Ministry for Regulation
Te Manatū Waeture

Artificial intelligence: A literature review of anticipated economic impacts and regulatory responses

Sam Le Quesne
Kevin Counsell
Isabelle Sin

August 2026

Published in August 2026 by the Ministry for Regulation, Wellington, New Zealand.

ISBN 978-1-991372-17-8 (Online)

Executive summary

This literature summary is a snapshot (as at early 2026) of the emerging economic literature on artificial intelligence (AI) and its implications for productivity, economic growth, labour markets, inequality, uncertainty, and market structure. It does not make any independent claims about government policy or regulation of AI in the NZ context. Rather, it synthesises recent theoretical and empirical work to assess how AI is likely to affect economic outcomes and where existing regulatory frameworks may be ill-suited to manage these changes. While most studies agree that AI will increase productivity, there is uncertainty about the magnitude, timing, and distribution of gains, with policy choices playing a central role in shaping outcomes.

Key takeaways for regulation and policy include:

- Growth impacts are uncertain: Estimates range from small near-term productivity gains to highly transformative long-run scenarios, depending on whether AI mainly automates existing tasks or accelerates innovation and research.
- Distribution matters as much as growth: Many papers emphasise that GDP growth and welfare may diverge if productivity gains are unevenly distributed between labour and capital or if growth is associated with social harms.
- Labour markets are the main transmission channel for AI's economic impacts: Adjustment is likely to be slow and uneven, creating transitional costs and equality issues even in scenarios with positive economic growth.
- Markets may misdirect AI development: Private incentives tend to favour labour-replacing or socially harmful applications, leaving a role for policy to steer innovation toward labour-complementing, productivity-enhancing, and socially beneficial uses.
- Uncertainty complicates regulation: The novelty and potential irreversibility of AI harms make static, ex-ante regulation ineffective. Regulators should utilise adaptive, risk-based approaches such as sandboxes and staged deployment.
- Market power risks are significant: AI markets exhibit strong tendencies toward concentration, implying that competition policy and oversight will be important.

Overall, the literature suggests that effective AI policy should focus less on slowing technological progress and more on shaping its direction, managing distributional impacts, addressing uncertainty, and monitoring market structure to ensure productivity gains translate into welfare improvements.

Introduction

This literature review aims to provide information and background of how artificial intelligence (AI) challenges existing policy and regulation. It focuses on impacts on economic growth predictions, labour markets, innovation, uncertainty, and market power. These topics appear throughout the literature as major areas of concern for regulators and policymakers. Most papers argue current regulation is ineffective or will be ineffective as AI continues to evolve. Most papers also tend to agree that the productivity gains from AI are positive, and that regulation is more about shaping the direction and distribution of the productivity gains.

In this literature review, AI tends to refer to large language models (LLMs). These are AI systems trained on massive amounts of data to understand and generate human-like language (Stanford HAI, 2026). Foundation models are LLMs that can be adapted for a wide range of downstream applications (Vipra and Korinek, 2023). Examples of well-known providers are OpenAI and Anthropic.

A few papers look at the effects of automation. This can refer to the use of AI to automate tasks, but also the use of robotics and machinery to conduct tasks (Acemoglu and Restrepo, 2018). Future-looking papers refer to artificial general intelligence (AGI). This is AI that can match human levels of intelligence across all tasks that humans can do (Korinek, 2024). While this level of AI has not been achieved, parts of the literature see it as a future possibility and try to predict its economic impacts.

Predictions for economic growth

The literature shows some disagreement over the impact of AI on economic growth. Estimates range from relatively modest near-term productivity gains to highly transformative long-run scenarios. These differences largely reflect contrasting assumptions about the mechanisms through which AI affects the economy, particularly whether AI primarily automates existing tasks or fundamentally accelerates innovation and knowledge creation.

More cautious projections suggest that AI will raise productivity and GDP, but by relatively small amounts over the next decade, broadly comparable to earlier waves of digital technology. Acemoglu (2025) aggregates task-level labour market effects to assess their macroeconomic implications and argues that, under plausible adoption and investment scenarios, AI is likely to deliver positive but modest productivity gains. His analysis emphasises slow diffusion, organisational adjustment, and the limited scope of current AI applications, leading to a relatively pessimistic outlook compared with other parts of the literature. Nordhaus (2021) similarly downplays the likelihood of an AI-driven “singularity”, the idea of a superintelligent AI that will cause exponential growth in productivity. He finds that recent data do not support the conditions required for such outcomes, citing diminishing returns on investment and bottlenecks in production, such as energy and resource constraints. As a result, he treats AI as a continuation of past general-purpose technologies rather than a “deus ex machina”, with GDP per capita growth expected to remain close to historical trends.

In contrast, more optimistic projections argue that AI could have much larger growth effects if it accelerates innovation itself rather than merely automating existing tasks. In these models, AI acts as a general-purpose technology that boosts idea generation, research productivity, and the creation of new tasks, potentially raising long-run growth rates. Jones (2026), for example, predicts that AI’s impact on growth could be an order of magnitude higher than that of the internet, driven by faster knowledge creation. Despite this more optimistic growth outlook, Jones emphasises concerns around inequality and labour market adjustment, suggesting that policy challenges remain central regardless of the scale of productivity gains.

From a more extreme perspective, Trammell and Korinek (2025) lay out what could happen in a highly transformative scenario in which artificial intelligence fully automates both production and research and development. In this setting, they predict AI will break the long-standing Kaldor facts of stable growth rates and labour shares, potentially leading to explosive growth while rendering human labour largely redundant. The authors are explicit that this is a thought experiment rather than a forecast, and they note that such growth could only be unconstrained until physical constraints or diminishing marginal returns bind it. These scenarios rely on highly speculative assumptions and are best interpreted as upper-bound possibilities rather than empirically grounded predictions.

Empirical evidence to date aligns more closely with the cautious view. Aggregate productivity statistics show no clear AI-driven surge despite rapid technical progress, which is consistent with historical patterns in which major technologies take time, complementary investment, and organisational change before translating into GDP growth. Forward-looking projections also remain constrained. The IMF World Economic Outlook Update (January 2026) suggests that AI is likely to have only a minor impact on the global economy over the medium term, while Filippucci, Gal and Schief (2024) similarly project low aggregate growth effects over the next decade for OECD economies.

Despite disagreement over magnitude, there is a broad agreement across the literature on several key risks. Growth gains from AI are likely to be unevenly distributed, with important implications for labour and capital incomes, discussed further in the labour market literature. Overall, the literature suggests that AI is likely to increase economic growth, but with substantial uncertainty around timing and scale. Regardless of aggregate outcomes, most studies emphasise that AI’s primary economic impacts will be transmitted through labour markets, inequality, and market structure, making distributional and adjustment policies central to the policy response.

Labour markets and inequality

Much of the literature identifies labour markets as a main channel through which AI-driven automation affects inequality. Acemoglu and Restrepo (2018) show that automation displaces workers by replacing specific tasks, with early impacts concentrated among routine and lower-skilled roles. Increased automation decreases demand for those workers and subsequently decreases their wage as well. Acemoglu and Restrepo predict from this a growing inequality between workers whose tasks are automated and those whose tasks remain complementary with technology.

This is built upon by their later work where they assume that automation can displace a much broader range of jobs (Acemoglu and Restrepo, 2020). Inequality in this case shifts to being between workers and capital owners, concentrating wealth at a higher level than before. The authors argue in both their papers that market adjustment will be slow and incomplete, meaning policy and regulation will be required to support displaced workers. They suggest focus should be on retraining, income support, and the creation of accessible jobs for humans.

Korinek and Stiglitz (2018) provide a useful framing for why AI-driven automation is politically and economically contentious. They argue AI is likely to be worker-replacing, so the central issue is not whether productivity rises, but who receives the gains. They identify two main inequality channels: surplus captured by innovators (rents/market power) and redistribution via factor price changes (wages falling relative to returns to capital and other scarce factors). They argue AI can still be welfare-improving if institutions and redistribution mechanisms are designed so winners compensate losers, including via relatively low-distortion taxes on windfall gains.

Recent empirical work strengthens the task-based theory by providing early causal evidence of AI's labour market effects. Johnston and Makridis (2025) use a difference-in-differences approach to show that occupations with higher AI exposure experience measurable gains in employment and wages. Complementary work by Massenkoff and McCrory (2026) and Brynjolfsson, Chandar and Chen (2025) similarly finds that AI impacts appear first in job content, hiring, and task composition rather than headline employment. Eckhardt and Goldschlag (2025) argue that this pattern reflects gradual adjustment along intensive margins, suggesting that current data may understate longer-run distributional consequences.

Korinek (2024) extends this to AGI in his working paper, arguing that when AGI systems can perform most cognitive tasks, demand for human labour will fall sharply across the economy. He predicts falling wages, rising unemployment, and failure of traditional labour market tools such as the minimum wage or employment-based welfare systems. Inequality is predicted to not just increase within a country between workers and capitalists, but also internationally as capital-abundant countries disproportionately benefit. Korinek therefore argues that policymakers and regulators will have to push for redistributive mechanisms that share the AI-generated productivity gains amongst workers displaced by AI, such as a universal basic income. Both Korinek (2024) and Acemoglu and Restrepo (2018) support the idea that market adjustment will be slow and incomplete.

Acemoglu, Kong and Ozdaglar (2026) offer a separate mechanism that may increase inequality, arguing that AI will substitute the human-learning effort, leading to a decline in human capital over time as humanity experiences an erosion of general knowledge. They believe as AI is used to provide accurate and context-specific recommendations, it reduces incentives for individuals to invest in learning that also generates shared general knowledge. Because this shared knowledge acts as a complement to labour, the quality of human capital may decline over time as workers become more dependent on AI to carry out their tasks. The working paper therefore suggests that AI-driven inequality will not only be introduced through the improvement of AI technology, but also through the regression of human skills and knowledge. This suggests a role for policy that preserves learning incentives for humans and supports skill development alongside AI adoption.

Building on the growth projections discussed earlier, Acemoglu (2025) believes even when AI raises aggregate productivity, GDP growth and welfare need not move together. He suggests there will be a widening gap between labour and capital income, even in the absence of large aggregate employment losses, as capital productivity increases at a faster rate than labour. His central argument is that AI applications that raise output through socially unproductive or harmful uses may increase measured productivity whilst reducing welfare. This argument reinforces the importance of policies that shape the direction of AI deployment rather than focusing solely on the pace of adoption.

Innovation incentives and direction

A consistent theme in the literature is that markets shape not only the pace of AI development, but also its direction. Private incentives tend to support forms of AI that are socially harmful or labour-replacing. Acemoglu (2021) shares his opinion that, without regulation, AI research moves toward automation, surveillance, and behavioural manipulation because these AI uses are privately profitable despite imposing broad social costs on workers, privacy, and the wider democratic process. He believes these issues come down to distorted incentives for AI producers rather than the inevitable direction for AI, leaving space for regulation and policy to help push AI in the socially preferred direction.

Focusing specifically on labour markets, Acemoglu, Autor and Johnson (2026) argue that the market undersupplies pro-worker AI: technologies that complement human skills and expand worker productivity and knowledge rather than replacing workers outright. Because firms do not capture the full social benefits of labour complementing innovation, only the private costs, they are incentivised to favour automation instead.

The authors therefore argue for policies that actively redirect innovation, such as competitive grants for pro-worker AI, stronger worker input into AI adoption, and laws governing the use of worker expertise. They suggest that workers should retain greater control over how their knowledge and skills are used to train AI, with compensation if appropriate. These measures aim to correct the private market distortions by making labour-complementing innovation more attractive relative to purely labour-replacing applications.

Agrawal, McHale and Oettl (2026) claim that AI will increase both human knowledge change and will become a substitute for humans, impacting two parts of the classic production function. They present a multi-stage model of scientific productivity. The paper emphasises that AI tools can support multiple stages while retaining an important role for human judgement. Their advice for policymakers is to focus on balancing different types of economic responses to AI-driven change: if AI automates most tasks, defensive policies should be pursued such as a universal basic income; and if AI functions as a tool for augmentation, a more offensive strategy focusing on increasing human capital investment will be needed to ensure productivity gains.

Regulating AI under uncertainty

A central theme in the literature is that uncertainty fundamentally complicates AI regulation. Due to its novelty, AI may generate harms that are difficult for policymakers and regulators to predict. This means that standard regulatory tools risk either allowing socially harmful over-adoption or imposing over the top restrictions.

Acemoglu and Lensman (2024) argue that markets adopt AI too quickly because firms face private costs that are lower than the social costs of the AI, which can be especially harmful when risks are uncertain or irreversible. They show that optimal policy should favour gradual and differentiated adoption rather than laissez-faire adoption or total bans. To achieve this, they propose sector-specific, risk-based measures, including taxes on AI use that reflect expected social harm and decline as uncertainty is resolved, alongside regulatory sandboxes that permit experimentation in low-risk contexts whilst delaying deployment in high-risk areas.

Guerreiro, Rebelo and Teles (2023) also argue for regulatory sandboxes, emphasising the information asymmetries that exist between regulators and AI providers. They argue that regulators often lack access to developers' true beliefs about AI risks, undermining the effectiveness of taxes or liability rules. To address this, they propose a two-stage regulatory framework in which AI systems are either immediately released or, if deemed risky, subjected to a mandatory, restricted sandbox experiment. Evidence from the testing phase is then used to determine whether broader deployment is permitted and what, if any, regulations should be applied.

Kretschmer et al. (2023) critique risk-based AI regulation under uncertainty, arguing that static risk classifications can quickly become outdated in the context of rapidly evolving AI systems. They warn that ex-ante regulations may shift rather than reduce harm if liability for downstream impacts is weak or misallocated. They instead push the importance of aligning responsibility with control and suggest the actors who adapt and deploy the AI systems should bear primary liability if harm occurs. This focus on liability provides a more flexible regulatory tool that can respond to unforeseen risks.

Agrawal, McHale and Oettl (2026) argue that AI will reshape the production of knowledge by both accelerating scientific discovery and substituting for human effort in parts of the research process. They present a multi-stage model of scientific productivity, decomposing research into question generation, idea generation, design, and testing, and show that AI can affect each stage differently. The paper emphasises that while AI tools can support multiple stages of scientific work, human judgment remains central, particularly in inference, interpretation, and validation.

Market power

An underlying theme in the literature is that AI market is structurally prone to concentration and market power. Vipra and Korinek (2023) look at leading AI models as natural monopolies, due to high fixed costs, low marginal costs, and reliance on scarce inputs such as computational resources, data, and specialised talent. These features imply that unconstrained AI markets are likely to produce standard monopoly outcomes, including restricted supply, inefficient pricing, and concentrated economic and political power.

They further warn AI firms may look to extend their market power into downstream markets through vertical integration, reinforcing barriers to entry and limiting competition. They propose a two-pronged approach to regulation. Regulators should ensure market contestability to limit market power issues, and quality and safety standards should be imposed to improve social welfare. The former aims to limit power concentration, and the latter recognises that competition alone may not be enough in the naturally monopolistic market to ensure quality and safe AI products.

They outline several ways they believe regulators can promote market contestability in practice. First, antitrust tools, like merger review, can be used to limit vertical integration by investigating acquisitions and partnerships that raise input costs or advantage firms. Regulators can also discourage anti-competitive labour practices, such as restrictions on worker mobility, to prevent incumbents from hoarding scarce talent. In addition, data governance and data protection rules can constrain the accumulation and reuse of proprietary data, for example through purpose limitation and restrictions on data sharing, reducing entry barriers for rivals. Alongside these measures to preserve contestability, the authors argue that dominant firms should be subject to direct regulation that imposes quality, price, access, and non-discrimination standards, to stop monopolistic behaviour.

Korinek and Vipra (2024) similarly argue that the market for foundational models exhibits structural features that generate strong tendencies towards concentration. They believe that rapidly rising fixed costs such as computational resources, data, and talent create significant economies of scale and scope. Whilst competition now remains intense, these forces raise the risk of market tipping and dominance over time, such as through vertical integration with cloud and computer providers.

Korinek (2024) reaches similar conclusions in the context of AGI, arguing that frontier AI systems are likely to generate winner-takes-all outcomes that existing antitrust frameworks are poorly equipped to address. He suggests that traditional, ex-post competition policy may be inadequate, and that instead more tailored interventions such as regulating access to key inputs like data may be necessary to prevent concentration and its distributional consequences.

Schmid, Sytsma and Shenk (2024) assess whether the market concentration would constitute a natural monopoly rather than a traditional antitrust problem. They develop criteria for identifying natural monopoly conditions in the foundation-model market, emphasising economies of scale, sunk costs, and scope across applications. They argue that if scaling laws persist and compute costs remain high, the production of pre-trained foundation models may become cost sub additive, strengthening the case that competition policy alone may be insufficient and that direct regulatory oversight could be required.

Policy and regulation recommendations

Altogether, the literature suggests that effective AI policy should focus less on slowing technological progress and more on shaping its direction and distributional consequences, limiting uncertainty, and monitoring market structure.

For labour market policy and regulation, the literature suggests that AI-driven task automation is likely to generate uneven adjustment, with early impacts concentrated in specific tasks and occupations. The policy objective is therefore to shape incentives around AI adoption toward augmenting workers rather than replacing them, and to support workers through slow and incomplete adjustment while preserving incentives for skill formation and labour market participation.

Some suggestions from the literature are income support and adjustment assistance for displaced or at-risk workers, active labour market policies focused on retraining and job transitions, and redistribution mechanisms to share productivity gains where labour demand falls sharply, particularly in more automation-heavy scenarios.

For policy shaping AI innovation, the literature suggests private incentives tend to favour labour-replacing and socially harmful AI applications over labour-complementing and knowledge-augmenting uses.

Design should therefore focus on incentivising the market and the direction of AI development toward socially beneficial, labour-complementing, and productivity-enhancing applications. Suggestions include public funding and procurement that favours proworker and augmentation-oriented AI, policies that strengthen worker voice and control in AI adoption, and support for human capital investment where AI complements rather than substitutes for human skills.

The literature suggests how regulation can be designed. AI development is characterised by uncertainty, learning, and potentially irreversible harms, making static ex-ante regulation fragile. Regulation should therefore enable controlled learning while preventing socially costly over-adoption under uncertainty.

The literature gives recommendations of ways to achieve these goals. General advice is that policy should be adaptive and sector-specific rather than uniform. Suggestions for specific regulatory methods include regulatory sandboxes and staged deployment. These should be used in cases of high uncertainty regarding the impact of AI. Another recommendation is for liability regimes to be used that align responsibility with control and remain flexible as technologies evolve.

One of the larger topics in the literature regards impacts of AI on competition in the market. AI markets, particularly for foundation models, exhibit strong tendencies toward concentration due to economies of scale, scope, and reliance on scarce inputs.

The literature suggests policymakers and regulators should look at AI through a market competition lens. They should prevent excessive concentration while recognising that competition alone may not deliver socially desirable outcomes. This will involve ensuring access to key inputs such as data and computational resources, and monitoring and limiting efforts that reinforce entry barriers, such as vertical integration. If natural monopoly conditions were to emerge, policymakers should use regulatory oversight rather than rely solely on ex-post antitrust enforcement.

References

Acemoglu, D. (2021). *Harms of AI (NBER Working Paper No. 29247).* National Bureau of Economic Research. doi:https://doi.org/10.3386/w29247.

Acemoglu, D. (2025). The Simple Macroeconomics of AI. *Economic Policy, 40*(121), 13-58. doi:https://doi.org/10.1093/epolic/eiae042

Acemoglu, D., & Lensman, T. (2024). Regulating Transformative Technologies. *AER: Insights, 6*(3), 359-376. doi:https://doi.org/10.1257/aeri.20230353

Acemoglu, D., & Restrepo, P. (2018). The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment. *American Economic Review, 108*(6), 1488-1542. doi:10.1257/aer.20160696

Acemoglu, D., & Restrepo, P. (2020). Robots and Jobs: Evidence from US Labor Markets. *Journal of Political Economy, 128*(6), 2188-2244. doi:https://doi.org/10.1086/705716

Acemoglu, D., Autor, D., & Johnson, S. (2026). Building pro-worker artificial intelligence. *The Hamilton Project*. Retrieved from https://economics.mit.edu/sites/default/files/2026-03/Building%20Pro-Worker%20Artificial%20Intelligence.pdf

Acemoglu, D., Kong, D., & Ozdaglar, A. (2026). *AI, Human Cognition and Knowledge Collapse (NBER Working Paper No. 34910).* National Bureau of Economic Research. doi:https://doi.org/10.3386/w34910

Agrawal, A. K., McHale, J., & Oettl, A. (2026). *AI in Science (NBER Working Paper 34953).* National Bureau of Economic Research.

Brynjolfsson, E., Chandar, B., & Chen, R. (2025). Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence.

Eckhardt, S., & Goldschlag, N. (2025). *The Final Word (Until the Next One).* Economic Innovation Group.

EU. (2026, 03 12). *EU Artificial Intelligence Act*. Retrieved from The EU Artificial Intelligence Act: https://artificialintelligenceact.eu/

Filippucci, F., Gal, P., & Schief, M. (2024). Miracle or Myth? Assessing the macroeconomic productivity gains from Artificial Intelligence. *OECD Artificial Intelligence Papers, 29*. doi:https://doi.org/10.1787/b524a072-en

Guerreiro, J., Rebelo, S., & Teles, P. (2023). *Regulating Artificial Intelligence (NBER Working Paper No. 31921).* National Bureau of Economic Research. doi:https://doi.org/10.3386/w31921

Johnston, A. C., & Makridis, C. A. (2025). The Labor Market Effects of Generative AI: A Difference-in-Differences Analysis of AI Exposure. doi:https://dx.doi.org/10.2139/ssrn.5375017

Jones, C. I. (2026). *A.I. and our Economic Future (NBER Working Paper No. 34779).* National Bureau of Economic Research. Retrieved from http://www.nber.org/papers/w34779

Korinek, A. (2024). *Economic Policy Challenges for the Age of AI (NBER Working Paper No. 32980).* National Bureau of Economic Research. doi:10.3386/w32980

Korinek, A., & Stiglitz, J. E. (2018). Artificial Intelligence and its Implications for Income Distribution and Unemployment. In J. Gans, A. Goldfarb, & A. Agrawal, *The Economics of Artificial Intelligence: An Agenda* (pp. 349-390). National Bureau of Economic Research, Inc. Retrieved from https://www.nber.org/books-and-chapters/economics-artificial-intelligence-agenda.

Korinek, A., & Vipra, J. (2024). *Concentrating Intelligence: Scaling and Market Structure in Artificial Intelligence (NBER Working Paper No. 33139).* National Bureau of Economic Research. doi:10.3386/w33139

Kretschmer, M., Kretschmer, T., Peukert, A., & Peukert, C. (2023). The risks of risk-based AI regulation: taking liability seriously (CEPR Discussion Paper No. 18517). Retrieved from https://cepr.org/publications/dp18517

Massenkoff, M., & McCrory, P. (2026). *Labor market impacts of AI: A new measure and early evidence.* Anthropic.

Nordhaus, W. D. (2021). Are We Approaching an Economic Singularity? Information Technology and the Future of Economic Growth. *American Economic Journal: Macroeconomics, 13*(1), 299-332. doi:10.1257/mac.20170105

Schmid, J., Systma, T., & Shenk, A. (2024). *Evaluating Natural Monopoly Conditions in the AI Foundation Model Market.* Santa Monica: RAND. Retrieved from https://www.rand.org/pubs/research_reports/RRA3415-1.html

Stanford HAI. (2026, 08 11). *What is a Large Language Model (LLM)?* Retrieved from https://hai.stanford.edu/ai-definitions/what-is-a-llm

Trammell, P., & Korinek, A. (2023). *Economic Growth under Transformative AI (NBER Working Paper No. 31815).* National Bureau of Economic Research. doi:10.3386/w31815

Vipra, J., & Korinek, A. (2023). Market Concentration Implications of Foundation Models: The Invisible Hand of ChatGPT (Center on Regulation and Markets Working Paper #9). Retrieved from https://www.brookings.edu/articles/market-concentration-implications-of-foundation-models-the-invisible-hand-of-chatgpt/

World Economic Outlook Update. (2026). *Global Economy: Steady amid Divergent Forces.* International Monetary Fund. Retrieved from https://www.imf.org/en/publications/weo/issues/2026/01/19/world-economic-outlook-update-january-2026