

Submission to Senate Select Committee on Adopting Artificial Intelligence (AI)

Authors:

Dr Darcy W.E. Allen

Professor Chris Berg

Dr Aaron M. Lane

Contact Author:

Dr Darcy W.E. Allen, Deputy Director, RMIT Blockchain Innovation Hub, RMIT University.

Email: [REDACTED]

Date:

3 May 2024

Dear Committee Members,

We welcome the opportunity to contribute to this important inquiry into the adoption of artificial intelligence (AI) in Australia. We are a team of academics from RMIT University who are deeply involved in research on the digital economy through the RMIT Blockchain Innovation Hub and RMIT Digital 3.¹ Our research examines frontier technologies — notably generative AI and blockchains — from the perspective of economics, law, and public policy. We have collectively spent decades studying, using and building the technologies and business models of the digital economy. The breadth of our research informs our views on how generative AI provides opportunities for Australians. Please note that our submission below is adapted from prior submissions and expert witness appearances, including the House of Representatives Inquiry into the Use of Generative AI in Education² and the New South Wales Inquiry into Artificial Intelligence.³ These inquiries are ongoing.

Generative AI is a new example of a general purpose technology (GPT) that can be applied and built on across a wide variety of Australian industries. While there is widespread understanding of the economic importance of general purpose technologies, identifying their

¹ RMIT Blockchain Innovation Hub, <https://rmitblockchain.io/>
RMIT Digital 3, <https://www.rmit.edu.au/partner/hubs/digital3>

² Our submission to that inquiry is here:
<https://www.aph.gov.au/DocumentStore.ashx?id=f88c769c-c0a0-42da-a2e9-d5ca9aba7f89&subId=745186>

³ Our submission to that inquiry is here:
<https://www.parliament.nsw.gov.au/lcdocs/submissions/83215/0048%20Dr%20Darcy%20W.E.%20Allen.%20Professor%20Chris%20Berg.%20and%20Dr%20Aaron%20M.%20Lane.pdf>

precise applications and economic benefits is hard.⁴ We simply don't know what these technologies are for, and those valuable uses are diffuse across an economy and across time. There is growing evidence that generative AI provides some productivity benefits in complex ways.⁵ Rather than simple productivity gains, they can also reduce transaction costs and may reorganise economic activity.⁶ We and others have laid those benefits out elsewhere, including in specific sectors, such as our submission to the House of Representatives inquiry into generative AI in education.⁷

The focus of this inquiry on Generative AI adoption is welcome. The benefits of generative AI will ultimately be discovered through a process of adoption by both entrepreneurs and consumers. Yet an underappreciated and critical factor of generative AI is that this process is comparatively low-cost, open and decentralised. Trying to discover valuable uses for this suite of technologies is aided by the low cost of entry, both in terms of monetary cost and the use of human language chatbot-style interactions. Unlike previous general purpose technologies (such as the steam engine or electricity), Australian users can easily innovate with these technologies in their own local context. Our recent research reveals how employees can covertly discover applications without the direct oversight of management – a process of shadow user innovation.⁸ To ensure that effective discoveries are more widely adopted, managers will have to grapple with misaligned incentives with employees, perhaps incentivising them to share their discoveries. This is as true for workings in the public sector as it is in the private sector, where there are significant opportunities for public sector productivity gain if the technology is managed well.

While we are optimistic that generative AI will provide significant benefits to Australia, our approach in this submission is to focus on some of the main concerns that have been raised. There have been widespread calls to slow down the acceleration and application of

⁴ On general purpose technologies see Bresnahan, T. and Trajtenberg, M. (1995) 'General Purpose Technologies 'Engines of Growth'' *Journal of Econometrics* 65(1): 83-108

⁵ See, for instance:

Dell'Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., ... & Lakhani, K. R. (2023). Navigating the jagged technological frontier: field experimental evidence of the effects of AI on knowledge worker productivity and quality. Harvard Business School Technology & Operations Mgt. Unit Working Paper, (24-013).

Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187-192.

⁶ Potts, J., Allen, D. W. E., Berg, C., and Ilyushina, N. (2023). 'Large Language Models Reduce Agency Costs'. *Available at SSRN*.

⁷ See Submission 18, RMIT Blockchain Innovation Hub Researchers

<https://www.aph.gov.au/DocumentStore.ashx?id=f88c769c-c0a0-42da-a2e9-d5ca9aba7f89&subId=745186>

⁸ See Waters-Lynch, J., Allen, D. W.E, Potts, J., & Berg, C. (2024). Managing Generative AI in Firms: The Theory of Shadow User Innovation. Available at SSRN.

generative AI.⁹ The initial response to a new disruptive technology, let alone one that has human-like characteristics, is to apply a precautionary principle.¹⁰ The precautionary principle around new technologies involves overweighting the current tangible harms of the technology, and underweighting the potential future benefits. This approach poses a significant potential risk to ensuring Australia is at the forefront of digital technologies.

Our contribution in this submission is to address five main areas of concern that are often raised regarding generative AI. We structure our submission around the following questions:

1. Doesn't generative AI hallucinate?
2. Aren't the models biased?
3. Will generative AI take our jobs?
4. Is the technology controlled by a few large firms?
5. Do these tools undermine democracy?

Our aim is to provide evidence on how to think about these genuine concerns that people have about generative AI. Understanding and adapting to these concerns is critical if we are to enable the necessarily entrepreneurial experimentation of generative AI into Australian business models.

Doesn't generative AI hallucinate?

A common misconception about generative AI revolves around its intended purpose and functionality. Unlike traditional search engines designed for delivering accurate, factual information, generative AI operates as a prediction engine.¹¹ This key distinction underscores its primary purpose: fostering creativity rather than ensuring accuracy. As non-deterministic systems, generative AI models excel in creativity.¹² This creative ability propels their applicability across many new domains as a general purpose technology. But while the non-determinism of generative AI models is the source of their benefits, it also contributes to what are often termed as 'hallucinations' in their outputs. These are instances where the AI

⁹ See 'Pause Giant AI Experiments: An Open Letter', Future of Life Institute, 22 March 2023, Available at <https://futureoflife.org/open-letter/pause-giant-ai-experiments/>

¹⁰ Thierer, A. (2016). *Permissionless Innovation: The Continuing Case for Comprehensive Technological Freedom*. Mercatus Center at George Mason University.

¹¹ Agrawal, A., Gans, J., & Goldfarb, A. (2022). *Prediction Machines, Updated and Expanded: The Simple Economics of Artificial Intelligence*. Harvard Business Press.

¹² On the non-deterministic nature of generative AI models, see: Berg, C., Davidson, S., and Potts, J. (2023). 'Institutions to Constrain Chaotic Robots: Why Generative AI Needs Blockchain'. Available at SSRN 4650157.

generates content that — while potentially unique, creative and even plausible — may not be factual.

The community of generative AI users must evolve their understanding of the technology to incorporate the potential for hallucinations. They might also seek to use additional features of the technology (e.g. the capacity to browse the internet) and other approaches (e.g. manual search, judgement and expertise) to mitigate some of the risks.

There is already increasing awareness and scepticism regarding some of the outputs from these models. This scepticism is positive and welcome. It is part of a healthy use of consuming and applying generative AI outputs in practice. Humans co-evolve with technologies. As users grow more accustomed to the idiosyncrasies of generative AI, they are learning to navigate its capabilities and limitations. Greater user discernment is crucial and fosters a more informed and critical approach to the adoption of generative AI.

Aren't the models biased?

Generative AI models are regularly benchmarked against some idealised, perfectly unbiased AI. Signs of bias are quickly identified so that the technology can be slowed, or each bias can somehow be corrected. While the bias of generative AI outputs is a major and legitimate concern for many users and observers of the technology, this benchmarking approach is unrealistic and unhelpful. Such a standard is unattainable for several reasons, including the complexities of both technology and humans. Addressing this bias requires an understanding of the limitations of generative AI models, as well as a comparison to alternative tools or technologies to solve particular problems.

The biases in generative AI models are, in part, a reflection of the biases inherent in humans. These models are trained on vast datasets, often from accessible internet. Unsurprisingly biases from the datasets become embedded in the models. This is capturing the prevailing tendencies, preferences, and prejudices of the data it has been trained on. Attempts to correct for these biases, such as reinforcement learning through human feedback, might mitigate some of these biases, but it may also integrate new ones (from the humans in the reinforcement process).

The critical question for users of generative AI is not whether bias exists, but rather what types of bias are acceptable in specific contexts. Users must consciously navigate the trade-offs between different sources of bias: the averaged biases of input data in large language models (LLMs), biases introduced through safety and human reinforcement interventions, biases inherent in individual or group perspectives, and so on. This role of

humans in the loop of generative AI is essential. Humans dictate the level and type of bias users are willing to tolerate and correct for in their interactions with AI.

Awareness and acknowledgment of these biases are crucial. Users must not only recognise the potential biases in AI models but also actively respond to them. This response may involve adapting our norms and behaviours in how we interact with these technologies. As we become more cognizant of the generalised biases present in AI models, our engagement with them must evolve. This approach is similar to dealing with the challenges of hallucinations as described above. This evolution in interaction and norm-setting is a vital step in responsibly harnessing the capabilities of generative AI, ensuring that its benefits are maximised while its biases are acknowledged and managed. Importantly, the approach we have laid out here differs widely from a model-level approach of assuming we can create some unbiased generative AI model itself.

Will generative AI take our jobs?

There are significant potential productivity benefits from generative AI in the Australian economy. While estimates are necessarily predictive the early results suggest staggering productivity gains and improvements. For instance, a recent paper studying the productivity benefits of consultants with generative AI (compared to those who completed tasks without) showed across-the-board productivity gains.¹³ Indeed, as the author Ethan Molluck describes, “Consultants using AI finished 12.2% more tasks on average, completed tasks 25.1% more quickly, and produced 40% higher quality results than those without. Those are some very big impacts.”¹⁴

But simply because generative AI increases worker productivity does not mean that robots will take our jobs en masse. Unlike technologies that purely automate, generative AI applications typically require a process between a prompting-human and the technology. Generative AI is applied as a process of *co-production*. Indeed, generative AI expert Ethan Molluck at Wharton’s recent book was titled *Co-Intelligence*.¹⁵ Human expertise is needed to craft effective prompts, and to identify valuable problems and applications that generative AI models might help with. Furthermore, effective co-production with generative AI typically involves feedback loops and responses from human prompters, including an

¹³ Dell’Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., ... & Lakhani, K. R. (2023). Navigating the jagged technological frontier: field experimental evidence of the effects of AI on knowledge worker productivity and quality. Harvard Business School Technology & Operations Mgt. Unit Working Paper, (24-013).

¹⁴ Molluck, Ethan (2023) ‘Centaur and Cyborgs on the Jagged Frontier’ *One Useful Thing*, 16 September, <https://www.oneusefulthing.org/p/centaurs-and-cyborgs-on-the-jagged>

¹⁵ Molluck, E. R. (2024). *Co-intelligence: Living and working with AI*. Penguin Random House.

almost-entrepreneurial process of making judgements over outputs and adapting to them.¹⁶

This suggests a significant role for humans remaining in the process. Many of the productivity improvements through generative AI will come through replacing tasks not jobs.¹⁷ Co-production is not merely about automating processes but enhancing them through a deep understanding of the nuances involved in each task.

Generative AI is fundamentally a bottom-up technology, requiring significant human input not only in guiding these systems but also in making critical judgments about their outputs.

These are not standalone solutions but tools that need to be applied to a diverse range of tasks, each requiring deep, contextual local knowledge. The human element in directing and interpreting the outputs of generative AI is indispensable, emphasising the technology as an augmentative tool rather than a replacement for human skills.

Is the technology controlled by a few large companies?

As in all sectors, competition is a crucial driver of innovation and growth in the generative AI sector. Competition in the generative AI sector manifests at different levels.

- Competition between the underlying foundational large language models.
- Competition between the general interfaces of those models.
- Competition in the “wrappers” and applications built on top of the models.

For instance, there is differentiation in products such as Open AI’s ChatGPT and Microsoft’s Bing, which, despite being based on the same underlying model, offer distinct user experiences and functionalities. This diversity in applications highlights the dynamic and competitive nature of the generative AI field.

The current regulatory debates around generative AI should consider the importance of these competitive dynamics. In the first instance, the current nascent stage of AI technology means that we may not fully understand the implications or the full scope of what we are regulating. This makes it difficult to regulate, and leads to prominent calls for major changes such as licensing AI models.¹⁸ Such approaches carry a significant risk of regulatory lock-in and capture. It is crucial for lawmakers and regulators to recognize that any regulatory

¹⁶ Potts, J. (2023). ‘The Use of Knowledge in a Digital Economy’. *Available at SSRN 4440771*.

¹⁷ On the impact on tasks in an Australian context see Walkowiak, E. and MacDonald, T. (2023) ‘Generative AI and the Workforce: What Are the Risks?’ *Available at SSRN 4568684*.

¹⁸ See, for instance, Bartz, D. and Dastin, J. ‘OpenAI chief goes before US Congress to propose licenses for building AI’, Reuters, 16 May 2023 Available at <https://web.archive.org/web/20230516122128/https://www.reuters.com/technology/openai-chief-goes-before-us-congress-propose-licenses-building-ai-2023-05-16/>

impositions, such as licensing requirements, can inadvertently lead to a concentration of power in the industry, thus stifling competition.

We expect that the market for generative AI models and products will be diverse. It will consist of both large closed corporate models and a rich, diverse array of open-source models. This diversity is desirable. Indeed, some strategic dynamics may lead to major models remaining open source. For instance, Meta recently released their Llama 3 model open source, arguably as a way to undermine the commercialisation options for competitors who have kept their models closed. We have also seen a suite of other models being released open source.¹⁹

Policymakers must understand that their regulatory actions could constrain development of generative AI and its application across the economy, distorting the market towards closed proprietary models. Thoughtful regulation should aim to protect consumers and society while also fostering an environment where diverse AI models and applications can thrive, thereby ensuring a competitive and dynamic generative AI ecosystem.

Do these tools undermine democracy?

As with all new communication technologies, there have been concerns that generative AI models will impact some of our foundational institutions, including democracy. But generative AI can have both positive and negative impacts on our democratic processes.

On the one hand, there are legitimate concerns around deepfakes (AI-generated original media or manipulated content such as images or videos produced for the purpose of misleading or deceiving) and disinformation more broadly. In the political landscape, there is a concern that this content will be used to mislead voters, undermine political processes, and potentially influence election outcomes.

A range of potential solutions to this problem exist. In the first instance, there is a suite of existing laws that aim to bolster integrity in our electoral advertising processes. For instance, under electoral laws material must be properly authorised and there are restrictions on political advertising. There are also existing laws dealing with foreign interference. While there might be a temptation for lawmakers to prohibit deepfakes altogether, such a ban would inhibit innovation in areas where the underlying technology has a range of legitimate uses, such as in the creative arts, film, entertainment and gaming sectors.

¹⁹ See <https://www.wired.com/story/metasp-open-source-llama-3-nipping-at-openais-heels/>

Other technologies also act as a countervailing force to the rise of deepfakes. Innovations in generative AI sit within a broader stack of technologies that are also rapidly evolving. Blockchain technology, for instance, provides a new type of digital infrastructure for the timestamping and verification of data, providing additional trust to information consumers. We anticipate that such mechanisms will become more deeply integrated into existing digital platforms, such as social media, and form the basis of new blockchain-native ones. Blockchains can enable censorship-resistant and time stamped digital signatures of data, making it harder for disinformation to spread. But there is of course a challenge here: voters will need to adapt their norms to seek out such trusted sources. Notably, consumers of digital communications are already becoming more vigilant, such as seeking out “authorised” accounts on platforms such as X (formerly Twitter). Nevertheless, we are optimistic that voters (alongside and aided by political candidates and journalists) will adapt to new challenges, as they have with other technological advancements.

On the other hand, generative AI can be a positive force for Australian democracy. It can help voters to better understand political debates, legislation, policy proposals and undertake data analysis. Rather than undermining the democratic process, these tools can be used as a positive development to create a more informed and engaged Australian electorate. We are already seeing the diffusion of these tools in frontier blockchain-based digital governance processes — such as summarising proposals in Decentralised Autonomous Organisations — and we anticipate that these tools will also be used to expand the effectiveness of Australian democracy.²⁰

We welcome the opportunity to expand on any of our contributions here by appearing before the committee. We look forward to seeing the result of this important inquiry and hearing from you in due course.

Regards,

Dr Darcy W. E. Allen

RMIT Blockchain Innovation Hub, RMIT University

Professor Chris Berg

RMIT Digital 3 and RMIT Blockchain Innovation Hub, RMIT University

Dr Aaron Lane

RMIT Blockchain Innovation Hub, RMIT University

²⁰ On blockchains and democratic governance see Allen, D. W., Berg, C., & Lane, A. M. (2019). *Cryptodemocracy: How Blockchain can Radically Expand Democratic Choice*. Rowman & Littlefield.