



BANK OF CANADA
BANQUE DU CANADA

Staff analytical paper / Document analytique du personnel—2026-4

Last updated: March 3, 2026

AI Paradox: Promise vs. Reality—What It Means for Monetary Policy

Joshua Brault

Data and Digital Services Department
Bank of Canada
jbrault@bank-banque-canada.ca

Maryam Haghghi

Data and Digital Services Department
Bank of Canada
mhaghghi@bank-banque-canada.ca

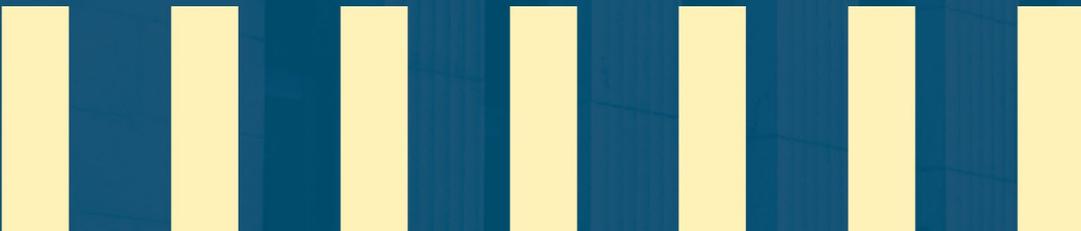
Jing Yang

Economic and Financial Research
Department
Bank of Canada
jyang@bank-banque-canada.ca

Bank of Canada staff research is produced independently from the Bank's Governing Council and may support or challenge prevailing views. The views expressed in this paper are solely those of the authors and may differ from official Bank of Canada positions. No responsibility for them should be attributed to the Bank.

DOI: <https://doi.org/10.34989/sap-2026-4> | ISSN 3110-9195

© 2026 Bank of Canada



Acknowledgements

We would like to thank Stefano Gnocchi, Oleksiy Kryvtsov, Jim MacGee and Alexander Ueberfeldt for their comments and suggestions in preparing this note.

Introduction

Technological progress has repeatedly reshaped work. From 19th-century mechanization to the late-20th-century computer revolution, each wave sparked anxiety about the future of jobs—followed, often with a lag, by adjustment and renewal. Artificial intelligence (AI) is the latest chapter in this history. Growth in the availability of granular and timely data, combined with advances in machine learning, natural language processing, and generative models have revived expectations of “intelligent automation” capable of executing tasks once considered uniquely human or out of reach for technology.

Thus far, however, economic outcomes trail the technological headlines. Despite rapid improvements in AI capabilities, measured productivity growth in most advanced economies remains subdued and aggregate employment has not collapsed. Firms report some efficiency gains from AI-enabled tools, but adoption is still narrow—concentrated among large, digitally capable enterprises and primarily in advanced economies. This gap between technological possibility and measured performance—the “productivity paradox of AI”—has become a central puzzle.

Early labour-market signals suggest adjustment is underway. Brynjolfsson, Chandar, and Chen (2025) document declines in entry-level hiring within occupations most exposed to generative AI, especially in marketing, design, and customer support. While aggregate effects remain modest, these shifts may be a “canary in the coal mine,” reflecting diffusion at the task margin where automation is easiest, with broader consequences emerging as adoption deepens and organizations learn the technology and focus their efforts on re-imagining their workflows with AI.

Understanding AI’s impact requires moving beyond simple substitution narratives. Building on task-based frameworks (e.g., Acemoglu and Autor 2011; Acemoglu and Restrepo 2018), AI acts on tasks, not occupations wholesale. Thus, automating some routine or information-processing tasks while augmenting others through better prediction, retrieval, and creative support. Net effects on employment and wages depend on whether new tasks and roles absorb displaced labour, and on the speed with which firms reorganize production to realize productivity gains.

In this article, we synthesize the emerging evidence on these dynamics. We situate AI within the literature on general-purpose technologies and outline the conceptual frameworks economists use to analyze its impact. We then review the empirical literature on the observed impacts of AI on productivity and employment, emphasizing the mechanisms of automation and augmentation that drive labour market reallocation. We conclude by discussing the macroeconomic and monetary policy implications of potential output (Y^*), inflation dynamics, and the natural rate of unemployment (U^*),

emphasizing that impacts will unfold gradually. The policy challenge is not to forecast a single equilibrium outcome, but to facilitate the transition. For monetary policymakers, better parsing structural dislocation from cyclical weakness in labour market in real time is crucial.

Conceptual Framework: AI as a General-Purpose Technology

General-purpose technologies (GPTs) are transformative innovations that diffuse broadly across a wide array of industries, improve continually over time, and give rise to complementary innovations and economy-wide spillovers. Historical examples include the steam engine, electricity, computers and the internet. Each of these episodes was characterized not only by technological breakthroughs, but also by prolonged periods of organizational adaptation, investment in complementary capital, and the emergence of new business models. AI is increasingly viewed as a GPT, with systems being deployed across an expanding range of sectors, improving rapidly, and generating substantial complementary innovations (e.g., collaborative robots, smart grids, diagnostic imaging systems, compliance systems, and autonomous vehicles and drones).

If history is any indication, the most transformative effects of AI are likely still ahead (Agrawal, Gans, and Goldfarb 2019b; Calvino, Haerle, and Liu 2025). The diffusion of GPTs within and across firms, sectors, and occupations tends to be gradual and uneven, driven by experimentation and learning among early adopters. Realizing the full benefits of GPTs requires substantial complementary investments, organizational changes, and human capital accumulation. Firms must redesign workflows, train workers, and integrate new technologies into production. For example, the full impact of electrification did not impact productivity growth in US manufacturing until the 1920s—roughly 40 years after the first central power station opened (David 1990). Similarly, ICT investment took around a decade to appear in industry-level productivity statistics (Basu and Fernald 2008). However, the broader effects of AI may materialize more rapidly than previous GPTs because AI is primarily software-based and can be distributed instantaneously via the internet and cloud infrastructure for near-zero cost (Baily, Brynjolfsson, and Korinek 2023).

Given the lag between advances at the technological frontier and their measurable economic impact, conceptual frameworks remain essential for interpreting early evidence, assessing the current stage of the transition, linking micro-level mechanisms to aggregate outcomes, and anticipating longer-run effects.

The dominant paradigm in recent academic work is the task-based approach (Acemoglu and Autor 2011; Acemoglu and Restrepo 2018, 2019). In this framework, production is

presented as a series of tasks that can be allocated to capital or labour. New technologies can increase the productivity of capital and labour in tasks they already perform, while also altering how these tasks are allocated across factors of production—offering a unified framework for analyzing a broad range of AI-related mechanisms. Within this framework, technological change can operate through two primary channels:

- **Automation:** Refers to the development and adoption of technologies that expand the extensive margin of tasks that can be performed by AI systems and capital. This process generates a **displacement effect**, as machines take over tasks previously carried out by labour.
- **Augmentation:** Refers to technologies that enhance the outputs of occupations (e.g., efficiency, quality, variety) but leave human input indispensable. Such innovations can also give rise to new tasks in which labour holds a comparative advantage—a process referred to as the **reinstatement effect**.

By making these mechanisms explicit, the task-based framework provides clarity on the net effect of multiple, often countervailing forces, helping to assess how AI may reshape the composition of work. AI transforms work both through the reallocation of existing tasks and the creation of new ones. For example, automation can displace existing workers, but by raising productivity, it also increases labour demand in other, non-automated tasks—a channel referred to as the **productivity effect**. The net effect on overall labour demand depends on the relative strength of these displacement and productivity effects. Moreover, the automation and augmentation mechanisms can occur concurrently, and their dominance on a specific job class and its tasks can vary over time.¹

Empirical Evidence: AI's Effects on Jobs and Productivity

AI and Jobs

The rise of AI, much like earlier GPTs, has renewed fundamental questions about the future of work. Will AI primarily automate human labour, or will it instead augment it? What will happen to labour demand—will the emergence of new types of occupations and tasks be sufficient to offset displacement from automation? How will this process

¹ Although the task-based approach has become a widely used paradigm, other overlapping and complementary frameworks are also employed to study AI. Aghion, Jones, and Jones (2019) adopt a task-based structure embedded within an endogenous growth framework, emphasizing how AI-driven automation affects both production and innovation, including the creation of new ideas. Other contributions build on neoclassical growth frameworks to study broader implications of AI, such as existential risks (Jones 2024) and the future evolution of the Kaldor facts (Trammell and Korinek 2025).

play out over the short-, medium-, and long-run? Should this process be guided, and if so, how?

Some forecasts suggest that AI could reshape the labour market on an unprecedented scale, largely because the technology has the potential to perform a wide array of tasks across many different domains and occupations. Task-based measures of occupational exposure consistently indicate that knowledge workers—especially those involving analysis, communication, and information-processing—are highly susceptible to AI capabilities, while in-person and manual jobs remain least exposed (e.g., Webb 2020; Felten, Raj, and Seamans 2021, 2023). This technology marks a significant departure from previous waves of automation, which primarily affected routine physical and cognitive tasks (Autor, Levy, and Murnane 2003). Moreover, in advanced economies, knowledge workers constitute a significant share of employment. For example, Eloundou et al. (2024) estimate that around 80% of U.S. workers have at least 10% of their tasks exposed to AI, and approximately 19% could have at least half of their tasks impacted. When anticipated software and complementary technologies are included, the projected impact rises dramatically—nearly half of all workers could see more than 50% of their tasks exposed. However, AI exposure does not necessarily imply job displacement, as the effect on labour demand depends on whether AI substitutes for or complements human labour—a distinction more recent work has sought to clarify (Handa et al. 2025; Tomlinson et al. 2025).

Despite forecasts of widespread job displacement, aggregate evidence to date suggests that the labour market effects of AI remain limited. Most advanced economies have not yet experienced significant shifts in overall employment or unemployment rates that can be directly attributed to AI. For example, Gimbel et al. (2025) document that in the 33 months following ChatGPT's release, U.S. employment composition has remained broadly stable, with no concentration of unemployment in highly exposed industries. Similarly, Chandar (2025) finds no systematic differences in U.S. employment changes between more and less AI-exposed professions in the CPS. Furthermore, while many advanced economies show signs of cooling labor markets, it is difficult to determine if this is due to AI or a number of other ongoing cyclical factors.

While widespread job displacement has not yet materialized, pockets of the labour market are beginning to show cracks, suggesting that adjustments are emerging beneath the surface of otherwise stable aggregate trends. One key example is the impact on early-career workers. Brynjolfsson, Chandar and Chen (2025) provide evidence of “canaries in the coal mine”, the emergence of displacement before possibly becoming more widespread. Using high-frequency administrative data from ADP, the largest U.S. payroll provider, they study labour-market outcomes by AI exposure across age groups. They find significant job displacement effects among early-career workers

(ages 22-25) in occupations most exposed to AI, such as software development and customer support. These effects manifest as stagnating employment levels rather than outright job losses—a pattern consistent with the “no-hire, no-fire” dynamic—while employment among older workers in the same occupations has continued to rise. Early-career employment in the most AI-exposed occupations experienced a *relative* decline of 13%. Furthermore, they show that this effect depends on how AI is deployed: early-career employment declines in occupations where AI primarily automates tasks, while the effects are much more muted when AI augments tasks.

While clear evidence of job displacement is perhaps only beginning to emerge, there is much stronger evidence that AI has already augmented many jobs. This has occurred through the reallocation of job tasks within roles, the addition of new types of tasks, and by evolving skill and competency demands from firms. For example, in an experimental setting, Dillion et al. (2025) show that after the adoption of Microsoft’s M365 Copilot, workers reallocated time from emails and document completion to collaboration and preparation. Humlum & Vestergaard (2025) find that 8 percent of employees in their Danish survey reportedly took on new tasks following the introduction of LLMs. Acemoglu et al. (2022) analyze establishment-level vacancy data and find evidence that AI adoption reduces hiring for non-AI positions and increases demand for AI-related roles. Job postings increasingly emphasized skills compatible with AI technologies. Similarly, Green (2024) documents a rising demand for social and language skills in AI-intensive firms, consistent with durable human advantages in social and interpersonal domains (Deming 2017).

What does this mean

Although aggregate employment remains largely stable, emerging evidence points to localized disruptions in certain segments of the labour market. The big question is whether these localized effects will broaden as AI adoption accelerates and diffuses across the broader economy. Combined, the above literature gives evidence for

1. **Augmentation baseline:** Near-term, most impacts manifest as within-occupation task mix changes, tooling, and workflow redesign rather than widespread job separations.
2. **Adjustment frictions:** Early-career workers are the “canaries in the coal mine”: fewer entry-level tasks, steeper skill thresholds, and slower job ladders in exposed fields.
3. **Heterogeneity by deployment:** Automation-first use cases depress junior hiring; augmentation-oriented adoption preserves or enhances roles, conditional on complementary skills.

Taken together, the existing evidence indicates that although AI has not yet caused widespread labour-market disruption, substantial uncertainty remains regarding its longer-term effects on employment. If current and forthcoming AI innovations are primarily augmenting, historical evidence suggests that new work—such as new job categories and tasks—will emerge to offset those displaced by technology, and that this new work may be especially valuable (Autor et al. 2022). Conversely, if AI innovations are largely automating, such countervailing forces may not materialize. In the extreme, the development of artificial general intelligence (AGI) may result in machines being able to perform essentially all human tasks. Such a scenario could lead to a decoupling of employment and growth, a substantial shift in the role and functioning of the labour market, and pose fundamental challenges for existing policy frameworks.

AI and Productivity

Historically, GPTs have been precursors to periods of sustained productivity growth and have profoundly influenced the organization of production, types of work, and the broader structure of the economy. Predictions regarding the impact of AI on productivity gains are particularly striking—even in comparison to previous GPTs—due both to the wide range of tasks that AI could automate and its potential to accelerate innovation. These innovation-driven gains, which might be imperceptible in any single year and masked by cyclical fluctuations, can accumulate and yield significantly larger long-run productivity growth (Aghion, Jones, and Jones 2019).

Recent projections suggest that AI could substantially accelerate productivity growth in advanced economies over the coming years.² For example, Filippucci et al. (2025) estimate that labour productivity growth could increase by approximately 0.15 to 1.30 percentage points (p.p.) annually across G7 countries over the next decade, depending on the speed of AI adoption, future technological capabilities, and sectoral composition. For Canada specifically, the projected annual gains range from 0.35 to 1.13 p.p. for the next ten years. Given that Canada’s labour productivity growth averaged under 1 percent over the past decade, these gains represent a significant increase—potentially doubling the rate of growth.

Although forecasts anticipate substantial productivity gains from AI, there is little empirical evidence that these gains have materialized at the aggregate level. Labour productivity growth in most advanced economies remains subdued with average annual growth rates in the range of 1%; TFP tells a similar story, albeit with lower magnitudes. This divergence between technological capabilities and corresponding evidence in aggregate productivity is described as the “productivity paradox of AI”. While aggregate

² Throughout the literature, the term ‘productivity’ is used to refer to both labour productivity and total factor productivity (TFP). Where possible, we distinguish between these two.

productivity growth remains modest, it is unclear whether this reflects a genuine absence of productivity gains from AI or simply underscores the difficulty of disentangling and quantifying AI's contribution amid broader structural and cyclical forces.

However, recent experimental studies provide direct causal evidence of sizable productivity gains from AI adoption. Experimental studies focusing on knowledge workers find large task-level productivity gains from AI tools. Across settings such as document preparation (Dillon et al. 2025), professional writing (Noy and Zhang 2023), and software development (Peng et al. 2023), access to generative-AI assistants reduced completion times by roughly 30 to 60 percent and improved output quality, with the largest benefits among lower-skilled and less-experienced workers.

Firm-level evidence paints a similar picture for productivity gains. Czarnitzki, Fernández, and Rammer (2023) analyze survey data covering German firms from manufacturing and business-oriented service sectors and find that AI adopters exhibited significantly higher productivity, sales, and value added, with both adoption and intensity of use being positively correlated with performance. Brynjolfsson, Li, and Raymond (2025) study the staggered rollout of a generative-AI conversational assistant across thousands of customer-support agents, showing that access to the tool increased issues resolved per hour by about 15 percent and raised customer sentiment. Consistent with findings from experimental studies, the largest gains occurred among lower-skilled and less-experienced workers, supporting the notion that AI captures and conveys some tacit organizational knowledge and best practices—insights that previously could only be gained through on-the-job experience. Dell'Acqua et al. (2024) analyze professional consultants performing complex knowledge-work tasks and find that those with access to AI completed 12 percent more tasks, were 25 percent faster, and produced work of over 40 percent higher quality than those without AI. However, for tasks that fell outside the AI's capabilities, performance declined—workers were 19 percentage points less likely to reach correct solutions—underscoring that AI's effects are highly task-dependent and contingent on the alignment between tool capabilities and task complexity.

While the micro-level evidence points to substantial productivity gains from AI adoption, the absence of similar gains in aggregate measures remains a subject of ongoing debate. Several factors may help to explain this apparent disconnection:

1. **Limits to extrapolate the speed and scope of adoption:** Early adopters are often atypical—large, digital, and technologically sophisticated firms that are already well positioned to integrate and benefit from AI. As a result, initial estimates may overstate the average impact and provide limited guidance on aggregate effects. Moreover, many current use cases remain narrow in scope,

and a substantial share of remaining tasks may be “hard-to-learn” for AI systems (Acemoglu 2024). Finally, even when automation is technologically feasible, it may not be economically attractive under current cost structures or productivity thresholds (Svanberg et al. 2024).

2. **Implementation and Restructuring Lags:** GPTs require firms to make large and time-consuming investments in intangible capital—such as new business processes, organizational redesign, worker training, and software integration—to realize their full potential. Aggregate productivity effects only become noticeable after firms have undergone these organizational changes, a sufficiently large stock of the new technology has been accumulated, and complementary innovations have emerged and diffused throughout the economy (Brynjolfsson, Rock and Syverson 2018). Additionally, the speed of diffusion may also be hampered by early adopters to maintain competitive advantages. Such implementation and restructuring lags can produce a “Productivity J-Curve”, where measured productivity initially stagnates or declines before rising sharply (Brynjolfsson, Rock, and Syverson 2021). McElheran et al. (2025) provide micro evidence consistent with this pattern. Analyzing tens of thousands of U.S. manufacturing establishments and firms, they find that those adopting AI experience short-term productivity losses before realizing longer-run gains.
3. **Measurement challenges:** The economic value being generated from AI is not fully captured in current measures of productivity and output. Since AI primarily impacts cognitive labour—who comprise a large share of workers in advanced economies and whose productivity is notoriously difficult to quantify—this may lead to a significant under-measurement or “silent productivity growth” (Bailey, Brynjolfsson, Korinek 2023). Additionally, large investments in intangible capital are often recorded as expenses, making them difficult to observe and measure as productive investment. Consequently, aggregate productivity may be growing faster than official statistics suggest, but mismeasurement leads to incorrect assessments about both the current state and recent past. A similar issue occurred during the 1990s ICT boom, when real-time labour productivity estimates underwent substantial upward revisions due to data limitations, measurement issues, and the omission of intangible complementary investments (Jorgensen, Ho, Stiroh 2008).

These explanations are not mutually exclusive, and the timelines for assessing the validity and relative importance of each are likely to differ. Ultimately, the productivity dynamics that emerge will depend on the relative magnitudes and persistence of these forces. If AI leads to broad-based automation and innovation, productivity growth could surge in line with some forecasts. However, productivity gains may fall short of these

projections for the reasons discussed above, or because of other constraints, including those related to data availability, ownership, and property rights.

Given the potential measurement challenges and implementation lags, it will be important to continue to monitor the impacts of AI. This includes tracking both micro-level indicators—such as firm-level adoption, investment flows, and task-level productivity—and macro-level trends in aggregate productivity, labour markets, and sectoral shifts.

Implications for Monetary Policy

Artificial intelligence (AI) represents a tectonic structural transformation that is likely to alter the long-run equilibrium of many economies. While the near-term aggregate labour-market effects are likely to remain limited, the diffusion of AI technologies will, over time, reshape production structures, relative prices, and the composition of labour demand. From a macroeconomic perspective, these changes imply a shift in the so-called “star variables”—the economy’s potential output (Y^*) and the natural rate of unemployment (U^*).

Long-run neutrality

In the long run, monetary policy is neutral with respect to real variables: adjustments in the policy rate cannot permanently influence Y^* or U^* .

The introduction of AI is expected to raise potential output through productivity gains. The long-term impact of AI adoption on the natural rate of unemployment (U^*) is widely expected to be neutral or to lower it, despite potential short-term increases due to job displacement during the economic transition. Historical precedents suggest that technological progress ultimately creates more jobs than it destroys, driving overall economic growth and labour demand. At the same time, more speculative scenarios consider the case of AGI, where machines are able to perform most tasks cheaper and more efficiently than humans, leading to a potential total collapse of the human labour market (Korinek and Juelfs 2022; Korinek and Suh 2024). However, the likelihood of such outcomes is generally viewed as low, but remains a subject of ongoing debate.

On inflation, AI and digitalisation affect inflation and price setting through two main channels: broad macroeconomic shifts in supply and demand, and direct changes in how firms calculate and update their prices. Over the long-term, broad-based productivity gains can lower production costs, and all else equal, reduce inflationary pressure. But over the short-term, high demand for AI infrastructure (chips, power, and data centers) and investment-driven demand can put upward pressure on inflation. Digitalisation has also fundamentally changed the speed and precision with which companies set prices. Digitally intensive firms adjust prices more frequently than

traditional businesses (Gorodnichenko and Talavera 2017). Algorithmic pricing monitors market demand, competitor prices, and input costs in real-time, allowing for near-instantaneous updates. Thus, prices are more responsive to shocks in a digital economy, and the relationship between economic activity and inflation (the Phillips Curve) may be steeper and more volatile than in the past. Additionally, market structure is likely to play an important role in determining how AI adoption influences firms' pricing behaviour. If these firms are primarily located in less competitive markets, such as oligopolies, lower costs may not translate into lower prices (Athey and Morton 2025).

Ultimately, such structural shifts are likely to unfold gradually, driven by capital deepening, diffusion lags, and institutional adaptation. Consequently, even as the steady-state equilibrium improves, the transition path may involve periods of labour-market dislocation and temporary output gaps.

The role of monetary policy in a transition

During the adjustment phase, monetary policy can play a crucial stabilising role. The introduction of AI is likely to alter production in ways that make the current allocation of labour and capital increasingly misaligned across sectors—routine-intensive occupations and industries with high potential for automation may contract, while AI-complementary activities expand. However, labour and capital do not move frictionlessly across sectors. The ensuing reallocation frictions—stemming from skill mismatches, retraining delays, geographic immobility, or regulatory barriers—can generate both unemployment and relative-price pressures.

In particular, these frictions can produce a two-speed economy. The contracting sectors experience excess supply and falling wages, while the expanding sectors face labour shortages, rising wages, and cost pressures. Because prices and wages are sticky downward, the adjustment occurs primarily through relative-price increases in the expanding sectors, rather than nominal declines in the shrinking ones. Aggregate inflation can therefore rise even as overall unemployment remains elevated. This type of mechanism, formalised in Guerrieri et al. (2021), implies that sectoral demand shifts driven by technological change behave like cost-push shocks: inflation emerges not from excess demand, but from the need for relative-price realignment. In this circumstance, monetary policy should be cautious about responding to inflation driven by structural shifts. Attempts to suppress inflation through contractionary policy may slow or distort the reallocation process prolonging unemployment and output losses.

In summary, during the AI transition, monetary policy should exhibit measured flexibility. It may have to tolerate temporarily higher inflation to facilitate the necessary relative-price adjustments. This does not entail abandoning price stability as an anchor,

but rather recognising that some inflationary pressure is an unavoidable symptom of structural transformation, not a signal of excess demand.

Distinguishing structural from cyclical developments

For policymakers, the challenge lies in parsing structural dislocation from cyclical weakness in real time. Weak labour-market indicators—such as rising youth unemployment or sectoral job losses—may reflect either insufficient aggregate demand or technology-induced reallocation. Distinguishing the two is critical for setting the appropriate policy stance. Here are some useful indicators:

1. **Breadth and duration of unemployment.** A rising share of long-term or occupation-specific unemployment indicates structural adjustment, while a broad-based, short-duration increase suggests cyclical weakness.
2. **Job-finding and separation rates.** Cyclical downturns typically reduce both job-finding and separation rates. Structural change, by contrast, raises separations in AI-exposed sectors while job-finding rates in AI-complementary or expanding sectors remain robust.
2. **Vacancy-unemployment relationship (Beveridge curve).** A rightward shift—high vacancies coexisting with high unemployment—signals mismatch consistent with reallocation rather than demand shortfall.
3. **Sectoral dispersion of wage and price inflation.** Broader, demand-driven inflation will lift prices across most sectors; structural inflation will be concentrated in AI-complementary industries showing productivity and wage growth.

Monitoring these indicators can help central banks judge whether observed slack warrants cyclical accommodation or reflects structural reallocation requiring patience and targeted support.

Taken together, these considerations imply that AI will temporarily complicate the inflation-unemployment trade-off. In the long run, AI raises potential output (Y^*) and likely lowers the natural rate of unemployment (U^*), thereby easing inflationary pressures. But in the short run, the reallocation of labour and capital could produce a mix of sectoral inflation and aggregate slack—an environment where conventional policy rules based solely on aggregate inflation or unemployment may misfire.

Hence, monetary policy during the AI transition must operate with a structural lens: recognising that some inflation reflects healthy adjustment, and that stabilising relative prices may require tolerating temporarily higher headline inflation. Ultimately, the goal is

not to resist structural change but to keep aggregate demand close to potential while allowing the economy to reach its new equilibrium efficiently.

References

- Acemoglu, Daron, and David Autor. (2011). "Skills, Tasks and Technologies: Implications for Employment and Earnings." In *Handbook of Labor Economics*, Vol. 4, Part B, edited by Orley Ashenfelter and David Card, 1043-1171. Amsterdam: Elsevier.
- Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo. (2022). "Artificial Intelligence and Jobs: Evidence from Online Vacancies." *Journal of Labor Economics*, 40(S1): S293-S340.
- Acemoglu, Daron, and Pascual Restrepo. (2018). "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment." *American Economic Review* 108 (6): 1488-1542.
- Acemoglu, Daron, and Pascual Restrepo. (2019). "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal of Economic Perspectives* 33 (2): 3-30.
- Acemoglu, Daron, and Pascual Restrepo. (2020). "Robots and Jobs: Evidence from US Labor Markets." *Journal of Political Economy*, 128(6): 2188-2244.
- Aghion, Philippe, Benjamin Jones, and Charles Jones. (2019). "Artificial Intelligence and Economic Growth." In Ajay Agrawal, Joshua Gans, and Avi Goldfarb (Eds.), *The Economics of Artificial Intelligence: An Agenda*, pp. 237-290. University of Chicago Press.
- Agrawal, A., Gans, J., & Goldfarb, A. (2019a). *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press. <https://www.nber.org/books-and-chapters/economics-artificial-intelligence-agenda>.
- Agrawal, A., Gans, J., & Goldfarb, A. (2019b). Economic Policy for Artificial Intelligence. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The Economics of Artificial Intelligence: An Agenda* (pp. 139-159). University of Chicago Press.
- Athey, Susan and Fiona Scott Morton. (2025). "Artificial Intelligence, Competition, and Welfare," NBER Working Paper 34444.
- Autor, David H. (2015). "Why Are There Still So Many Jobs? The History and Future of Workplace Automation." *Journal of Economic Perspectives* 29 (3): 3-30.
- Autor, David, Caroline Chin Anna M. Salomons Bryan Seegmiller. (2022). "NEW FRONTIERS: THE ORIGINS AND CONTENT OF NEW WORK, 1940-2018." Working Paper 30389, National Bureau of Economic Research.
- Autor, David H., Frank Levy, and Richard J. Murnane. (2003). "The Skill Content of Recent Technological Change: An Empirical Exploration." *Quarterly Journal of Economics* 118(4): 1279-1333.

- Baily, Martin Neil, Erik Brynjolfsson, and Anton Korinek. (2023). *"Machines of Mind: The Case for an AI-Powered Productivity Boom."* Brookings Institution, May 10.
- Basu, Susanto, and John G. Fernald. (2008). "Information and Communications Technology as a General-Purpose Technology: Evidence from U.S. Industry Data." *Economic Review*, Federal Reserve Bank of San Francisco, 1–15.
- Brynjolfsson, Erik, R. Chandar, and L. Chen. (2025). "Canaries in the Coal Mine: Early Labor Market Effects of Generative AI." Stanford Digital Economy Working Paper, August 2025.
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. (2019). "Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics." In Ajay Agrawal, Joshua Gans, and Avi Goldfarb (Eds.), *The Economics of Artificial Intelligence: An Agenda*, pp. 23–57. University of Chicago Press.
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. (2021). "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies." *American Economic Journal: Macroeconomics*, 13: 333–372.
- Brynjolfsson, Erik, Li, D., & Raymond, L. (2025). "Generative AI at Work." *Quarterly Journal of Economics*, 140(2): 889–942.
- Calvino, Flavio, Daniel Haerle, and Sarah Liu. (2025). "Is Generative AI a General-Purpose Technology? Implications for Productivity and Policy." OECD Artificial Intelligence Papers No. 40.
- Chandar, R. (2025). "Tracking Employment in AI-Exposed Jobs." Tech. rep., Social Science Research Network.
- Czarnitzki, D., Fernández, G. P., & Rammer, C. (2023). Artificial intelligence and firm-level productivity. *Journal of Economic Behavior & Organization*, 211, 188–205.
- David, Paul A. (1990). "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox." *American Economic Review* 80 (2): 355–361.
- Dell'Acqua, Fabrizio, Edward McFowland III, Ethan R. Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Kraymer, François Candelon, and Karim R. Lakhani. (2023). *"Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality."* Harvard Business School Technology & Operations Management Unit Working Paper No. 24-013.
- Deming, David J. (2017). "The Growing Importance of Social Skills in the Labor Market." *Quarterly Journal of Economics*, 132(4): 1593–1640.
- Dillon, E., S. Jaffe, S. Peng, and A. Cambon. (2025). "Early Impacts of M365 Copilot." Tech. Rep. MSR-TR-2025-18, Microsoft.

- Eloundou, T., S. Manning, P. Mishkin, and D. Rock. (2024). "GPTs are GPTs: Labor Market Impact Potential of LLMs." *Science*, 384: 1306–1308.
- Felten, E., M. Raj, and R. Seamans. (2021). "Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses." *Strategic Management Journal*, 42: 2195–2217.
- Felten, E., M. Raj, and R. Seamans. (2023). "How will Language Modelers like ChatGPT Affect Occupations and Industries?" arXiv preprint arXiv:2303.01157.
- Filippucci, Francesco, Peter Gal, Katharina Laengle, and Matthias Schief. (2025). "Macroeconomic Productivity Gains from Artificial Intelligence in G7 Economies." OECD Artificial Intelligence Papers, No. 41.
- Gimbel, Martha, Molly Kinder, Joshua Kendall, and Maddie Lee. (2025). "Evaluating the Impact of AI on the Labor Market: Current State of Affairs." Oct 01. <https://budgetlab.yale.edu/research/evaluating-impact-ai-labor-market-current-state-affairs>.
- Gorodnichenko, Yuriy, and Oleksandr Talavera. (2017). "Price Setting in Online Markets: Basic Facts, International Comparisons, and Cross-Border Integration." *American Economic Review* 107 (1): 249–82.
- Green, Andrew. (2024). "Artificial Intelligence and the Changing Demand for Skills in the Labour Market." OECD Artificial Intelligence Papers, No. 14, OECD Publishing.
- Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Iván Werning. (2021). "Monetary Policy in Times of Structural Reallocation." *Proceedings of the Jackson Hole Symposium*.
- Handa, Kunal, Alex Tamkin, Miles McCain, Saffron Huang, Esin Durmus, Sarah Heck, Jared Mueller, Jerry Hong, Stuart Ritchie, Tim Belonax, Kevin K. Troy, Dario Amodei, Jared Kaplan, Jack Clark & Deep Ganguli. (2025). "Which Economic Tasks are Performed with AI? Evidence from Millions of Claude Conversations." arXiv preprint arXiv:2503.04761.
- Humlum, A. and E. Vestergaard. (2025). "Large Language Models, Small Labor Market Effects." Working Paper, National Bureau of Economic Research.
- Jones, Charles I. (2024). "The AI Dilemma: Growth versus Existential Risk." *AER: Insights*, 6(4): 575–590.
- Jorgenson, Dale W., Mun S. Ho, and Kevin J. Stiroh. (2008). "A Retrospective Look at the U.S. Productivity Growth Resurgence." *Journal of Economic Perspectives* 22 (1): 3-24.
- Korinek, Anton, and Maximilian Juelfs. (2022). "Preparing for the (Non-Existent?) Future of Work." In *Oxford Handbook of AI Governance*, edited by Justin Bullock et al., Oxford University Press. [Forthcoming].

- Korinek, Anton. (2023). "Scenario Planning for an A(G)I Future." *IMF Finance and Development Magazine*, 60(4): 30-33.
- Korinek, Anton, and Donghyun Suh. (2024). "Scenarios for the Transition to AGI." NBER Working Paper No. 32255.
- Kristina McElheran, Mu-Jeung Yang, Zachary Kroff, Erik Brynjolfsson. (2025). "The Rise of Industrial AI in America: Microfoundations of the Productivity J-curve(s)." Center for Economic Studies, US Census Bureau Working Papers, March 2025.
- Noy, S. and W. Zhang. (2023). "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence." *Science*, 381: 187–192.
- Peng, S., E. Kalliamvakou, P. Cihon, and M. Demirer. (2023). "The Impact of AI on Developer Productivity: Evidence from GitHub Copilot." arXiv preprint arXiv:2302.06590.
- Svanberg, M. S., Li, W., Fleming, M., Goehring, B. C., & Thompson, N. C. (2024). Beyond AI Exposure: Which Tasks are Cost-Effective to Automate with Computer Vision? Working Paper, Massachusetts Institute of Technology.
- Tomlinson, Kiran; Sonia Jaffe; Will Wang; Scott Counts; Siddharth Suri. (2025). "Working with AI: Measuring the Applicability of Generative AI to Occupations." arXiv preprint arXiv:2507.07935.
- Trammell, P. and Korinek, A. (2025). "Economic Growth under Transformative AI." NBER Working Paper 31815.
- Webb, M. (2020). The Impact of Artificial Intelligence on the Labor Market. Working paper.