



UTS

Human
Technology
Institute

e51
INSTITUTE

AI, Productivity, and Australia's Choice of Regulatory Framework

September 2025

Key Messages





As the Productivity Commission has recognised in its *Interim Report on Harnessing Data and Digital Technology*, regulatory certainty is a foundational enabler of AI-driven productivity. Without credible, durable policy signals focused on responsible and productive AI use, firms have weaker incentives to invest in the data, infrastructure, applications, complementary capabilities and organisational change that productivity requires.

Regulatory ambiguity delays adoption, may distort investment, and risks driving unsustainable or low-value reorganisation. But not all regulatory approaches support productivity equally. Australia faces a strategic choice between:

- a deliberately permissive approach to regulation;
- a prescriptive approach focused on AI-specific regulation; and
- a pragmatic approach that prefers technology-neutral regulation.

Although each claims to give markets confidence, they make different trade-offs in terms of investment incentives, risk mitigation, and diffusion, and will shape firm behaviour across the five productivity pathways. We suggest that the pragmatic, technology-neutral approach to regulation provides more sustainable levels of regulatory certainty.



Authors of the report:

e61 Institute

Michael Brennan, Theo Gibbons, Silvia Griselda,
Greg Kaplan, Matthew Olckers and Ewan Rankin

UTS Human Techonology Institute

Nicholas Davis, Edward Santow and Sarah Vallee

Contents

Executive summary	06
1. AI-driven productivity growth and why it matters	08
2. Accessing these gains depends on getting regulation right	10
3. Regulatory settings interact with the channels of productivity growth from AI	12
4. Not all regulatory approaches support productivity equally, and not all generate certainty	14
5. The path forward: reducing regulatory ambiguity to unlock AI-driven productivity	18

Executive summary

1. AI's productivity potential is real but contingent.

AI is a general-purpose technology that offers at least five interlinked pathways to boost productivity, from task-level augmentation to knowledge generation, reorganisation and economy-wide reallocation.

However, these gains are not automatic. Realising AI's full productivity dividend will likely require sustained complementary investment in skills, capital, and organisational change.

2. AI investment is sensitive to uncertainty in the regulatory context.

AI requires large, complementary and partly irreversible investments with long payback periods. Even small rule changes can materially change expected returns and design choices.

Without a clear national regulatory approach, firms face legal ambiguity on data use, liability, intellectual property, employment law and professional obligations.

Investment is sensitive to both the quality of rules and their predictability over time, particularly in regulated sectors where compliance confidence is essential.

Uncertainty raises perceived risk on long-horizon, cross-functional AI investments. It delays and distorts capital spending and encourages short-termism.

Uncertainty increases the risk of coordination failure: firms may wait for others to act, mitigate risks inefficiently, or design systems to align with foreign regulatory regimes. The result is less experimentation, slower diffusion, and missed productivity gains.

3. Not all regulatory approaches support productivity equally. Australia faces a strategic choice among different models.

Three regulatory models are emerging globally. These different approaches carry distinct trade-offs for AI adoption, risk mitigation, and long-term investment:

- **Deliberately permissive approach:** A minimal-intervention approach that refrains from AI-specific rules and relies on existing law and market behaviour. It can promote rapid AI adoption, but may under-regulate potential harms, and lead to long-term uncertainty and risks prompting reactive future regulation.
- **Prescriptive, AI-specific approach:** An interventionist approach, focused on AI-specific regulation for existing high-risk domains or use cases. Given rapidly evolving AI capabilities, it risks becoming inflexible and burdensome as technologies and use cases evolve.
- **Pragmatic, technology-neutral approach:** An approach that favours the application of technology-neutral law, and uses guidance, standards and stress testing to enhance certainty and enable a nationally consistent approach. The effectiveness of this approach in achieving regulatory certainty depends on consistent interpretation and proactive, well-resourced regulators, and a willingness to depart from technology-neutral regulation only where there is a strong case for doing so.

The chosen model will shape incentives, risk management and diffusion across the five productivity pathways. Durable certainty is created by credible signals, a clear institutional posture and the capacity to evolve, enabling firms and workers to invest with confidence when governance is predictable.

4. Australia should act decisively to reduce regulatory ambiguity if it is to seize the productivity benefits of AI.

While Australia's current legal environment is not lawless, its lack of direction and credible commitments has created real uncertainty for firms and workers.

International approaches offer lessons but not blueprints; Australia must define its own strategy. The goal is not to predict every use case, but to create regulatory certainty through clear goals, institutional capability, and credible application over time.

Reducing regulatory ambiguity can be seen as a strategic lever for long-term growth: at both the firm and economy levels, governance should not be positioned as a constraint on AI, but the foundation for investment in productive transformation.

A large, stylized graphic of the letters 'Of' in a dark teal color, positioned in the upper left quadrant of the page. The background is a solid teal color with a pattern of vertical lines in varying shades of teal and light blue, creating a digital or data-like aesthetic.

AI-driven productivity growth and why it matters

Improving productivity—doing more with less—is critical to achieving economic growth and rising living standards. Sustained productivity growth means that businesses can pay higher real wages, workers can enjoy higher incomes or choose to work less, consumers can access better-quality products at lower prices, and governments can deliver improved public services without raising tax rates.


These pathways translate higher productivity into some of the key building blocks of wellbeing: the consumption of more and better things, including improved public services, and the option to enjoy more leisure.

Productivity is about working smarter, not longer. It is how much we produce from a given amount of inputs, such as hours of labour and capital resources like machinery, buildings and information technology.

AI represents a broad collection of increasingly capable technologies that can help individuals and organisations work smarter. It closely resembles a ‘general-purpose technology’, one with broad applicability across the economy, spurs complementary innovations, and transforms how economic activity is organised. Historical examples include electricity, the steam engine, and the internet¹.

An emerging body of research points to five principal channels through which AI could contribute to productivity growth. Each is both promising and uncertain, with the degree of uncertainty generally increasing in the order that the channels are described below.

Channels for how AI could impact productivity



1. Labour-improving tools

AI can act as a new form of capital that complements workers by helping them perform their existing tasks more efficiently and effectively. Both narrow (analytical AI) and more recent AI tools like coding assistants or generative AI chatbots are increasingly widely used, and new tools are in development.

Consider a software developer working on a new app. Previously, writing each line of code required careful thought and manual typing. With AI coding tools, the developer now quickly reviews AI-generated suggestions and integrates them with a single keystroke.

2. Automation of tasks

AI can also act as a form of capital that substitutes for labour. It has the potential to fully automate tasks previously done by people, allowing firms either to produce the same output with fewer workers or to expand output. In addition, AI can improve the efficiency of already-automated tasks. It is possible that AI will automate some relatively complex cognitive and non-routine tasks, unlocking broader productivity gains than previous waves of automation².

3. Reorganisation of processes and firms

AI can enhance productivity by enabling firms to restructure how work is organised and how internal resources are deployed. Beyond automating tasks, AI may prompt changes in workflows, decision rights, information flows, and team structures, thereby allowing firms to experiment with new ways of operating, something already occurring in many sectors, particularly manufacturing. For example, faster prototyping with AI tools can support more iterative and decentralised product development, while AI-enabled forecasting or decision support can shift how firms coordinate across functions. These forms of organisational redesign do not replace labour, but instead aim to unlock greater efficiency, flexibility, and responsiveness, especially when paired with investment in complementary systems, training, and management practices. This reorganisation has driven much of the productivity gains for previous general-purpose technologies³.

4. Sector disruption

The economy-wide effect of AI depends not only on productivity growth within individual firms, but also on reallocation: how market share shifts toward more efficient AI adopters. To be sure, in a best-case scenario there is widespread adoption of AI. But inevitably, some firms will adopt AI more effectively and tend to grow faster, increasing average productivity across the economy through a compositional shift. AI may also increase the entry rate of new firms by reducing barriers to entry, and these new firms may adapt to AI particularly well, further contributing to productive reallocation⁴.

5. Greater knowledge and innovation

AI-driven invention is the most speculative—but potentially the most transformative—channel of potential productivity gains. Even small improvements in idea generation could compound over time into faster GDP growth⁵. New ideas are the most powerful ingredient of long-run economic growth because they can be ‘non-rivalrous’ — i.e. many people and firms can use new ideas without prejudice to their use by others⁶. There is evidence that new ideas have been harder to find in recent decades, requiring a greater number of researchers to achieve the same rate of innovation in technology, medicine and agriculture⁷. There is the possibility that AI could compensate for this and lead to new breakthroughs that create broad public benefit.

02

Accessing these gains depends on getting regulation right

How AI is regulated in Australia will have implications for our ability to achieve long-term productivity gains and minimise harms. Two dimensions are important: the quality of regulatory approach (how well it addresses well-defined harms at least economic cost), and its predictability over time.

Regulatory quality and predictability are linked: a high-quality regulatory framework is more likely to prove enduring over time. But there can be trade-offs in the context of fast-evolving technology like AI. The optimal (or ideal) rules could shift rapidly with technological change, while firms and workers need a degree of stability in approach. Hence, regulatory certainty is valuable in its own right.

This need not imply rigid maintenance of specific rules, but rather a framework within which regulatory decisions adhere to consistent and well-understood principles that align with appropriate standards and regimes.

Accessing the productivity benefits described above requires significant investment by firms, and with it, a willingness to take on risk. In the past, in addition to the direct costs of investing in capital equipment and operating costs of new machinery, general-purpose technologies have required significant complementary investments in new processes and skills before productivity saw large gains⁹. These include retraining workers, redesigning workflows, upgrading capital equipment, and adapting organisational structures and management practices.

General economic uncertainty reduces firms' incentives to invest, hire, and expand. Uncertainty acts like a 'pause button' on business activity: when future conditions are unclear, firms delay irreversible decisions such as capital investment or entering new markets⁹. Said differently, higher uncertainty increases the value to firms of waiting¹⁰. This effect has been observed across sectors and in response to past technological transitions and policy and regulatory changes. It has also been documented in the Australian context, with past increases in economic uncertainty associated with materially slower growth in employment and business investment¹¹.

Policy-induced uncertainty, as distinct from general macroeconomic volatility, is also associated with lower investment and employment, especially in policy-sensitive sectors¹². At the macro level, periods of high policy uncertainty are correlated with weaker economic growth and foreshadow declines in investment, industrial production, and employment¹³. Policy uncertainty also depresses other forms of corporate reorganisation, with lower M&A activity, slower deal completion, and fewer IPOs¹⁴.

Finally, in relation to general economic uncertainty, there is some evidence of an important asymmetry: that increases in uncertainty can have bigger negative economic effects than the benefits from equivalent decreases in uncertainty¹⁵. To the extent this operates in the regulatory context, it could imply that a key priority is to avoid the perception and likelihood of large adverse future regulatory changes for business.

It is thus likely that the scale and complexity of complementary investments across skills, capital, organisational change, and coordination will make AI adoption by firms sensitive to the surrounding regulatory environment. Firms undertaking long-horizon investments in AI-enabled transformation must navigate legal uncertainty around data access, intellectual property, liability, and employment law.

Regulatory certainty can act as an economic coordination device: it enables firms, workers, and investors to act with confidence and align their efforts around AI adoption that delivers long-run gains. Seen in this way, just as complementary investments amplify the productivity potential of general-purpose technologies, they also raise the economic value of regulatory certainty. Therefore, unlocking AI's productivity dividend requires not just innovation in technology, but coherence and credibility in the institutional frameworks that govern its use.



03

Regulatory settings interact with the channels of productivity growth from AI

There is a clear productivity case for the Australian government to reduce regulatory uncertainty. Uncertainty around the future regulatory environment for AI means workers and firms don't know what their rights and responsibilities will be by the time their investments are 'up and running'. This could have a chilling effect on investment and reduce the potential upside of productivity growth from AI.

How much can the government expect to reduce uncertainty in practice? For a fast-evolving technology like AI, a degree of uncertainty is inevitable given the pace of change. It is not practical, or indeed credible, to offer a detailed regulatory blueprint for all future use cases. But some more detailed guidance should be possible, beyond the high-level response to the *Supporting Responsible AI* consultation. An absence of regulatory direction compounds the already high level of uncertainty facing Australian workers and firms as they adapt to AI. Regulatory uncertainty affects the channels for productivity growth in different ways, and the connection is more direct in some than others.

The first two channels discussed in Section 1 pertain to decisions taken within the firm: to augment and to automate human labour. It is here that the empirical and theoretical literature is strongest: that policy-induced uncertainty reduces the incentive to invest in technology and new business methods.

- **Labour-improving tools** depend primarily on individual or firm-level decisions about training and experimentation. Here, uncertainty around data use, IP protection, and acceptable AI-human collaboration can affect how confidently workers and firms adopt such tools, especially in highly regulated industries (e.g., legal, healthcare). Ambiguity about what is allowed can discourage experimentation.
- **Automation of tasks** requires significant capital investment. Firms are less likely to invest in AI systems if the regulatory treatment of liability, workforce requirements, or algorithmic accountability is unclear. For example, if future rules might limit the extent of AI-led decision-making or impose costly compliance, firms may delay automation investments.

The third and fourth channels discussed in Section 1 go to broader effects: new business models and the reallocation of resources across the economy. The effective working of these channels is highly affected by existing, broader product and labour market regulations. But AI regulation can affect these if it creates new, technology-specific rules that inhibit the creation of new businesses.

- **Reorganisation within firms** may be slowed by ambiguity around employment laws or other regulations. Reorganisation within firms is often time-consuming and disruptive at first. Regulatory clarity can help reduce the risk of needing to reverse or revise AI-enabled restructuring.
- **Broader reallocation across firms** is shaped by market structure and competition policy. If there's uncertainty about future regulatory barriers to scaling, data access, or competitive behaviour, firms may hesitate to expand. Conversely, clear guardrails can ensure that AI-enabled market leaders don't stifle diffusion or innovation, thereby enabling productive reallocation.

The fifth channel discussed in Section 1 deals with the growth of economy-wide knowledge. AI will likely improve the efficiency of idea-generation under multiple different regulatory approaches. Nonetheless, as with other forms of investment, a lack of regulatory certainty can inhibit some innovation.

- **Investment in new knowledge** through R&D is inherently risky, and ambiguity about IP protection for AI-generated content, data-sharing rules, or the legality of using AI in higher-risk contexts like health can significantly shape the expected return on innovation. Clear frameworks can lower the cost of innovation and reduce duplication of effort.



04

Not all regulatory approaches support productivity equally, and not all generate certainty

While regulatory certainty is essential for unlocking the productivity potential of AI, the form that certainty takes matters greatly. The durability and restrictiveness of different regulatory approaches carry distinct trade-offs for investment incentives, risk mitigation, and the diffusion of AI technologies across the economy. The international landscape illustrates three broad models currently under consideration.

Possible regulatory approaches to AI

A. Deliberately permissive approach to regulation



No new regulation; avoid applying existing regulatory goals to AI

B. Prescriptive, AI-specific approach to regulation



New regulation written specifically for AI development and use

C. Pragmatic, technology-neutral approach to regulation



Apply existing regulatory goals deliberately to AI

A. Deliberately permissive approach to regulation

This approach is reflected in recent US federal policy. It aims to prioritise innovation by deliberately minimising regulatory interference—sometimes even to the extent of creating specific exceptions to certain regulation for AI use. The US AI Action Plan, for example, directs agencies to cut back on existing rules that may slow AI adoption, while the Productivity Commission has floated the idea of amending the Copyright Act to include a fair dealing exception for text and data mining for AI systems. The purported goal is to avoid creating additional disincentives related to AI investment.

B. Prescriptive, AI-specific approach to regulation

A prescriptive, AI-specific approach, as adopted by the European Union in its 2024 AI Act, provides detailed rules tailored to different levels of risk. This includes outright bans on some uses, strict obligations in high-risk domains like healthcare, and regulatory sandboxes to support innovation. The purported goal is to provide certainty through clear obligations, aligned with the unique characteristics of AI.

C. Pragmatic, technology-neutral approach to regulation

A pragmatic, technology-neutral approach deliberately avoids AI exceptionalism and adapts existing regulatory goals such as privacy, fairness, and consumer protection to the use of AI. The goal is to focus on safeguarding outcomes rather than regulating inputs and departs from general law only when existing frameworks clearly fall short.

Australia's choice among which of these three approaches to follow matters not just because of their regulatory content, but because they influence whether and how firms and workers feel confident investing in the capabilities, systems, and transformations that productivity growth requires. For a dynamic, general-purpose technology like AI, regulatory certainty is not just a constraint on harm, but an economic enabler of productive change.

Table 1 (on page 17) takes the productivity pathways for AI and outlines how they are affected by each high-level regulatory approach in terms of regulatory certainty. In particular, we note that:

A. *Deliberately permissive approach—seductive in the short-term, uncertain and risky beyond*

A deliberately permissive approach could take one of two main forms. It could involve consciously adopting a laissez-faire approach, declining to create any new regulation or update existing law to address AI-related risks. Alternatively, it could take the more radical stance of exempting certain activities associated with the development and use of AI from complying with otherwise-applicable law.

Either form could minimise the short-term regulatory burden for firms and maximise the initial uptake of AI systems. If credible and consistent, this could promote productivity-enhancing reorganisation across the economy. But this approach would face significant challenges in maintaining regulatory certainty in practice.

One such challenge from this approach is state and territory governments stepping into what they perceive as a regulatory vacuum, creating a patchwork of sub-national regulation. A recent example of this is in AI-enabled workplace surveillance, where several states, including NSW and Victoria, are in the process of reforming their own laws.

Another risk is that this approach fails to guard against a range of consequential emerging harms, resulting in long-term instability in government policy. For example, without clear liability regimes, developers and users are unlikely to mitigate second-order societal impacts and even some direct harms flowing from AI, increasing both the risk of harm and the likelihood of costly regulatory reversals.

Over time, policy reversals or reactive interventions to emerging, AI-related risks and harms could undermine both business and consumer confidence and weaken AI adoption and investment. It could thus fail to build the institutional foundations to foster significant long-term productivity gains.

B. Prescriptive, AI-specific approach—selective certainty, rigid rules

A prescriptive, AI-specific regime aims to provide clarity for AI developers, deployers and users, particularly in high-risk contexts. By creating risk-based rules and setting compliance standards specific to AI systems, it can deliver short-term certainty about obligations, supporting adoption where both use cases and related risks are clear. Because productivity gains depend on predictable rules that let firms invest and deploy at scale, this promise of certainty is the main channel to higher productivity.

However, this approach risks being both inflexible and increasingly burdensome in fast-changing contexts. Detailed rules can struggle to remain durable as technology evolves. As a result, the effort of interpreting rules and implementing compliance, such as human oversight, may overwhelm both firms and regulators, potentially creating a disproportionate burden on smaller firms. When rules date quickly or diverge, planning becomes less predictable, undermining the very certainty needed for investment and diffusion. Thus, while such a framework may enhance trust in narrow settings, it is unlikely to sustain long-term regulatory certainty.

A further concern is that this approach would create a new class of technology-specific regulation, applying across existing regulatory frameworks in product and labour markets. Being overly prescriptive about how technologies should be designed and implemented, regardless of the actual risk, could become a drag on innovation and could limit the productivity growth from new business models and reallocation of resources across the economy.

C. Pragmatic, technology-neutral reform—flexible, but reliant on strong signals across key policy areas and regulators

A pragmatic, technology-neutral approach relies primarily on existing rules, while encouraging technological innovation and remaining adaptable to emerging risks. Its flexibility and focus on outcomes rather than inputs means this model may be particularly suited to Australia's likely downstream role in the AI value chain, enabling adoption without locking in specific technologies. This approach still accepts the need for reform, some of which may be technology specific, but it prioritises the more effective application of existing law.

A key strength of this approach is that it embeds a degree of discipline around the identification of harms. It encourages policymakers and regulators to rely first on current regimes that seek to address risks that arise regardless of the technology in question and depart from these only where there is a case for doing so, such as where those laws cannot be effectively applied or where a genuinely new issue arises that merits a novel regulatory response.

Compared with a permissive approach that could prove unstable in practice, or a prescriptive approach that risks rigidity, a pragmatic, technology-neutral model can offer greater regulatory certainty. It avoids speculative rulemaking aimed at anticipating every possible high-risk use case, and it does not signal a doctrine of non-interference. Instead, by orienting firms toward compliance with outcome-focused laws already in force, it establishes a stable baseline of expectations: businesses know which obligations apply, regulators can act consistently, and both parties can adapt incrementally as new risks emerge.

However, its success depends on consistent interpretation, proactive guidance, and adequate regulatory resourcing. Without these, firms would still face ambiguity in areas like privacy, intellectual property law, algorithmic accountability, and the application of employment and consumer law in novel contexts.

The differences evident among these three approaches demonstrate that regulatory certainty for AI-driven productivity investment is not just about choosing a particular framework. Beyond clarity, the framework must be structured and enforced in ways that signal reliability, predictability, and long-term institutional commitment.

Productivity pathway	A. Deliberately permissive approach to regulation	B. Prescriptive, AI-specific approach to regulation	C. Pragmatic, technology-neutral approach to regulation
Labour-improving tools and automation: AI augments workers by making existing tasks more efficient (e.g. coding assistants) or substitutes for labour in routine and some complex tasks (e.g. customer service chatbots).	May support the adoption and use of existing AI tools and productivity benefits in the short term. Concerns about the stability of this regulatory approach (i.e., the risk of ad hoc regulatory reaction to potential harms) could suppress investment. Varying AI-specific regulation across states may compound regulatory frictions.	Provides clarity initially (particularly for a subset of defined AI use cases) but is likely to be inflexible in allowing for changing uses. Could deter beneficial experimentation or investment in responsible, low-risk uses of AI due to 'over-compliance' or confusion.	Offers a degree of clarity via prioritising the enforcement of existing, technology-neutral laws. Where appropriate, changes (such as displacement from automation) are dealt with through existing regulatory schemes, allowing for gains from implementation to be realised.
Reorganisation within and across firms: AI enables firms to restructure workflows, decision-making, and internal operations, and more productive firms grow and new entrants emerge, shifting resources to where AI is most effective.	Unimpeded uptake of AI could lead to greater dynamism through greater reorganisation within firms and the entry of new firms and business models. Dynamism—and the emergence of AI-related risks—will be driven by how existing regulatory frameworks affect firm entry and exit, labour mobility and competition.	May create rigid compliance requirements that inhibit reorganisation within firms. Regulatory requirements may disproportionately burden new entrants relative to incumbents. Dynamism will be primarily driven by existing regulatory frameworks on firm entry and exit, labour mobility and competition, but there is a risk of new limits to firm dynamism based on tech-specific rules.	Encourages productivity-improving reorganisation aligned with existing employment and competition regulatory goals. There will likely be additional regulatory burden to accommodate new harms from AI use and this may be more burdensome for smaller firms. The tech-neutral approach helps clarify whether new business models should be regulated differently from existing regulation of harms.
Knowledge & innovation: AI supports invention, idea generation, and long-run growth through R&D and innovations which shift the productivity frontier.	Weak IP rules and unpredictable liability could discourage long-term research and development. To the extent that policy uncertainty matters as much or more than policy settings themselves, any doubts over the credibility and stability of this approach could limit innovation.	Encourages R&D in narrow, regulated domains where rules are clear but may stifle exploratory or cross-cutting innovation due to complexity.	Provides a degree of confidence for R&D through stable rights and obligations, but is heavily reliant on suitable law reform to clarify and update the status of key AI inputs and outputs.

05

The path forward: reducing regulatory ambiguity to unlock AI-driven productivity

Australia does not yet have AI-specific laws. A number of existing technology-neutral laws of general application do apply to the development and use of AI. However, incomplete application of these existing laws has wrongly led some to conclude that Australia operates in a 'digital wild west' where the law doesn't apply to AI¹⁶.

The path ahead for AI regulation in Australia remains highly uncertain, a fact that is reflected in surveys and data regarding concern by CEOs about regulation for AI¹⁷. Australian policymakers are exploring AI-specific regulation amid an evolving range of international approaches for how best to regulate AI systems.

While a number of consultation processes have been run by the government, these do not yet amount to a clear signal from the Australian Government on how it will seek to regulate AI. By contrast, other jurisdictions such as the European Union and the United States have laid out frameworks or legislation that more clearly signal their regulatory intent, which vary in line with the three approaches described above. These choices reflect different views about the risks and opportunities posed by AI, and about the role regulation should play in shaping them.

Because of this variance, international experience does not provide a blueprint for Australian workers and firms. Investors, boards, executives and workers are left without a clear sense of whether the domestic regime will follow a precautionary, innovation-focused, or sector-specific model.

Australia cannot eliminate all ambiguity, especially with regard to a general-purpose technology as dynamic as AI. But it can and should provide the kind of regulatory certainty that matters most: clear direction, credible commitments, and institutions capable of interpreting and applying evolving rules.

Regulating appropriately for AI also strengthens investment certainty. Where upstream jurisdictions have clearly set out requirements for AI use, and Australia's regulatory settings are aligned, compliance costs can be reduced and firms can scale more easily across markets. Well-tailored regulation provides guardrails for responsible AI adoption, while alignment with technology-neutral approaches from major jurisdictions ensures that Australian businesses are not burdened with duplicative or conflicting rules.

This form of certainty does not require predicting every future use case in advance. Rather, it reduces the economic cost of waiting by enabling firms and workers to invest in the skills, systems, and organisational change that AI adoption demands.

The choice between a passive model, a prescriptive regime, or a technology-neutral approach is a critical determinant of whether Australia creates the conditions for AI-driven productivity or stalls its potential. Each model signals something different about the role of the state, the allocation of risk, and the expected pace of diffusion. But even the most appropriate regulatory framework will fall short without consistent interpretation, timely guidance, and strong coordination across governments and sectors. These are not operational details. They are the institutional infrastructure that enables long-horizon, productivity-enhancing investment.

Regulatory certainty, in this context, is not about fixing the rules. It is a property of the system as a whole. It arises from the interaction of clear goals, transparent processes, and a regulatory system with the capacity to respond to a rapidly evolving technological context. When firms and workers can understand and anticipate how and where the rules apply, and trust that those rules will evolve reasonably over time, they are more likely to invest in the complementary capabilities that unlock productivity, including labour-augmenting tools, automation, organisational reorganisation, reallocation, and innovation.

Australia's opportunity may not be to lead the world in building foundation models, but rather to lead in building the institutional capability to adopt AI in ways that boost productivity, economic dynamism, and long-term wellbeing. If AI is to deliver on its promise, it will do so by enabling workers to do more with less, by supporting reorganisation within firms, by allowing more efficient firms to grow, and by accelerating the generation and diffusion of new ideas. These outcomes all rely on the same thing: the willingness and ability of firms, workers, and institutions to invest in change. That willingness is shaped by the regulatory environment. Governance is not a constraint on productivity. It is the enabler of investment in the complementary systems and capabilities that make productivity growth possible.

The question is not whether Australia should regulate—our laws always apply to new technologies. The challenge here is to ensure that our laws strike an appropriate balance between encouraging AI innovation in a way that benefits the Australian community and economy as a whole beyond the very short term.

Reducing regulatory ambiguity is not simply a matter of legal clarity. It is a strategic lever for increasing the pace and quality of AI adoption across the economy. With the right regulatory settings and institutional commitments, Australia can unlock AI's full productivity potential, thereby delivering higher incomes, more efficient services, and broader access to the tools and opportunities of technological progress.

What is productivity?


There are two common ways to measure productivity, depending on which inputs are considered:

- **Labour productivity** measures output per hour worked. It captures improvements that come from better skills, new equipment, or smarter ways of doing things. For example, when farmers use new tractors to harvest more wheat per hour worked, labour productivity increases.
- **Multi-factor productivity** (also called total factor productivity) measures how efficiently labour and capital inputs—such as buildings, machinery, and information technology—are combined. It reflects productivity growth from innovation and better organisation rather than from simply increasing inputs. For example, if a factory reorganises its workflow to produce more goods without needing extra workers or new machines, multi-factor productivity has improved.

Historically, labour productivity growth explains most of Australia's increases in real incomes. But productivity and income are only means to an end, not ends themselves. In principle, wellbeing could fail to rise if the extra income is captured by a narrow slice of society, if workers displaced by technology are left behind, or if progress comes at the cost of environmental degradation. Some analysis suggests that benefits from productivity are broadly shared: that more than 90 per cent of society's benefits are received by consumers through higher quality and cheaper products, rather than by producers¹⁸. Others have argued that technological progress has not always translated into broad welfare gains, and the result has depended on institutions that support worker bargaining power¹⁹.

Productivity & Income Growth

Growth since 1995



* Average hourly earnings (national accounts) deflated by the CPI

** Measured in the market sector only

Source: ABS National Accounts

Why has productivity growth slowed?

Productivity growth has been slowing since the mid-2000s, in Australia and in other higher-income countries. There is debate about whether AI can turn this trend around. One way to understand the productivity slowdown is growth accounting, which identifies three key contributors.


First, the earlier surge in investment in IT hardware and software that powered growth in the 1990s and early 2000s has faded. Second, growth in other types of capital per worker — like machinery — has also slowed and even declined around the COVID period.

In combination, these two factors reflect that business investment has not kept up with growth in the workforce.

There has also been an important historical role for innovations in practices—i.e., ways of doing things rather than technology you can drop on your foot. Important advances in practices include double-entry bookkeeping, limited liability companies, financial markets, secure property rights and regulation of anti-competitive practices. These practices promote economic dynamism by increasing the means and the rewards to innovation and investment.

Sources of Labour Productivity Growth

Contributions to annual average growth by period*



* Market sector; relies on an experimental ABS methodology. Source: ABS Estimates of Industry Multifactor Productivity

Third, multi-factor productivity (MFP)—which captures improvements in how efficiently labour and capital are used together—has grown more slowly since the mid-2000s. In other words, Australia has not been generating much more output from the same inputs. This is true even after accounting for gradual gains in workers' education and experience, which are treated separately to MFP in this version of growth accounting²⁰.

But growth accounting doesn't tell the full story of where productivity growth really comes from. Why do businesses invest, or not? Where do new ideas and technologies come from? We have some answers to these questions, but there is also a degree of mystery.

We do know that substantial productivity growth is largely a phenomenon of the past two centuries. The industrial revolution was a turning point, and the invention of new technologies has been a key driver since. Many new technologies solve economic problems or create new economic value. The most important have been 'general-purpose technologies' (GPTs) with broad applicability across the economy, like electricity, computers and the internet.

Part of the slowdown in productivity growth since the mid-2000s may reflect a period with less transformative technological progress²¹. There is also evidence of a broader decline in economic dynamism and competition since the mid-2000s²². This suggests that part of the slowdown in productivity growth is because resources are being reallocated to the most productive firms at a slower rate. There is also evidence that Australian firms are falling further behind the global productivity frontier²³.

Other explanations for the productivity slowdown include a slowdown in global trade and the rise of the non-market sector. Trade surged in the 1990s following a dismantling of trade barriers in Australia and elsewhere, boosting productivity through specialisation of production and competition. But that had substantially run its course by the mid-2000s²⁴. More recently, growth in government-funded services (i.e., the non-market sector) has weighed on measured productivity due to a combination of measurement issues and genuine challenges to achieving productivity gains in labour-intensive care and education services²⁵.

Sources of potential harms from AI and the role of regulation

There is broad recognition of AI's potential to raise productivity. However, like past general-purpose technologies, AI also poses significant risks and challenges²⁶.

As outlined by the Human Technology Institute, potential harms from AI can be grouped into three categories:

- Malicious or misleading deployment of AI. AI makes a range of existing damaging activities cheaper and potentially more effective, in a perverse parallel of AI's potential positive effects on productivity. This includes concerns about the use of AI for cyberattacks, misinformation, or the development of bioweapons. New methods of harm may be developed, akin to the malicious end of new business models. This category can also include less overtly malicious activities, such as the use of AI to manipulate consumer behaviour. However, it is important to delineate whether applications such as targeted advertising cause harm.
- Overuse, inappropriate or reckless use of AI without regard to its second-order effects. Individuals may deploy AI without malicious intent but fail to account for the second-round effects of these actions at scale. Two examples are threats to privacy from the use of AI facial recognition technologies and threats to creative and innovation professions if existing regulation fails to protect intellectual property.
- Failures of AI to operate in the way, or to the level of quality, required. An example of this is when AI model errors are systematic in ways that create biased treatment by race or sex.

Some potential harms may be resolved through market forces. For example, in most markets, it is acceptable to have a degree of quality segmentation, and firms adopting AI have an incentive to adopt AI that operates in the way or to the level of quality that delivers for them. This could work to manage issues of AI quality in parts of the market sector where consumers can judge their desired quality and firms can respond to price signals.

However, in many contexts, thoughtful and effective regulation will be critical to addressing the potential for harm while also empowering the social benefits from AI. A central objective of regulation should be the maximisation of overall public benefit—including productivity gains, distributional outcomes, and the minimisation of harms. Importantly, these objectives are not always in conflict: well-designed regulation can mitigate risks while supporting innovation, competition, and inclusive growth.

Endnotes

- 1 Timothy F. Bresnahan and M. Trajtenberg, "General Purpose Technologies 'Engines of Growth'?", *Journal of Econometrics* 65, no. 1 (1995): 83–108.
- 2 Michael Webb, *The Impact of Artificial Intelligence on the Labor Market*, 2020.
- 3 Paul A. David and Gavin Wright, "General Purpose Technologies and Surges in Productivity: Historical Reflections on the Future of the ICT Revolution," in *The Economic Future in Historical Perspective*, ed. Paul A. David and Mark Thomas (British Academy, 2006), <https://doi.org/10.5871/bacad/9780197263471.003.0005>.
- 4 Samuel G Goldberg and H Tai Lam, *Generative AI in Equilibrium: Evidence from a Creative Goods Marketplace*, 2024.
- 5 Philippe Aghion et al., "Artificial Intelligence and Economic Growth," *NBER Working Paper*, 2017.
- 6 Paul M. Romer, "Endogenous Technological Change," *Journal of Political Economy* 98, no. 5 (1990): S71–102.
- 7 Nicholas Bloom et al., "Are Ideas Getting Harder to Find?," *American Economic Review* 110, no. 4 (2020): 1104–44.
- 8 Paul A. David, "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox," *The American Economic Review* 80, no. 2 (1990): 355–61.
- 9 Nicholas Bloom, "The Impact of Uncertainty Shocks," *Econometrica* 77, no. 3 (2009): 623–85.
- 10 Nick Bloom et al., "Uncertainty and Investment Dynamics," *The Review of Economic Studies* 74, no. 2 (2007): 391–415.
- 11 Angus Moore, "Measuring Economic Uncertainty and Its Effects," *Economic Record* 93, no. 303 (2017): 550–75.
- 12 Huseyin Gulen and Mihai Ion, "Policy Uncertainty and Corporate Investment," *The Review of Financial Studies* 29, no. 3 (2016): 523–64.
- 13 Nicholas Bloom, "The Impact of Uncertainty Shocks"; K Jurado et al., "Measuring Uncertainty," *American Economic Review* 105, no. 3 (2015): 1177–216.
- 14 Alice Bonaime et al., "Does Policy Uncertainty Affect Mergers and Acquisitions?," *Journal of Financial Economics* 129, no. 3 (2018): 531–58; Nam H. Nguyen and Hieu V. Phan, "Policy Uncertainty and Mergers and Acquisitions," *The Journal of Financial and Quantitative Analysis* 52, no. 2 (2017): 613–44; Gonul Colak et al., "Political Uncertainty and IPO Activity: Evidence from U.S. Gubernatorial Elections," *The Journal of Financial and Quantitative Analysis* 52, no. 6 (2017): 2523–64.
- 15 Andrew T. Foerster, "The Asymmetric Effects of Uncertainty," *Economic Review*, no. Q III (2014): 5–26; Giovanni Caggiano et al., "Uncertainty Shocks and Unemployment Dynamics in U.S. Recessions," *Journal of Monetary Economics* 67 (October 2014): 78–92.
- 16 Lauren Solomon and Nicholas Davis, *The State of AI Governance in Australia* (Human Technology Institute, 2023).
- 17 KPMG, "Generative AI Is Top Investment Priority despite Economic Uncertainty," <https://kpmg.com/au/en/media/media-releases/2024/09/generative-ai-top-investment-priority-despite-economic-uncertainty.html>.
- 18 William D. Nordhaus, "Schumpeterian Profits in the American Economy: Theory and Measurement," *NBER Working Paper*, April 2004.
- 19 D Acemoglu and S Johnson, *Power and Progress: Our Thousand-Year Struggle over Technology and Prosperity* (PublicAffairs, 2023).
- 20 Angelina Bruno et al., "A (Closer to) Real Time Labour Quality Index," *RBA Bulletin*, Reserve Bank of Australia, 2025.
- 21 Z Duretto et al., "Understanding Productivity in Australia and the Global Slowdown," *Treasury Roundup*, 2022.
- 22 Dan Andrews et al., "The State of Competition in Australia," E61 Research Note, 2023; J Hambur, "Product Market Power and Its Implications for the Australian Economy," *Treasury Working Paper*, Department of the Treasury, 2021.
- 23 Dan Andrews et al., *Reaching for the Stars: Australian Firms and the Global Productivity Frontier*, 2020.
- 24 Ian Goldin et al., "Why Is Productivity Slowing Down?," *Journal of Economic Literature* 62, no. 1 (2024): 196–268; Jarkko Jaaskela and Thomas Mathews, "Explaining the Slowdown in Global Trade," *RBA Bulletin*, 2015.
- 25 M Maltman and E Rankin, "What If We Didn't Care? Implications of the Growth in the Care Economy for the Broader Macroeconomy," *E61 Research Note*, e61 Institute, 2024.
- 26 Solomon and Davis, *The State of AI Governance in Australia*, 2023.

