

Checking for Mist**AI**kes

The hidden negative externalities in the AI
productivity debate

Discussion Paper
August 2025



Table of Contents

	Executive Summary	2
	AI productivity benefits are unclear and will take a long time	3
	AI creates extra work and generates new tasks	9
	AI creates new issues in the workplace and in the workforce	19
	AI leaves a lot of unanswered questions in the present	23
	The way forward requires clarity and consideration, not wishful thinking	24
	Conclusion	25

Executive Summary

Approaches to regulating AI have turned from discussions around safety, ethics and human rights, into ones around productivity, efficiency and worker outputs. This shift is being felt globally as technology firms gain traction with their lobbying efforts around the world. In Australia, where the government has yet to finalise its approach to regulating AI, the focus on productivity comes as the country experiences sluggish productivity growth.

But AI is not the magic bullet to a lukewarm economy. Although there are some genuine benefits being discovered, most of the bullish claims around productivity benefits have yet to be realised, more aspiration than inevitability.

Many claims of productivity benefits are derived from controlled, short-term 'lab' studies rather than real-world use cases, are produced or commissioned by technology companies who are agitating only for their own interests, or ignore significant considerations around negative externalities.

This report is a literature review around current AI productivity claims. It looks at studies which provide insights and analysis from real-world use cases, are from independent, credible institutions, and consider a holistic understanding of the AI environment, including negative effects and harms. This list is not meant to be exhaustive, rather provide key insights and compelling, evidence-based counterpoints to current claims around AI productivity.



There are four main themes which emerge from the literature:

- 1. Productivity benefits are uneven, unclear and will likely take many years to materialise**
- 2. In many cases, AI is creating extra workload, or generating new tasks as a result of its use**
- 3. Workers face new issues with AI use, including cognitive decline or debt, extra pressure and disadvantaging a new generation which could cause productivity issues in the medium to long term**
- 4. There are many unresolved issues – like worker displacement, unclear liabilities, privacy breaches and copyright breaches that remain unanswered and should be factored in**

This report recommends that further study and independent analysis is needed to qualify claims around the productivity benefits of AI, that negative externalities should not be ignored and instead be critical factors for consideration, and that AI is a national concern with diverse impacts requiring a holistic approach and not just filtered through a productivity lens.

AI productivity benefits are unclear

Many headlines capturing attention claim that AI will turbocharge productivity in astounding ways¹, will help workers boost their output and efficiencies by a factor of 10 or more², and that AI will transform the economy over the next few years.³

In fact, tech companies are using a sense of urgency and ‘fear of missing out’ as a key lever to influence governments and the public implying that not acting on AI now, and within the tech companies’ preferred terms of reference will be a cause of significant disadvantage.⁴ The real picture of course, is far more complex.

The Organisation for Economic Cooperation and Development (OECD) report titled ‘The Impact of Artificial Intelligence on Productivity, Distribution and Growth’ provides a comprehensive analysis of how AI impacts economic productivity. The report acknowledges AI’s promise as a potential transformative technology, but describes several complex factors which make predicting its impact on productivity unclear and multifaceted.⁵

¹ OpenAI & Mandala, 2025, AI in Australia, OpenAI’s Economic Blueprint

² Okemwa, K., July 2025, ‘Sam Altman says AI will make coders 10x more productive, not replace them – Even Bill Gates claims the field is too complex’, MSN Australia, <https://www.msn.com/en-us/money/careersandeducation/sam-altman-says-ai-will-make-coders-10x-more-productive-not-replace-them-even-bill-gates-claims-the-field-is-too-complex/ar-AA1Crkgg>

³ Altcheck, A & Perkel, S. May 2025, ‘Anthropic CEO says AI could wipe out half of all entry-level white collar jobs’, Business Insider, <https://www.businessinsider.com/anthropic-ceo-warning-ai-could-eliminate-jobs-2025-5>

⁴ Perez, S., July 2025, ‘Zuckerberg says people without AI glasses will be at a disadvantage in the future’ Techcrunch, <https://techcrunch.com/2025/07/30/zuckerberg-says-people-without-ai-glasses-will-be-at-a-disadvantage-in-the-future/>

⁵ OECD Publishing, April 2024, *The Impact of Artificial Intelligence on Productivity, Distribution and Growth, Key Mechanisms, Initial Evidence and Policy Challenges*

The report identified AI as a potential general-purpose technology, like the steam engine, electricity, computers and the internet before it. General purpose technologies are characterised by their ubiquitous nature, and as critical features within a broad range of economic and technical applications.

The report also posits AI could be different from previous general-purpose technologies in its ability to potentially self-improve and speed up the pace of innovation. This distinction comes with both opportunities and risks.

While the report found initial micro evidence of productivity boosts, the impact on longer-term and aggregate productivity remains uncertain, and are subject to several complex and unresolved factors, including whether the gains of early adopters will effectively trickle down to other users, whether labour reallocation due to AI can sustain aggregate growth, and whether broader social issues caused by AI will effectively harm productivity if not properly addressed. These are questions that require policy and will be influenced by political and social sentiment, which can determine the course of AI's development.

Aside from the complexity of measuring AI's impact at a macroeconomic level, and of the uncertain political and social factors which are likely to influence it, there are also known limiting effects to its development, such as the rising use of data and its need for commercially valuable knowledge. AI's reliance on high quality, proprietary, valuable data as critical ingredients for its development is a key battleground, as supporters argue that AI technologies should have unrestricted access to these data points⁶, while critics and owners/ data holders demand negotiation and compensation for the assets which they've worked hard to produce, and which have genuine and quantifiable commercial value.⁷ An existential question presents itself - as to whether AI development should be permitted, but at the cost of whole sections of the economy and at the expense of an enormous group of workers and their valuable intellectual property.

“An existential question presents itself - as to whether AI development should be permitted, but at the cost of whole sections of the economy and at the expense of an enormous group of workers and their valuable intellectual property.”

The report also describes how the adoption of AI technology to date is uneven, with larger companies and those in knowledge and information industries being early adopters, while smaller firms and other industries being laggards, widening productivity gaps.

⁶ Taylor, J., Aug 2025, 'Scott Farquhar thinks Australia should let AI train for free on creative content. He overlooks one key point'. Guardian Australia, <https://www.theguardian.com/technology/2025/aug/14/scott-farquhar-thinks-australia-should-let-ai-train-for-free-on-creative-content-he-overlooks-one-key-point>

⁷ Dudley-Nicholson, J., Apr 2025, 'Aussie authors launch bid to close the book on AI theft', AAP News, <https://aapnews.aap.com.au/news/aussie-authors-launch-bid-to-close-the-book-on-ai-theft>

The hyper concentration of AI development and AI technology owners within only a handful of technology companies may limit the broader diffusion of AI benefits, and create further risks. These risks include reducing the dynamic environment around AI owners, developers and users, increasing inequality and hampering inclusion, and broader social issues like intensified misinformation, biases, privacy violations and longer-term existential risks.




Figure 1 shows how general-purpose technologies like electricity and computers took decades to reach wide adoption and impact, in contrast with claims of AI transforming the economy in the next few years. Source: OECD report 'The Impact of Artificial Intelligence on Productivity, Distribution and Growth'

While the US President announced a tech firm-friendly and brash AI plan⁸, the country's most important economic institution paints a different picture. A paper from the US Federal Reserve called 'Generative AI at the Crossroads: Light Bulb, Dynamo or Microscope?' declares that the timeline for AI productivity will be "inherently slow" and will be "fraught with risk".⁹



⁸ Jamali, L. & Debusmann Jr, B. July 2025, 'Trump unveils AI plan that aims to clamp down on regulations and 'bias'', BBC News, <https://www.bbc.com/news/articles/c4g8nxrk207o>

⁹ Neil, Byrne, et.al., 2025, Generative AI at the Crossroads: Light Bulb, Dynamo or Microscope? Finance and Economics Discussion Series 2025-053, Washington, Board of Governors of the Federal Reserve System

“ The timeline for AI productivity will be “inherently slow” and will be “fraught with risk”. ”

The paper draws parallels between AI and previous transformative technologies like the lightbulb, suggesting that while there may be temporary short-term boosts in productivity, the longer-term impacts may diminish once the market becomes saturated.

It looks to history to remind that transformative technologies and their integration into everyday use across the broader economy is usually a slow and protracted process, dependent on many factors including workforce reconfiguration and people’s willingness to use the technologies.

The paper calls out that while there have been notable benefits to productivity in specific areas, such as writing, programming and customer services, it remains unclear whether these isolated successes will translate into widespread economy-wide applications.

It also provides a warning around the infrastructure and investment requirements needed to power AI. Building too quickly or impulsively could have “disastrous consequences”¹⁰ and likens it to the boom-and-bust cycles of the 19th century US railroad expansions which ultimately led to an economic depression.

It also gives warning around the availability of electrical power that is needed to run data centres and the infrastructure around AI.

Technology	Initial Impact
Domestication of plants	9000–8000 BCE
Writing	3400–3200 BCE
Iron	1200 BCE
Waterwheel	Early medieval period
Three-masted sailing ship	15 th century
Printing	15 th century
Factory system	Mid 18 th century
Steam engine	Late 18 th century
Railway	Mid 19 th century
Internal combustion engine	Late 19 th century
Electricity	Early 20 th century
Motor vehicle	Early 20 th century
Mass production, continuous process factory	Early 20 th century
Lean production	Late 20 th century
Computer	Late 20 th century
Internet	Late 20 th century

Source: Adapted from Lipsey, Carlaw and Bekar (2005).

Figure 2 shows examples of general-purpose technologies over history and the timeline of its impact. The US Federal Reserve describes that the timeline for AI productivity will be “inherently slow”. Source: Federal Reserve Board paper ‘Generative AI at the Crossroads: Light Bulb, Dynamo, or Microscope?’

Researchers from Columbia University argue that AI will be a “normal technology” much like electricity and the internet, and that integration into broader society and therefore impact on productivity will occur gradually and still within the control of humans.¹¹ Their paper attempts to present a view of AI as unremarkable, away from the more sensationalist narratives of AI both as a harbinger of utopia, and dystopia.

They suggest that AI’s impact will happen over decades, differentiating between different phases of AI development, from methodology, applications and adoption. Slow integration would also suggest that benefits to productivity will happen incrementally and not immediately.

¹⁰ Neil, Byrne, et.al., 2025, Generative AI at the Crossroads: Light Bulb, Dynamo or Microscope? Finance and Economics Discussion Series 2025-053, Washington, Board of Governors of the Federal Reserve System

¹¹ Narayanan A. & Kapoor S., Apr 2025, AI as Normal Technology, an alternative to the vision of AI as a potential superintelligence, Knight First Amendment Institute at Columbia University

The paper describes how AI diffusion in safety-critical areas is and will continue to be slow. For areas with important safety considerations, like health, law enforcement, or children's issues, AI becomes subject to the safety protocols and techniques used within those areas which are necessarily highly regulated and often using the same systems that have been in place for a long time. Progress in these areas will be slow as AI developers will not be able to properly anticipate many scenarios within those areas while testing and training is being done, given the complex real-world scenarios associated.

It also argues that diffusion is highly impacted by adoption, ultimately whether people can integrate AI into their everyday lives. Adoption is not just about availability, and even when new AI products become readily available, it still takes time for people to change their workflows and habits as well as test the risks involved in new products. Further, adoption is not just determined by individuals, but by the systems, organisations and institutions around them, which is an inherently slow process. The paper shows that this slow diffusion trend is consistent with previous general-purpose technologies, which took decades to fully materialise.

The researchers state that AI will continue to be at the hands of humans and human management. They reject the notion that AI will become a runaway agent, capable of self-determination, and instead describe how AI and its development, including any impacts on productivity will be determined and shaped by policy, regulations and social choices.




Figure 3 shows the various stages in AI development, and how AI impact materialises when its improvements are translated into applications across a broad section of the economy, with speed bumps/limit likely at every stage. Source: Knight First Amendment Institute at Columbia University paper 'AI as Normal Technology'

Another paper which reinforces the slow-paced diffusion and adoption of AI comes from the Oxford Review of Economic Policy. It looks back at previous general-purpose technologies such as steam, electricity and information and communications technologies like computers and smartphones and tries to extrapolate these to the impacts of AI.¹²

The paper shows that while general-purpose technologies eventually led to longer term productivity gains, this was only realised after decades of development and integration into the economy.

AI's productivity impact will depend on several different factors such as adoption rates, complementary investments and the broader economic environment.

The paper also highlights how job displacement and wage pressures could potentially lead to a reduction in labour's overall share of national income, thus complicating the impact of AI on productivity.

	Capital deepening	TFP	Total	% Whole economy
Steam (UK)				
1760–1830	0.011	0.003	0.014	5.6
1830–1870	0.18	0.12	0.30	19.0
1870–1910	0.15	0.16	0.31	29.2
Electricity (USA)				
1899–1919	0.04	0.06	0.10	5.6
1919–1929 (1)	0.07	0.07	0.14	3.6
1919–1929 (2)	0.07	0.30	0.37	9.5
1929–1941	0.04	0.16	0.20	8.0
ICT (USA)				
1974–1995	0.41	0.36	0.77	49.4
1995–2004	0.78	0.72	1.50	49.0
2004–2012	0.36	0.28	0.64	41.0

Note: In 1919–1929 estimate (1) does not take account of TFP spillovers but they are included in estimate (2).
Sources: Crafts (2004), Byrne *et al.* (2013), Crafts and Woltjer (2021).

Figure 4 shows milestones during the development of general-purpose technologies and the length of time it took, as well as capital, total factor productivity and percentage of the whole economy. Source: Oxford Review of Economic Policy, 'Artificial intelligence as a general-purpose technology: an historical perspective'

Therefore, while there is broad consensus in AI's potential to benefit productivity, the literature demonstrates that the road ahead and the road along the way is bumpy and uneven, with many uncertain and unpredictable factors to consider. Further, productivity benefits will likely take a long time, some decades before they fully materialise.

¹² Crafts, N., Sept 2021, Artificial intelligence as a general-purpose technology: an historical perspective, Oxford Review of Economic Policy, vol 37, issue 3

Looking back at previous waves of general-purpose technologies confirm this, and shows that the way forward can be fraught with risk, and is dependent on the choices made by both decision makers and the general populace of the day.

AI creates extra work and generates new tasks

Claims that AI boosts individual worker productivity and efficiency is far from unanimous, and is certainly to date, not a universal experience. In fact, many studies and trials suggest that AI can increase workloads, adding extra strain on workers and generating new tasks that were not previously required.

Workforce company Upwork surveyed 2,500 workers across the US, UK, Canada and Australia. Their research reveals that the majority of high expectations around AI boosting productivity come from C-suite leaders, and are disconnected from the reality of worker experiences, with 77% of employees reporting that AI has increased their workloads.¹³

“The majority of high expectations around AI boosting productivity come from C-suite leaders, and are disconnected from the reality of worker experiences.”

Further, despite mandates around the use of AI tools, around 47% of employees say they don't know how to achieve the desired productivity gains AI is supposed to deliver. The vast majority of employees (71%) are experiencing burnout and 65% are struggling with their employer's productivity demands.

This suggests that there is a huge gap in expectations between company owners and management who are optimistic and demand a boost in productivity through AI, and their employees who have to figure out how to deliver on this promise.

This disconnect between the promises of AI productivity and the reality of its delivery is a common theme.

When Australia's national science agency the Commonwealth Scientific and Industrial Research Organisation (or CSIRO) conducted a formal six-month trial of Microsoft CoPilot, they found mixed results, and that it was only helpful for simple, structured tasks. The trial was conducted with 300 participants, across the full spectrum of roles within the agency, ensuring a representative and diverse group. These included senior leaders and management, legal advisors, researchers, IT staff and operational staff. Knowledge workers such as those employed by CSIRO are prime candidates for AI 'exposure' and disruption.¹⁴

¹³ Upwork, Jul 2024, From Burnout to Balance: AI-Enhanced Work Models, <https://www.upwork.com/research/ai-enhanced-work-models>

¹⁴ Bano, Zowghi et.al., Dec 2024, Survey Insights on M365 Copilot Adoption, CSIRO

Before the trial, over 80% of participants expected CoPilot to help them save time, create efficiencies and enhance productivity. This was very much tempered post-trial, with only 21% of participants strongly agreeing that productivity improved, while 11% strongly disagreed.

Further, pre-trial, there were few concerns raised around ethics and ethical use. Post-trial, participants flagged concerns around data privacy, unauthorised access, security risks and biases, which correlated to lower willingness to recommend the tool.

While CoPilot showed some promise in assisting with certain basic, repetitive tasks, the result ultimately showed limited productivity benefits, and instead raised more awareness around ethical concerns in AI use. The results also showed the huge gaps in expectations versus reality when it comes to expected productivity benefits.

Coders and programmers are another key group of knowledge workers in assessing whether AI will bring real productivity benefits. As a technical and digitally savvy cohort, they are primed to use technological tools to assist with their jobs, and are naturally receptive towards software and new tools. But studies and research on AI impacts on coding show that AI in fact increase their workloads and hinder productivity.

In a randomised controlled trial conducted by frontier AI research agency METR, developers with

moderate (average 5 years) experience conducted 246 real tasks, some with AI assistance and some without. Prior to the trial, developers believed AI would reduce task completion time by 24%. The results were the opposite, with AI increasing completion time by 19%, contradicting both the developer's expectations as well as popular forecasts.¹⁵

“ *Developers believed AI would reduce task completion time by 24%. The results were the opposite, with AI increasing completion time by 19%.* ”

The study authors noted several factors which contributed to the time delays. This includes the fact that experienced developers were highly familiar with their own codebases, and found AI's superficial knowledge of these as not helpful. AI also produced results that were not acceptable after code review, and resulted in developers spending around 9% of their time cleaning up AI generated code. AI also lacked contextual awareness of tasks and codebases, forcing developers to spend time filling in the gaps or correcting misaligned suggestions.

This field experiment provides a compelling real-world counterpoint to lab studies showing AI productivity boosts. In the hands of seasoned developers working on mature, complex databases, AI added work rather than alleviated it.

¹⁵ Becker, Rush et.al., 2025, Measuring the Impact of Early -2025 AI on Experienced Open-Source Developer Productivity, Model Evaluation & Threat Research (METR)




Figure 5 shows how developers substantially over-estimated the time savings AI would bring, instead tasks took 19% more time to complete. Source: METR study, 'Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity'




Figure 6 shows the difference in forecast developers made between tasks that include AI, and tasks that don't, with developers over-estimating non-AI task times, but underestimating AI task times in terms of completion. Source: METR study, 'Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity'




Figure 7 shows the cumulative distribution functions of observed implementation times showing nearly all quantiles of AI-allowed tasks taking longer. Source: METR study, 'Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity'

Software knowledge company Stack Overflow conducts an annual survey around AI. This is one of the most comprehensive developer surveys available, which polls over 65,000 professional developers.

It found that while AI tools boost initial code output, they often impose a 'productivity tax' resulting from the need to debug AI generated code that is "almost right" but not quite. This generates a significant hidden cost in time and effort from developers needing to review and/or redo this code.¹⁶

While adoption of AI tools for developers were high with around 76% of developers using them, 66% have reported that their productivity has slowed due to additional work needed to validate and correct AI's outputs.

The biggest frustration developers had was with dealing with "AI solutions that are almost right, but not quite" at 66%, and debugging AI generated code that is more time-consuming (than writing their own) at over 45%.

¹⁶ Stack Overflow, 2025, 2025 Developer Survey, AI section, <https://survey.stackoverflow.co/2025/ai>




Figure 8 shows the top frustrations developers have when using AI. Source: Stack Overflow survey 2025

Atlassian is one of Australia's largest and most influential technology companies. They are also one of the companies who have publicly declared AI's productivity benefits. Some controversy surrounded this move as a mass redundancy announcement was made by founder Mike Cannon-Brookes due to AI around the same time as they announced their own company focus on the technology.¹⁷

As part of their 2025 State of Developer Experience Report, Atlassian uncovered an 'unexpected paradox'. While their developers report saving time due to AI tools, they also reported losing more time, up to more than ten hours a week due to "organisational inefficiencies".¹⁸ The time savings found were mostly with AI-supported coding tasks, but the vast majority of developer work, including documentation, meetings, group coordination and administration were still inefficient where AI offered little support. This shows that AI cannot offer widespread relief and still experience other organisational hurdles. AI also cannot be disentangled from broader workforce systems and processes which continue to suffer from inefficiencies.

The report also highlighted a growing disconnect between management and everyday staff, which is a common challenge around the AI productivity debate with conflicting expectations between the two groups.

While AI tools appear to have some benefits to productivity in the Atlassian report, these are being neutralised or negated by entrenched organisational obstacles. In fact, they may amplify pressure by raising expectations around aspirational efficiencies without resolving underlying systemic issues.

This shows that even the most technical and AI-savvy companies like Atlassian have yet to fully benefit from the promises of AI productivity.

¹⁷ Seeto, T., July 2025, 'Atlassian billionaire spruiks AI as CEO fires 150 workers over video in scary digital revolution, Yahoo Finance, <https://au.finance.yahoo.com/news/atlassian-billionaire-spruiks-ai-as-ceo-fires-150-workers-over-video-in-scary-digital-revolution-052623412.html>

¹⁸ Atlassian, 2025, State of Developer Experience Report 2025, <https://www.atlassian.com/teams/software-development/state-of-developer-experience-2025>

“ This shows that even the most technical and AI-savvy companies like Atlassian have yet to fully benefit from the promises of AI productivity. ”




Figure 9 shows that half of developers lose more than 10 hours a week due to organisational inefficiencies, despite use of AI tools.

Source: Atlassian State of Developer Experience Report 2025

AI's impact on schools is another notable area, with many unclear and unanswered questions on what it means for students and their learning, and teachers' roles in that process. One area of assessment is whether teachers can gain efficiencies from using AI much like other white-collar knowledge workers do, with document reviews, document creation, summaries and administration.

A study conducted with teachers across Australia and Sweden focused on their use of AI for school-related administrative tasks. Many teachers in the study found that using AI tools did not reduce the time they spent on administrative tasks. In some cases, it actually increased their workload.¹⁹

This was due to the large amount of effort required in reviewing, correcting and sometimes completely re-working AI outputs that were sub-par. Teachers often found that AI outputs provided false information, dated sources, and poor-quality replies, particularly for those with non-English language requirements.

¹⁹ Selwyn, Ljungqvist et.al., Apr 2025, When the prompting stops: exploring teachers' work around the educational frailties of generative AI tools, Lund University

Teachers also found that AI had a poor grasp of the detailed requirements of school curricula, and were too generous with student markings and feedback, not capable of the critical thinking and analysis required and expected.

The study demonstrates that far from assisting teachers, let alone replacing them, AI has been found to be a burden to teachers, creating extra work by not being able to fully grasp the concepts and requirements of their daily tasks.

A very significant longitudinal economic study from the University of Chicago and the University of Copenhagen looked at data from 25,000 employees across 7,000 workplaces in Denmark over two years. These included 11 occupations 'exposed' to AI such as accountants, customer support specialists, financial advisors, HR workers, IT support, journalists, legal professionals, office clerks, software developers and teachers. Detailed survey data were linked with enterprise records around hours, earnings and wages to paint a detailed, thorough picture of real-world impacts of AI on jobs.

Researchers Anders Humlum and Emilie Vestergaard concluded that "AI chatbots have had no significant impact on earnings or recorded hours in any occupation". Their findings challenged "narratives of imminent labor market transformation due to Generative AI".²⁰

“ AI chatbots have had no significant impact on earnings or recorded hours in any occupation. ”

Despite large adoption rates (47% of workers using AI unprompted), which significantly increased after employers encouraged usage (up to 83%), the economic impact was negligible. After two years of AI integration and widespread workplace adoption, outcomes remain essentially unchanged.

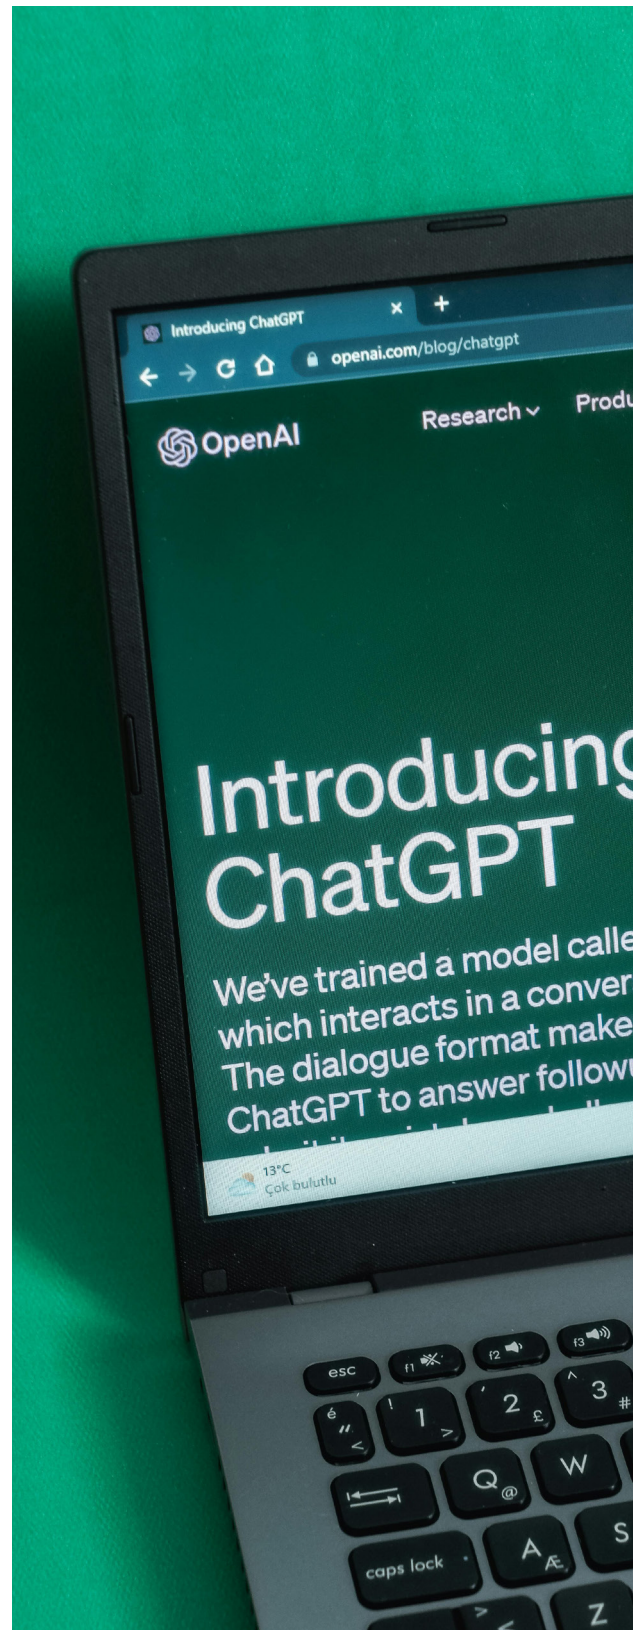


²⁰ Humlum, A. & Vestergaard, E. May 2025, Large Language Models, Small Labor Market Effects, National Bureau of Economic Research, Massachusetts Institute of Technology

Workers do self-report time savings averaging 25 minutes a day, but calculating for frequency and per-day savings this only translated to roughly one hour a week. The study also found that these did not translate into higher earnings. These deeper considerations are important, as many studies which claim productivity benefits are 'lab-controlled' experiments that don't measure realistic, real-world impacts.

The study also highlights the significance of workplace policy and employer intervention for workplace AI integration. When encouraged by employers, adoption rose by 36%. When companies invest in training and configure AI systems internally, results improve across all metrics. Employer intervention also helped to narrow demographic gaps. Without support and encouragement, women were less likely to adopt AI than men, and older workers and their adoption rates narrow by 40%. This underscores the significance of human policy and intervention in any AI workplace initiatives, which challenge the narratives around AI as inevitable or automatic.

These studies provide compelling, evidence-based counter-claims around headlines declaring AI productivity as inevitable across workplaces. With diverse and significant methodologies used across different sectors in different countries, these studies show that productivity benefits are either negligible or don't translate into real-world impacts. Outcomes are further influenced by the complexity and reality of workplaces, including workplace policies around adoption, enterprise configuration, technical specialisations, and employee assistance.



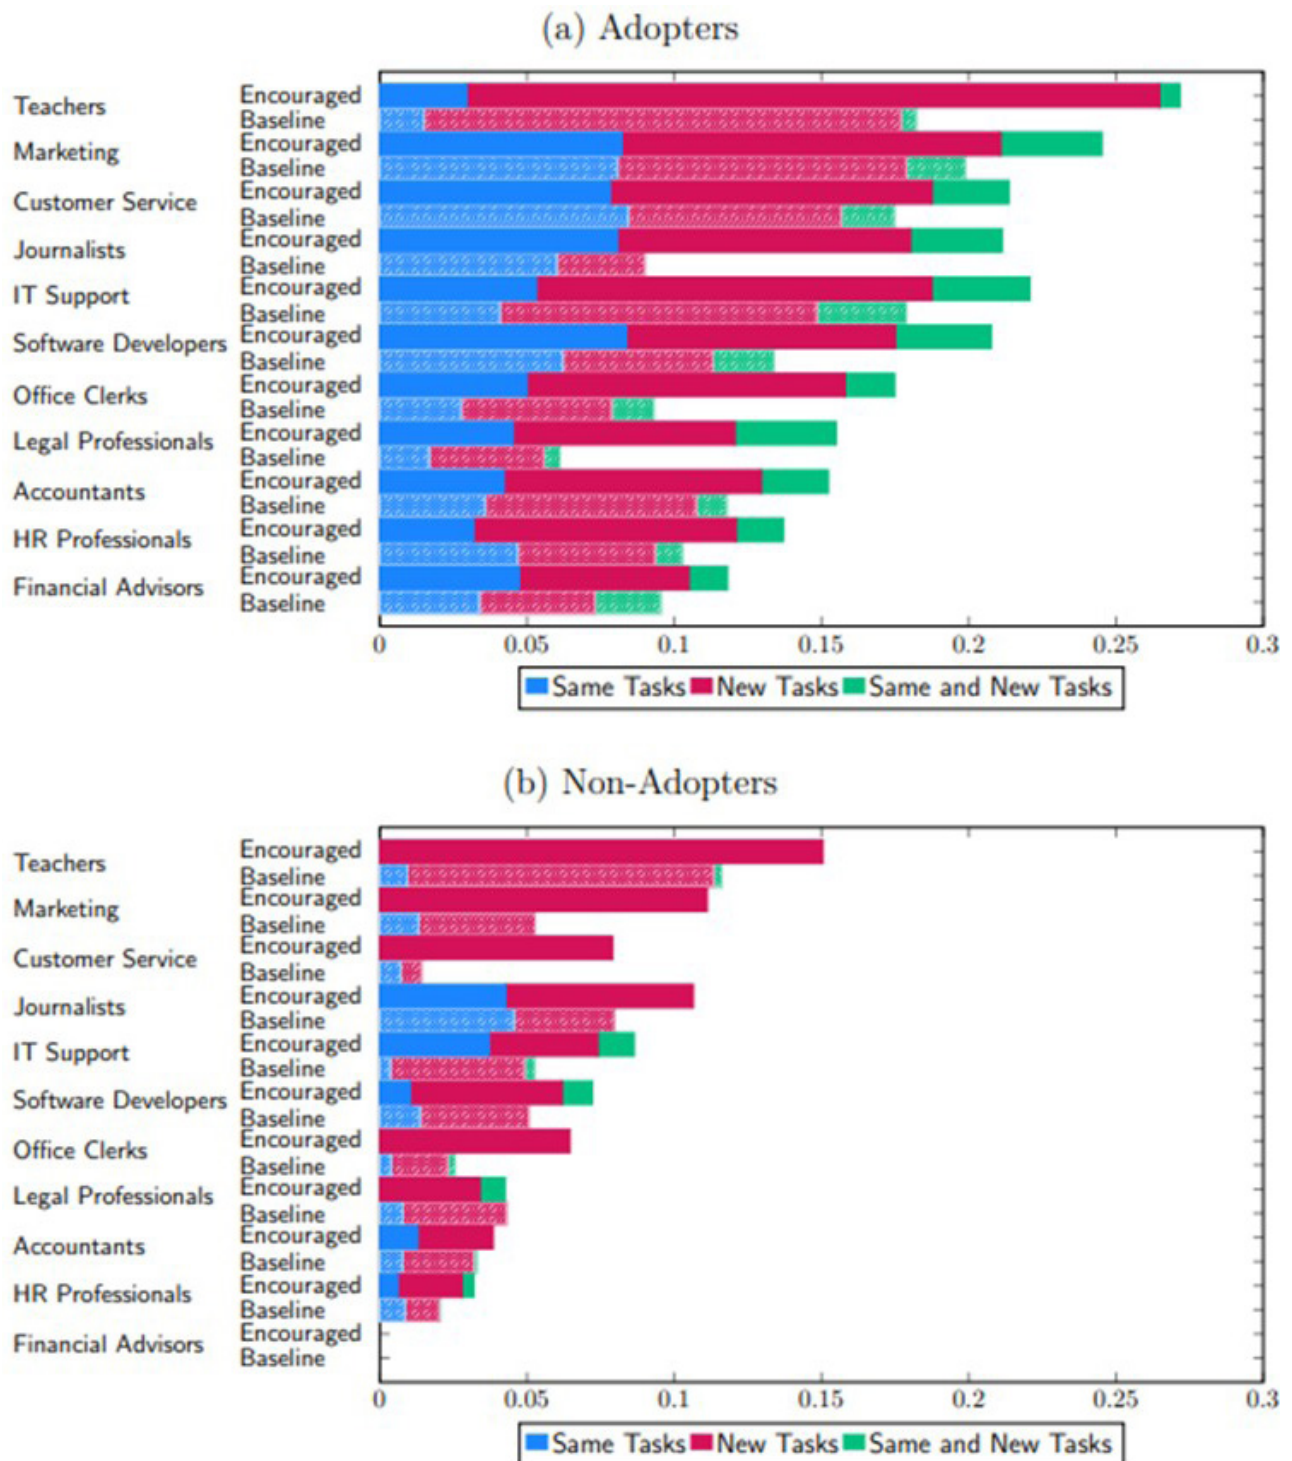



Figure 10 shows that AI chatbots have led to the creation of new tasks across all 11 occupations in the study. Source: National Bureau of Economic Research, 'Large Language Models, Small Labor Market Effects'

(a) Average Impacts (Dynamic Diff-in-Diff)



(b) Occupation-Level Impacts (Pooled Diff-in-Diff)




Figure 11 shows the negligible impact of AI tools in earnings and hours across occupations. Source: National Bureau of Economic Research, 'Large Language Models, Small Labor Market Effects'

AI creates new issues in the workplace and in the workforce

There are several unintended consequences created by AI adoption in the workforce which could negatively impact productivity.

Analysis from the Wharton School at the University of Pennsylvania highlights the 'AI Efficiency Trap' which creates heightened expectations from employees around their productivity and outputs, creating perpetual pressure with several adverse impacts on human labour.²¹

The AI Efficiency trap increases employee workloads and stress as employees are expected to maintain or increase output without corresponding adjustments to compensation or work/life balance. The constant pressure may lead to diminished job satisfaction and motivation as employees feel overwhelmed and undervalued.

“ *The AI Efficiency trap increases employee workloads and stress as employees are expected to maintain or increase output without corresponding adjustments to compensation or work/life balance.* **”**

As each real or perceived productivity improvement becomes integrated and locked in, this improvement becomes the new baseline. Deadlines compress, volumes increase, complexity expands, all while employees increasingly rely on AI tools, and may not receive extra support or even training for this 'new normal'. This could also entrench harmful new standards around digital overload and perpetual availability, where workers feel that they need to be online at all times, and available to their workplaces without material breaks.


The AI Efficiency trap enables the classic "Red Queen" hypothesis,²² adopted from studies in evolutionary biology, where continuous and accelerating adaptation becomes essential to stay competitive and stave off extinction. The concept comes from the classic story of Alice in Wonderland, where the Red Queen informs Alice that "here, you see, it takes all the running you can do to keep in the same place".²³ Thus the relentless pace of adaptation and 'innovation' becomes a race of no return.

As AI becomes embedded and workers feel more pressure to rely on them, the 'agency decay' phenomenon takes place, where workers may feel less capable of independent judgement, even in areas or tasks where they previously have not needed AI assistance.

²¹ Walther, C. June 2025, 'The AI Efficiency Trap: When Productivity Tools Create Perpetual Pressure', Knowledge at Wharton series from the Wharton School of the University of Pennsylvania

²² Science Direct, 2025, Red Queen Hypothesis, <https://www.sciencedirect.com/topics/earth-and-planetary-sciences/red-queen-hypothesis>

²³ Carroll, L. 1865, 'Alice's Adventures in Wonderland', Macmillan Press



Red Queen Hypothesis: Predator and prey must be constantly evolving to avoid extinction.

Figure 12 illustrates the concept of the 'Red Queen Hypothesis' where species need to keep evolving or face extinction'.

Source: Wikipedia

The decay in human agency and skill is validated by a landmark study from the Massachusetts Institute of Technology, which looked at the impact of using ChatGPT for writing and how this affected users' neural engagement, cognitive performance and sense of ownership over their work.

The findings are startling and describes the concept of "cognitive debt" whereupon an overreliance on AI reduces a person's internal cognitive engagement and hampers processes related to learning such as memory encoding, critical reflection, and sense of ownership over tasks.²⁴

“ An overreliance on AI reduces a person's internal cognitive engagement and hampers processes related to learning such as memory encoding, critical reflection, and sense of ownership over tasks. ”

These findings have significant implications for education, information workers and the knowledge economy, as well as productivity in general. The study found that frequent use of AI for writing and information retrieval tasks may undermine critical thinking and deep learning, erode long term memory retention and mastery, and impair the internal cognitive mechanisms required for creativity and original thought. These effects were found to persist even after stopping AI use.

²⁴ Kosmyna, N., Hauptmann, E. et.al. 2025, Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task, Massachusetts Institute of Technology Science Direct, 2025, Red Queen Hypothesis, <https://www.sciencedirect.com/topics/earth-and-planetary-sciences/red-queen-hypothesis>




Figure 13 shows that the percentage of people using AI tools struggled to quote anything from their work, far more than those participants who didn't use AI. Source: Massachusetts Institute of Technology, 'Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task'




Figure 14 shows the difference in feeling a sense of ownership over their work between those who used AI and those who didn't. Source: Massachusetts Institute of Technology, 'Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task'

As well as undermining our cognitive abilities, which impacts workforces and workplace productivity, AI driven automation may also create a generation of lost workers, who have been displaced by AI at the critical entry levels, and may harm productivity in the longer term.

A study from the University of Navarra highlights that while AI-driven automation may offer immediate, short-term productivity gains through automation of entry-level tasks, it also risks harming long-term economic growth by interrupting the transfer of tacit knowledge between junior and senior workers.

Skills which are typically learned through mentorship and apprenticeship, hands-on experience and direct engagement with more senior workers are potentially being lost by displacing junior staff for AI, impacting the employee pipeline and affecting workforce planning.²⁵

“*The skills which are typically learned through mentorship and apprenticeship, hands-on experience and direct engagement with more senior workers are potentially being lost by displacing junior staff for AI, impacting the employee pipeline and affecting workforce planning.*”

AI is best in automating repeatable, routine tasks, those typically given to junior workers. By removing this layer, AI reduced the opportunities for junior workers to learn these skills and gain valuable experience. This risks undermining the interpersonal, experiential learning pathways required to progress with their jobs and careers. Thus, it impacts the pipeline and progression of workers throughout companies' lifecycles, and interrupts medium to longer-term productivity.

Modelling done by the study predicts annual declines in long-term growth of up to 35%, depending on the extent and severity of automation initiatives. Over time, these losses accumulate significantly.

The study provides a cautionary tale against accepting the short-term claims of AI-driven productivity without proper considerations around workforce planning, succession and training.

Integrating a general-purpose technology like AI will not be a simple process and is bound to create new issues and challenges. Already, there is evidence of issues around the AI efficiency trap, the loss and decay of human agency, cognitive debt and decline, and creating a whole new generation of displaced workers which will ultimately dampen productivity growth and benefits. These issues need to be addressed in order to gather the full picture of AI's impacts on productivity.

²⁵ Ide, E., Aug 2025, Automation, AI and the Intergenerational Transmission of Knowledge, University of Navarra

AI leaves a lot of unanswered questions in the present

Beyond the specific topics highlighted in earlier sections, which provide detailed and evidence-based counter claims against AI-driven productivity, there are many significant and complex questions which remain unanswered around the impact of AI in society, the workforce and productivity overall.

The OECD report provides some key themes:²⁶

- Jobs displacement and task redefinition could impact entire industries, causing widespread disruption and worker disaffection in the near term. Even if workers retain their roles, task redefinition could see many tasks being automated, reducing job satisfaction and underemployment.
- Wage inequality could be an issue as workers with AI skills or who are able to navigate changes successfully see their wages increase, while workers whose jobs or tasks become automated experience significant stagnation or decline
- Management by algorithm, or through AI could lead to outcomes which while in theory seem productive and efficient, could in reality mean increased work intensity, reduced worker autonomy, and more difficult working conditions for employees without any humans to negotiate with or appeal to.
- Existing biases and discrimination could be exacerbated or increased by AI, as they become codified in automated systems around hiring, performance management and evaluation, promotions and incentives, and termination, undermining workplace diversity and inclusivity
- Uneven adoption and regional/global disparities could be significant with AI adoption and skills concentrated in certain regions or sectors, leaving those in low-or-middle income regions or countries, as well as smaller less prepared firms at significant disadvantage.

Furthermore, there are live debates around copyright breaches²⁷, privacy breaches²⁸, funding and investments, regulatory capture²⁹, vested interests and other complex topics which are influencing the AI landscape in the here and now. The outcome of these debates remains uncertain, with diverse groups vying for influence and lobbying for their own approaches.

²⁶ Ide, E., Aug 2025, Automation, AI and the Intergenerational Transmission of Knowledge, University of Navarra

²⁷ Riga, J. Jun 2025, 'Disney and Universal sue Ai firm Midjourney for copyright infringement', ABC Australia, <https://www.abc.net.au/news/2025-06-12/disney-and-universal-sue-ai-firm-midjourney-for-copyright/105407174> Science Direct, 2025, Red Queen Hypothesis, <https://www.sciencedirect.com/topics/earth-and-planetary-sciences/red-queen-hypothesis>

²⁸ Guan, L. Feb 2025, 'Widespread GenAI use raising data breach concerns', ARN, <https://www.arnnet.com.au/article/3827334/widespread-genai-use-raising-data-breach-concerns.html>

²⁹ Williams, t. Mar 2025, 'Big Tech wants Trump's help to weaken Australian laws', Information Age, <https://ia.acs.org.au/article/2025/big-tech-wants-trumps-help-to-weaken-australian-laws.html>

This regulatory, political, and commercial uncertainty will play a material impact on the productivity debate which follows and the broader economic landscape that results. In other words, AI productivity benefits are not a forgone conclusion, and are not inevitable. Whether it happens or not, and to what extent it is done fairly and equitably remains to be seen.

The way forward requires clarity and consideration, not wishful thinking

The debate around AI's impact on productivity has intensified, with Australia seeing it as a critical pillar in lifting its productivity issues. However, AI is a general-purpose technology which will have wide impact across the country, the region, and indeed globally. It also creates many new issues and exacerbates existing ones, with diverse impacts across broader society.

Looking at AI through a productivity filter alone, and funnelling Australia's AI regulatory scheme under the lens of productivity is short-sighted and ignores significant negative externalities which could result.

Instead, AI should be treated as a national issue with national scope, governed by a dedicated Australian AI Act to account for its wide impact and complexities.

More work needs to be done to properly determine not only AI's real impact on productivity, but across Australian society overall.

Policy considerations need to include diverse voices, in particular from the public and workers who will be most impacted by AI, and in many cases, will be the ones who will have to deliver on AI's promises.

Decision makers in government and industry should take care not to be dazzled by the promises of incumbent Big Tech companies, who are ultimately only agitating for an environment which serves their interests, and not necessarily those of the Australian public.

Some significant work has already been done from the last two years in the formation of Australia's AI regulatory regime. These should be incorporated and implemented, as valuable work which considers AI safety, AI ethics and human rights to produce an AI landscape that is responsible and serves the needs of all Australians, not just technology companies.

A dedicated Australian AI Act, governed by a central body like an AI Commissioner will ensure that the complexity of AI's impacts are captured, its harms are mitigated and its benefits are evenly distributed.

Conclusion

Claims around the productivity benefits of AI needs a reality check. Studies from top universities, software companies, government agencies and research organisations provide compelling evidence that AI's impact on productivity is complex, unclear and uneven. There are new issues and harms resulting from AI which should be factored into AI regulation, and not simply ignored as required by-products or negative externalities.



Report author:

Jordan Guiao, Director of Responsible Technology at Per Capita's Centre of the Public Square



Contact:

Inquiries pertaining to this report may be directed to Jordan Guiao at j.guiao@percapita.org.au



Website:

Learn more about Per Capita's Centre of the Public Square here: <https://percapita.org.au/>
