Amazon Commercial Services Pty Ltd and Amazon Web Services Australia Pty Ltd (AWS) (collectively, Amazon) provide the following responses to questions on notice which arose during a hearing held by the Senate Select Committee on Adopting Artificial Intelligence (Committee) on 16 August 2024, as well as additional written questions Amazon received from the Committee's secretariat on 20 August 2024.

Questions from Senator Pocock, sent to Amazon by the Committee's secretariat on 20 August 2024

Questions 8, 9 and 10

Question 8: Does Amazon foresee a risk where the economic benefits of AI could translate into disproportionate political influence, thereby skewing democratic processes in favour of those who control AI technologies?

Question 9: How does Amazon propose that global AI standards and regulations be coordinated across different jurisdictions to prevent conflicting governance approaches that could be exploited to weaken democratic institutions?

Question 10: Considering Amazon's involvement in international AI initiatives, how will the company ensure that these standards prioritise the protection of democratic values over commercial interests?

Security and trust are essential to fair democratic elections, and Amazon is committed to secure, safe and fair democratic processes. This is why we joined other technology companies in a pledge to work together to identify, counter and prevent deceptive AI content to protect the integrity of elections. This pledge is known as the Tech Accord to Combat Deceptive Use of AI in 2024 Elections (also referred to as the 'Munich Tech Accord').

The Munich Tech Accord is a set of voluntary commitments to develop and deploy technology countering harmful AI-generated content meant to deceive voters in democratic processes. These commitments complement Amazon's existing efforts through the White House Voluntary AI Commitments, the US Artificial Safety Institute Consortium, the UK and South Korea's AI Safety Summits, and more, to advance safeguards against deceptive activity and to protect the integrity of elections. We also leverage a range of third-party organisations to help us coordinate across industry, civil society, and government on issues related to election disinformation and other areas of AI safety, including the Frontier Model Forum, and the Global Challenge to Build Trust in the Age of Generative AI, a flagship project of the G7 to fight misinformation, in cooperation with the UN, the Organisation for Economic Co-operation and Development, and the Global Partnership on AI. Amazon is a member of the National Institute of Standards and Technology AI Safety Institute Consortium, in which we participate in the working group developing safety guidelines and standards around synthetic media safety.

Amazon is also undertaking a number of other efforts to ensure that the integrity of elections and the democratic process is preserved. A key tool is the development of provenance signals to identify the origin of AI generated content and enabling individuals to detect those signals. To that end, Amazon has

focused on building tools for our AI products and services that could be at higher risk for creation of content that could be used for election disinformation. Importantly, our image models are designed for the purpose of creative content generation, and not for the depiction or information retrieval of real-life people or events. All images generated by Titan Image Generator contain an invisible watermark by default, which is designed to help reduce the spread of disinformation by providing a mechanism to identify AI-generated images. AWS's Service Terms prohibit the alteration or removal of the watermark. Additionally, we are working to implement the Coalition for Content Provenance and Authenticity (C2PA) standard in our generative AI models too.

We also have content policies across our services that mitigate the risk of disinformation and other harms to our customers. For example, AWS's Responsible AI Policy prohibits use of our AI tools to depict a person's voice or likeness without their consent or other appropriate rights, including unauthorised impersonation. We continue to find new ways to engage in and encourage information-sharing among companies and governments regarding this topic, as AI technologies continue to evolve.

Questions 11 and 12

Question 11: What are Amazon's views on the potential disruption AI could cause to labour markets, and how might this economic upheaval affect democratic stability, especially if certain groups feel disenfranchised by the rapid technological changes?

Question 12: How does Amazon plan to mitigate the risks associated with AI-driven job displacement, and what role should government policy play in ensuring that AI advancements do not exacerbate social and economic divides that could destabilize democracies?

At Amazon, we believe we have a responsibility to help prepare people to be a part of the digital economy. We strongly believe that the future of a prosperous Australia is tied to its digitally skilled workforce, and that digital literacy skills are fundamental for modern Australian society. Australian workers who use advanced digital skills are propelling the country's growth - they add an estimated AU$41 billion to the country's annual gross domestic product (GDP). A 2024 study by Access Partnership found that 90 per cent of Australian employers expect their organisations to use AI-related solutions by 2028, while workers anticipate having AI skills would have a positive impact on their careers, increased efficiency, and job satisfaction. Still, while hiring AI-skilled talent is a priority for two in three employers in Australia, 75 per cent are struggling to find the AI talent they require. Similarly, more than seven in ten (71 per cent) of Australian employers say it is challenging to find people with the right digital skills. Building a future-ready workforce is key to overcoming hiring challenges and being able to rapidly adapt to future technologies, including generative AI.

There is a looming AI skills gap that needs to be addressed so that Australia is positioned to unlock the full benefits of AI. This requires collaboration between governments, industries, and educators to help employers implement AI training programs, guide workers in matching their AI skillsets to the right roles to harness their newly acquired AI capabilities, and make the most of this evolving technology. For example, our submission to, and appearance before, the House of Representatives' Standing Committee on Employment, Education and Training's inquiry into the use of generative AI in the Australian education system, we recommended that federal and state governments: work with the tertiary education sector and appropriate industry partners to design and develop digital literacy and technology competency training within the teaching qualification and broader teacher accreditation and professional developments; create a National Advisory Group examining teaching, learning and assessment practices alongside current and emerging technologies in Australia; and ensure technology-neutral consumer protections and sector-specific regulations remain fit-for-purpose, and require the adoption of risk-based AI governance practices where appropriate.

Amazon shares the government's vision of addressing the growing tech skills shortage in Australia to have 1.2 million people in tech jobs by 2030. AWS is committed to bridging the cloud skills gap in Australia, and have trained more than 400,000 people across Australia with cloud skills since 2017. This underpins our commitment to provide free cloud computing skills training to 29 million people globally by 2025. We're also focused on making AI, ML, and generative AI skills training and education accessible to everyone through our new 'AI Ready' initiative, which offers a suite of free AI and generative AI training course, aligned to both technical and non-technical roles, so that anyone can build AI skills. With the rapid adoption of cloud-enabled technologies like generative AI, more work needs to be done to build an AI-skilled workforce to unleash a culture of innovation and drive productivity in this dynamic region of the world.

Questions 13, 14 and 15

Question 13: What training do those classifiers have that are looking at misinformation?

Question 14: How often are their guidelines updated?

Question 15: Has Twitter/X turning reducing access to their "firehose data" feed limited your ability to identify sources of misinformation.

At Amazon, we build foundation models with responsible AI principles in mind at each stage of its comprehensive development process. We engage in rigorous testing, assessments and improvements to our services, and have implemented safeguards into our generative AI offerings, including combatting risks around synthetic content by offering built-in watermarking with Titan Image Generator; and applying filters on user inputs and model outputs for our Titan models to reduce the likelihood and spread of harmful content. This aligns with AWS's commitment to promoting responsible AI by reducing the spread of misinformation.

To help prevent potential misuse, Amazon Bedrock implements automated abuse detection mechanisms to identify potential violations of AWS's Acceptable Use Policy (AUP) and Service Terms, including the Responsible AI Policy or a third-party model provider's AUP. This automated abuse detection includes:

- Categorising content — We use classifiers to detect harmful content (such as content that incites violence) in user inputs and model outputs. A classifier is an algorithm that processes model inputs and outputs, and assigns type of harm and level of confidence. We may run these classifiers on both Titan and third-party model usage. The classification process is automated and does not involve human review of user inputs or model outputs.
- Identifying patterns — We use classifier metrics to identify potential violations and recurring behavior. We may compile and share anonymised classifier metrics with third-party model providers. Amazon Bedrock does not store user input or model output and does not share these with third-party model providers.

Once our automated abuse detection mechanisms identify potential violations, we may request information about our customers' use of Amazon Bedrock and compliance with our terms of service or a third-party provider's Acceptable Use Policy. In the event that they are unwilling or unable to comply with these terms or policies, AWS may suspend their access to Amazon Bedrock. We can't comment on other services.