cihon & Mert Demirer**, *The Impact of AI on Developer Productivity: Evidence from GitHub Copilot*, ARXIV PREPRINT ARXIV:2302.06590 (2023), https://arxiv.org/pdf/2302.06590.

Reports results from a controlled trial evaluating GitHub's Copilot AI coding assistant. Finds that developers using the tool completed programming tasks 55.8% faster than control-group participants. Identifies heterogeneous effects across user groups. Less-experienced developers, older programmers, and those who code more hours per day realized the largest productivity gains. Demonstrates substantial efficiency improvements in software development and highlights AI's potential to lower skill barriers and support workforce entry into programming occupations.

- **Schrepel, Thibault, & Jason Potts**, *Measuring the Openness of AI Foundation Models: Competition and Policy Implications*, 34 INFO. & COMM. TECH. L. 279 (2025).

Evaluates the openness of 11 prominent AI foundation models—including GPT-4, Llama 3, and Gemini—by analyzing licensing structures through an innovation-commons framework, rather than purely technical criteria. Constructs an index of 18 variables grouped into knowledge sharing, anti-opportunism, and governance dimensions. Finds that openness operates along a spectrum, rather than a binary "open" versus "closed" classification. Most models cluster in the middle-to-lower range, with limited differentiation between proprietary systems and those marketed as open source. For example, Meta's Llama 3 and OpenAI's GPT-4 differ by only two index points. Shows that higher openness scores correlate with stronger upstream access to code and datasets. At the same time, most models lack robust participatory-governance mechanisms and safeguards against opportunism, regardless of whether they originate from large technology firms or smaller developers.

- **Schwarcz, Daniel, Sam Manning, Patrick Barry, David R. Cleveland, J. J. Prescott & Beverly Rich**, *AI-Powered Lawyering: AI Reasoning Models, Retrieval Augmented Generation, and the Future of Legal Practice*, J. L. & EMPIRICAL ANALYSIS (forthcoming 2026) (Dec. 6, 2025), https://ssrn.com/abstract=5162111.

Reports results from a randomized controlled trial assessing how advanced AI tools affect the quality and efficiency of legal work produced by law students. Compares two systems—a specialized legal platform using retrieval-augmented generation (Vincent AI) and a general-purpose reasoning model (OpenAI's o1-preview)—against a no-AI control group. Finds that both tools significantly improved legal-document quality across most tasks, with gains concentrated in clarity, organization, and professional tone. The reasoning model produced the largest improvements in analytical depth, while the retrieval-augmented system reduced hallucinations and fabricated citations. Documents substantial productivity gains. Students completed assignments 50% to 130% faster with AI assistance, indicating that newer reasoning and retrieval architectures mitigate quality constraints observed in earlier-generation models.

- **Siddik, Abu Bakkar, Yong Li & Anna Min Du**, *Unlocking Funding Success for Generative AI Startups: The Crucial Role of Investor Influence*, 69 FIN. RSCH. LETTERS 106203 (2024).

Examines the determinants of funding success across 556 generative AI startups operating between 2010 and July 2024. Uses principal-component analysis and regression modeling to evaluate how financial and technological factors shape capital formation. Finds that investor influence—captured through variables such as investor count, lead-investor participation, and funding-round depth—exerts a significant positive effect on total funding raised. In contrast, technological influence, including IT spending and patent activity, shows no statistically significant relationship with funding outcomes. Concludes that investor networks and financing structures play a more decisive role in capital accumulation for generative AI ventures than technology metrics alone.

- **Stout, Kristian**, *Comments of the International Center for Law & Economics: Request for Information on the Development of an Artificial Intelligence (AI) Action Plan*, INT'L CTR. FOR L. & ECON. (Mar. 14, 2025), https://laweconcenter.org/wp-content/uploads/2025/03/OSTP-AI-2025-comments-v-1.pdf.

Advocates a regulatory framework that balances risk management with the imperative to sustain U.S. innovation leadership. Warns that burdensome or premature rules could constrain the AI sector's dynamic growth trajectory. Critiques international regulatory models, including the EU's AI Act, and highlights the practical difficulty of defining artificial intelligence for legal purposes. Argues that overly broad or fragmented definitions risk undermining economic competitiveness and regulatory coherence. Calls for an evidence-based policy approach that accounts for open-source development dynamics and copyright-law considerations. Frames this calibrated governance model as necessary to preserve both technological safety and competitive vitality.

- **Stout, Kristian**, *Federal Preemption and AI Regulation: A Law and Economics Case for Strategic Forbearance*, WLF LEGAL PULSE (May 30, 2025), https://www.wlf.org/wp-content/uploads/2025/05/053025Stout_LP.pdf.

Argues for federal preemption of conflicting state-level AI regulations to prevent market fragmentation and excessive compliance costs. Examines Commerce Clause constraints on state laws that produce extraterritorial effects. Contends that a unified federal framework enabling interstate

AI-service deployment would support broader adoption than a patchwork of incompatible state requirements. Applies a law & economics framework to questions of regulatory federalism. Clarifies legal constraints on state AI governance and articulates the economic case for federal harmonization.

- **Struta, Iuri**, *GenAI Funding on Track to Set New Record in 2024*, S&P GLOB. MKT. INTEL. (Nov. 4, 2024), https://www.spglobal.com/market-intelligence/en/news-insights/research/genai-funding-on-track-to-set-new-record-in-2024.

Quantifies venture-capital investment trends in generative AI, reporting more than $20 billion invested in startups through Q3 2024, surpassing the pace recorded in 2023. Highlights record-scale fundraising by frontier-model firms, including funding rounds exceeding $8 billion for companies such as Anthropic as they compete with large technology incumbents. Finds that a substantial share of startups concentrate on application-layer and vertical-specific solutions, rather than foundation-model development. Provides market-structure evidence showing capital dispersion across the AI stack. Reflects strong investor optimism despite high capital-burn rates and uncertain profitability timelines among many entrants. Offers a financing and capital-allocation lens for assessing competitive dynamics in the AI sector.

- **Teeselink, Bouke Klein**, *Generative AI and Labor Market Outcomes: Evidence from the United Kingdom*, KING'S COLL. LONDON (Dec. 21, 2025), https://ssrn.com/abstract=5516798.

Uses a difference-in-differences framework centered on ChatGPT's release to analyze U.K. labor-market outcomes. Finds that firms with higher large-language-model exposure reduced overall employment, with losses concentrated among junior and entry-level roles. Documents hiring contractions in technical and creative occupations, including software engineering and design, while demand for interpersonal roles—such as sales—remained stable or increased. Indicates that AI adoption reshaped workforce composition, rather than uniformly reducing labor demand. Shows that average firm-level compensation rose, driven by the shedding of lower-paid junior staff. At the same time, advertised salaries for exposed occupations declined, signaling weaker market demand for those skills. Finds that displacement effects concentrate in higher-wage labor segments, diverging from earlier automation waves that primarily affected low- to middle-skill routine work.

- **Weinstock, Lida R., & Paul Tierno**, *The Macroeconomic Effects of Artificial Intelligence*, CRS In Focus IF12762, CONG. RSCH. SERV. (Apr. 1, 2025), https://www.congress.gov/crs_external_products/IF/PDF/IF12762/IF12762.4.pdf.

Surveys estimates of AI's macroeconomic effects, finding that projected GDP impacts vary widely but remain "typically positive." Cites benchmark projections, including a Goldman Sachs estimate of a 0.9% cumulative GDP increase over 10 years and another study projecting long-run output gains of up to 35% above baseline. Emphasizes that realized economic effects depend on adoption rates and diffusion dynamics, noting that widespread integration may unfold over "multiple decades," mirroring the trajectory of personal-computer deployment. Addresses labor-market implications, concluding that current evidence shows AI replacing tasks rather than entire jobs. Frames these dynamics as likely to influence productivity growth and income distribution over time.

- **Yusuf, Shahid**, *The Macroeconomic Consequences of AI: The Next Great Divergence*, U.N. DEV. PROGRAMME (2025), https://www.undp.org/sites/g/files/zskgke326/files/2025-12/the-macroeconomic-consequences-of-ai.pdf.

Investigates AI's macroeconomic implications, focusing on its capacity to counter declining productivity and reshape income distribution across the Asia-Pacific region. Documents a sustained global slowdown in economic growth and total factor productivity since the early 2000s, noting that many high-growth Asian economies have relied more on capital accumulation than on efficiency gains. Frames AI as a potential response to this "productivity drought" while emphasizing uneven and time-lagged benefits. Finds that adoption effects will likely vary across sectors and labor groups, with gains materializing gradually rather than immediately. Warns that AI diffusion could widen within-country income disparities, reversing the decline in global inequality observed from the late 1990s through 2015. Positions distributional divergence as a central policy risk accompanying AI-driven productivity growth.