UNSW CANBERRA | Public Service Research Group

Australian National University

Research Snapshot

The future of Generative AI in policy work

Jade Hart, Helen Dickinson, Kate Henne, Vanessa McDermott, Shibaab Rahman & James Connor

Report Overview

Purpose of this report

The UNSW Public Service Research Group offers contemporary research-based thinking about topical themes for public service and the public administration community.

This report outlines high-level findings from a national research study exploring Generative Artificial Intelligence (GenAI) and its use in government policy work in Australia It seeks to summarise perspectives on this important topic in terms of how these applications are being used and what some of their risks are. The report sets out future priorities for research and practice.

Through this report we seek to help translate research into practice and build academic debate. It provides a short and accessible summary of the findings and implications pending the preparation of academic publications.

Acknowledgements

The authors of this report would like to acknowledge the traditional custodians of the lands on which this report has been written, reviewed and produced. We pay our respects to their Elders past, present and future.

Thank you to the participants from state, territory, and federal governments who participated in interviews. We greatly appreciate the time these stakeholders took to provide their valuable insights to us.

We would like to also express our thanks to Paul Hubbard of AI CoLab who reviewed the report as a critical friend.

This research was funded by the School of Business UNSW Canberra.

We asked Copilot to identify any errors in the report and any sentences that were long or difficult to follow. Copilot did not identify any errors but did identify some areas that the language could be edited for simplicity.

Suggested citation

Hart, J., Dickinson, H., Henne, K., McDermott, V., Rahman, S., & Connor, J. (2025) The future of Generative AI in policy work. Canberra: Public Service Research Group.

Contact

For further information on this paper please contact UNSW Canberra.

Contact: General Enquiries, T: +61 2 5114 5670 and E: PSRGinfo@unsw.edu.au

Disclaimers

This report serves as a summary to prompt a wider conversation regarding GenAI. It is written for a non-technical audience. This report was conducted in accordance a methodology designed to elicit interviewee perspectives on specific questions. It is not necessarily a comprehensive review of the growing literature relating to the topic area. It is reproduced for general information and third parties rely upon it at their own risk.

Executive Summary

This report by the Public Service Research Group summarises national research examining the current and anticipated use of Generative AI (GenAI) in policy work. Between November 2024 to January 2025, 37 senior public servants from 22 state, territory, and federal government agencies participated in research to understand how GenAI is used in policy work, including perspectives of risks and enablers for adoption.

Current adoption of Generative AI

Many government agencies are using GenAI to support their goals. However, opinions on GenAI's potential varies among senior public servants. Some believe it can transform policy work, while others are cautious due to perceived risks. Past challenges with technology implementation, including learnings from policy failures identified from the Royal Commission into the Robodebt Scheme, are influencing how senior public servants perceive advancements in the application of GenAI. There is widespread agreement that the use of GenAI in policy work requires adequate human oversight.

It is for these reasons that GenAI is likened to a resource that can build on and augment human capabilities. The ethical and moral complexities within matters of policy mean that the use of GenAI in the early stages of adoption has focused on internal administrative and analytical functions, particularly those that rely on administrative data (e.g. population or expenditure datasets). However, there remain sensitivities regarding the application of GenAI to certain activities that are perceived as potentially compromising citizen trust and confidence, such as government external-facing service delivery activities. These risks emerge in tension with demand for additional public service capacity to address other areas of need.

There are currently several GenAI tools being used in policy work. The availability of tools via websites at no direct cost to users has increased access to and adoption of GenAI. Agencies have refined perspectives of explicitly authorised use of GenAI tools. Alongside these workplace activities, the public sector workforce has been building GenAI skills using free tools on their personal devices and realising its benefits in their everyday lives. Yet these skills do not always readily transfer to the government workplace, creating forms of not explicitly authorised or even prohibited forms of GenAI use.

Current applications of GenAI tools in policy work include:

- Administrative applications using Microsoft Copilot.
- New content and idea generation using GenAI tools such as ChatGPT as a contemporary resource akin to a website search to more tailored GenAI tools such as those produced by Dragonfly Thinking for problem-solving and analytical policy tasks.
- Building an agency-specific knowledge base trained on agency data, policies, legislation, procedures, and other guidance.

- Summarising volumes of information or documentation. Within agencies, policymakers are using GenAI to summarise and retrieve information. Service delivery applications include the use of Grammarly, Lyrebird AI, Heidi AI, or ChatGPT by practitioners to prepare documentation such as client records.
- Quantitative and computer programming functions, such as integrating GenAI coding within econometric models as an extension of existing AI analytical approaches.

The GenAI adoption path has so far taken an incremental approach through trials around specific technologies within relatively constrained projects. Consequently, there are questions as to whether transformational change can be realised through disparate projects.

Risks and concerns surrounding Generative AI

The use of GenAI tools has been cautious to address risks and concerns. GenAI outputs can be biased and inaccurate, which can affect policy quality. Publicly available GenAI tools can be like a 'black box,' making it hard for the government to ensure transparency. These tools can also reinforce existing biases in the data. There is also a need to ensure GenAI adoption of mitigates the risks to data sovereignty, particularly First Nations data sovereignty.

The adoption of GenAI is guided by a need to realise value for the given investment. The costs of GenAI adoption can be considerable, spanning change management, licenses and infrastructure, and workforce development. The environmental costs are also seen as significant, with additional energy and water consumed to a degree that escalating patterns of use is expected to result in environmental impact.

Enablers for Generative AI adoption in policy work

There are several enablers that are identified to support the adoption of GenAI in policy work. These enablers seek to establish the conditions for success by:

- realising short-term potential benefits
- establishing the long-term innovation/experimentation that has safeguards, and
- mitigating risks and potential for harm.

Firstly, there is a need to formalise the intended role of GenAI in policy work. A statement regarding why governments are investing in GenAI is critical to building understanding of its intended contribution to high quality policy work. Through a GenAI strategy, measures of value can be further identified, monitored, and assessed. Secondly, it is important to enhance critical AI literacy amongst the senior executive and workforce so that it aligns with other public service values. There remains a need to maintain a pipeline of expertise and skills in policy craft, including the features that reflect distinctively human forms of analysis and reasoning, while also understanding these tools have an impact on

how public servants complete work-related tasks. The widespread use of these tools is expected to contribute to shifts in how individuals and teams work. Pathways for career development need to be recalibrated to consider whether and how GenAI can augment policy work. Thirdly, we emphasise the necessity of ensuring that GenAI adoption occurs in line with regulatory and governance arrangements that account for the breadth of agencies involved in these functions.

Current GenAI tools predict content based on patterns observed within the historical data; they provide plausible answers to questions based on available data. Some complex policy issues, however, require solutions that cannot emerge from what has been done before. Further research is necessary to understand whether GenAI tools can provide public servants with tools and opportunities to refine and transform policy craft; or a 'black box' that might function in ways that are counter to transparency and accountability.

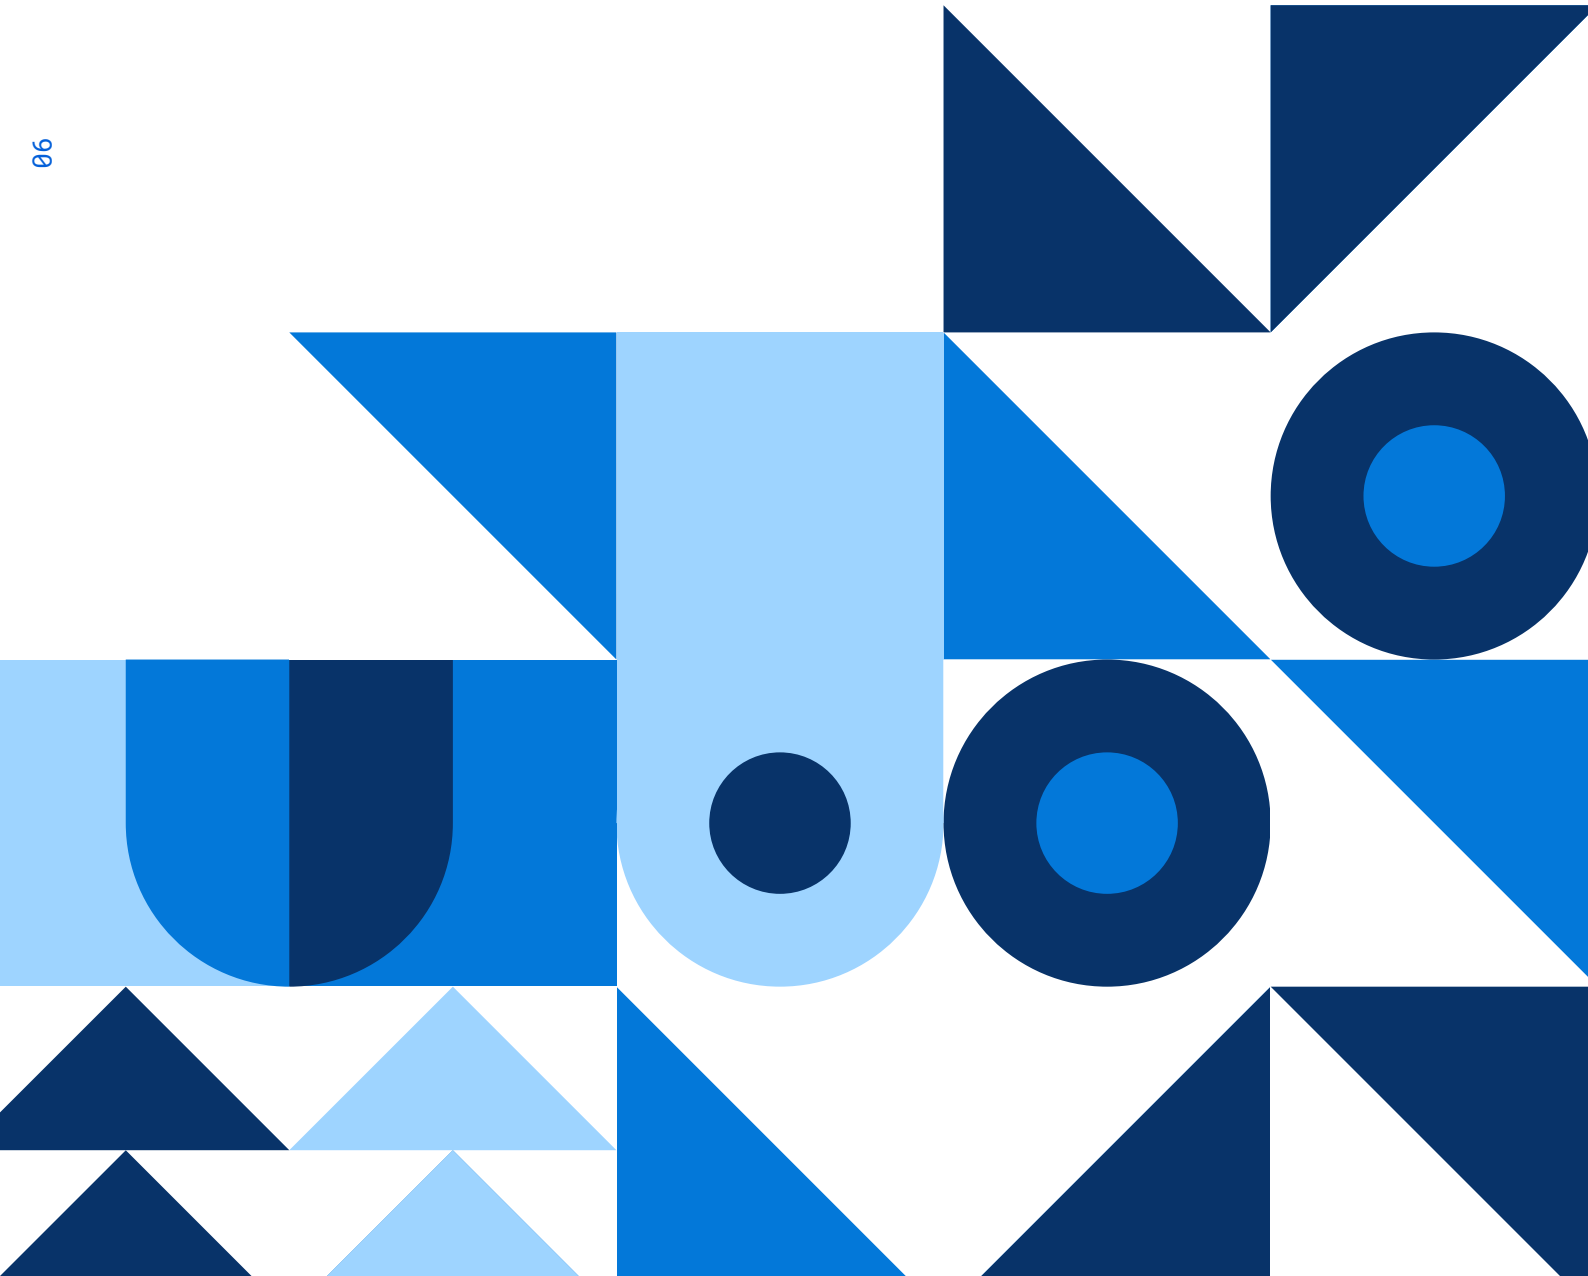

Contents

1. Introduction 8

1.1 Context 8

1.2 Aims and Approach 9

2. Findings 11

2.1 Current adoption of Generative AI 11

Views of GenAI's transformative potential spans hesitance to enthusiasm 11

GenAI will augment human capabilities and not replace them 12

Public servants prefer GenAI being kept away from external-facing activities 13

Availability of GenAI tools can lead to unauthorised forms of use 14

Examples of current use of GenAI tools in policy work 14

The challenge of achieving transformational change via incremental projects 16

2.2 Risks and concerns surrounding Generative AI 17

Data accuracy and risk of bias within the context of obligations 17

Threats to government's ability to retain sovereignty and control over its data 18

Challenges in achieving return on financial investment in GenAI 18

There are concerns over GenAI's environmental impact given sustainability aims . 19

2.3 Enablers for the adoption of generative AI in policy work 20

Strategies are needed to build clarity, trust, and confidence in GenAI 20

Senior executive capability is needed to deliver GenAI oversight 21

Effective use of GenAI needs workforce training, development, and support 22

There needs to be alignment across regulatory and governance arrangements . . 22

3. What can we take from this research 24

3.1 Implications for policy and practice 24

3.2 Areas for future research 25

3.3 Limitations 26

4. Summary 27

5. References 28

1. Introduction

1.1 Context

In recent years, we have seen significant developments in digital technologies. Increased processing power and decreased prices for many technologies make them more accessible and provide more potential for use in a wider range of contexts. The field of Artificial Intelligence (AI) has made significant impact on several areas of public services, such as through chatbots and virtual assistants, data analysis, visualisation, and text recognition and processing (1).

> **Definition:** Artificial intelligence (AI) – an umbrella term that encompasses the broad field of research and constellation of technologies in computer science focused on creating systems and technologies that mimic the level of human skill and ability. Along with deep learning, robotics, expert systems, and natural language processing, machine learning is a subset of AI (1).

This research study focuses on one form of AI, known as Generative AI (GenAI). As Feuerriegel et al. outline, the term generative in the context of AI refers to a system's ability to generate new material independent of human involvement (2). Content generated might be text, image, video or audio. There are a broad range of potential applications for GenAI in the work that governments do, including, for example, supporting the development of chatbots, improving data management, and supporting employee productivity (3).

> **Definition:** Generative AI – A further subset of machine learning focused on the generation of outputs, such as text, images, video, or audio based on learned patterns and data inputs. Applications of generative AI seek to output content that is contextually relevant and similar to outputs created by humans given the same task (1).

It is clear there is potential for GenAI to be used in several ways in policy work. By policy work, we refer to the full expanse of activities involved in the design and implementation of public policy. In this research, we sought to understand perspectives of senior public servants who are using or who see potential to use these tools in policy work. While there are expectations that the use of GenAI may lead to positive gains, there has been significant discussion about the potential risks of these activities (3). Moreover, some applications of GenAI may be challenging in the context of policy work, especially where public servants work with sensitive data. These challenges interface with bigger questions about the delivery of public good, which can be notoriously difficult to define (4).

The potential of GenAI to transform work practices in different sectors has received

widespread attention. While some commentators believe the use of GenAI will be transformative for governments and allow them to do work more efficiently and effectively, others have significant concerns about the use of these applications in terms of their accuracy, safety, and potential for bias (4-6). In this study, we were specifically interested in the ways that government agencies and individual public servants do and might use GenAI to undertake policy work, including how they perceive GenAI-informed policy work. We wanted to explore whether public servants see GenAI as a set of technologies with the potential to disrupt aspects of how governments undertake policy work.

Some state, territory, and federal governments are progressing legislative, regulatory, and policy reforms in support of GenAI. The Commonwealth Government stated its primary focus is to ensure the "rollout of digital applications and AI systems across Australian workplaces is carried out in a fair, ethical, safe, and responsible manner" (1, p. 10). The Digital Transformation Agency (DTA) is leading the Australian Government's digital and ICT transformation, and provides strategic, policy, and expert investment advice for government agencies, including the development of whole-of-government guidance on safe and responsible use of GenAI (1). Alongside these efforts, state and territory governments are independently developing governance arrangements. This means that regulatory and governance approaches are not centralised with several approaches to jurisdictional guidance on GenAI use in government policy across the states, territories, and the Australian Public Service.

Given the recency with which GenAI tools have become available, there is a lack of research that explores their use in policy work in Australia. Against this background, we undertook research that is exploratory in nature to examine some early experiments with GenAI. We sought to understand how senior public servants perceived the strengths and weaknesses of these applications, the potential of GenAI for policy work, and what needs to happen to make GenAI safe to use in policy work.

This report sets out findings of this national research study, specifically:

1. Current adoption of GenAI
2. Risks and concerns surrounding GenAI
3. Enablers for the advancement of GenAI adoption

This report provides a high-level view of responses to these questions. We conclude by setting out implications for policy and practice and areas for future research.

1.2 Aims and Approach

This exploratory research involved interviews with senior public servants holding executive roles within state, territory, and federal governments. Senior public servants were eligible based on their knowledge of the use of GenAI in government policy work. Several government publications were also examined for the purposes of gaining familiarity with public sector developments in Australia. This included a range of strategies, guidance,

and evaluation reports across a several state, territory, and federal governments. We used convenience sampling and snowball sampling techniques to recruit participants for this research. This process involved reviewing publicly accessible information and websites to identify individuals that met the eligibility criteria, referral to others who fit the participant eligibility criteria, and circulation of general information about the study through professional networks.

In total 22 state, territory, and federal government entities participated in this research. Across the Australian Commonwealth Government, nine entities were engaged, and agencies from six Australian states and territories were also represented. Entities spanned a mix of departments and agencies, including those with technology and non-technology specific portfolios. Ethical approval for this study was acquired through the UNSW Canberra Research Ethics Committee (iRECS7260). In accordance with ethical approvals, the details of interviewees and associated agencies have been kept anonymous.

The final sample comprised 32 interviews with 37 individuals. Individuals retained a mix of policy and technological expertise. Interviews were conducted online between November 2024 and January 2025. Interviews were an average of 33 minutes duration (23-52 minutes).

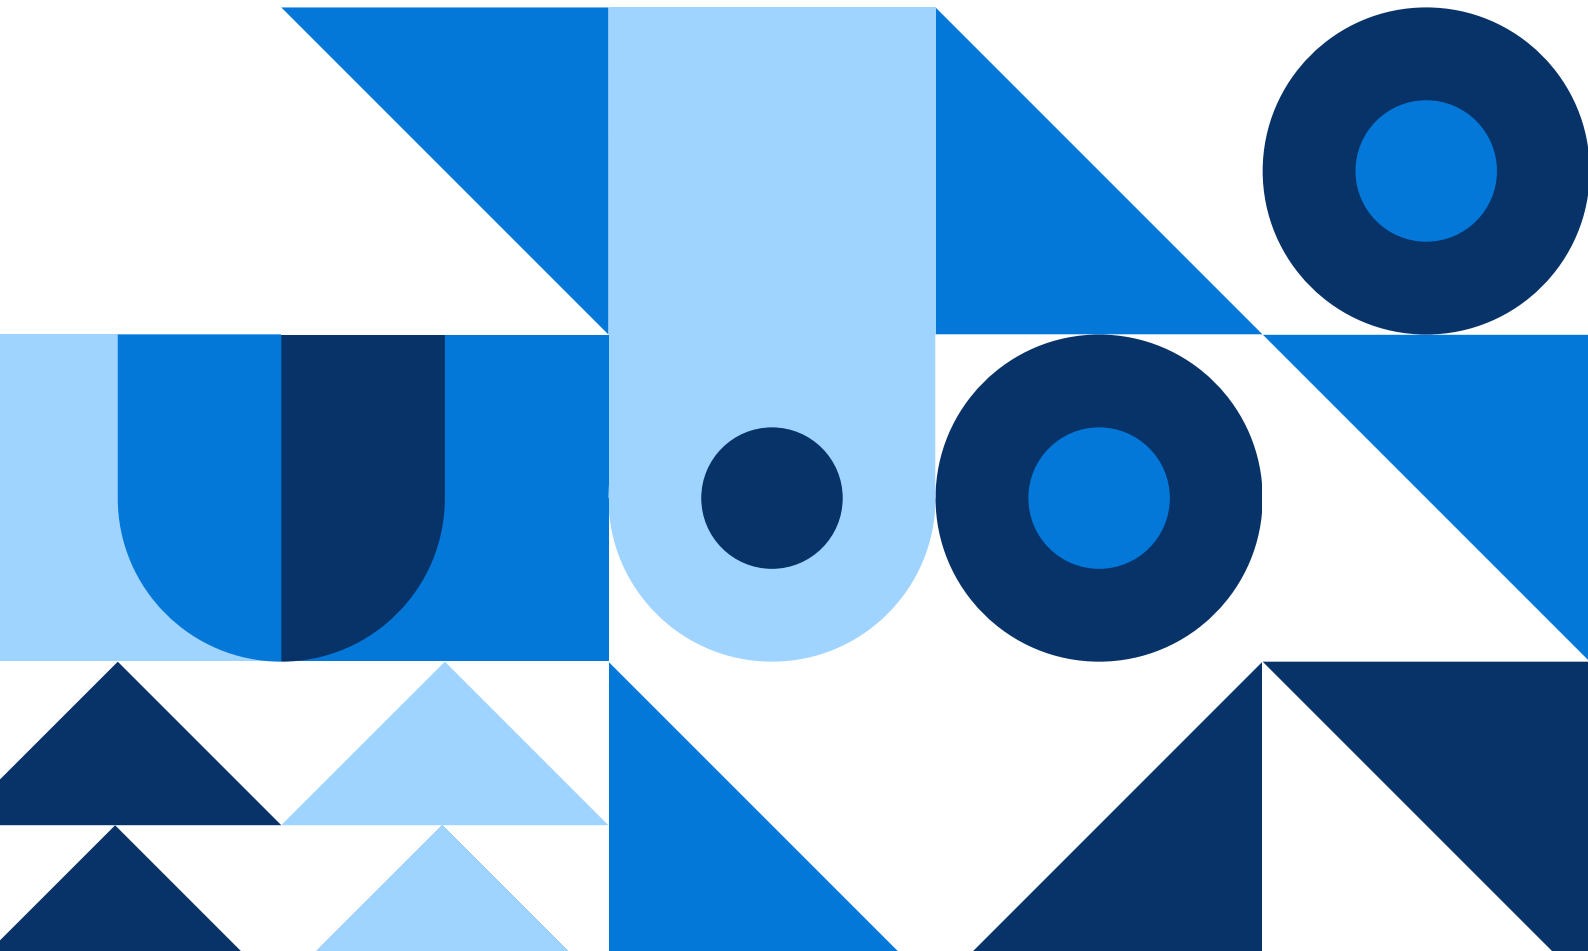

2. Findings

In this section we present key themes that emerged through analysis of data undertaken by the research team. We first discuss how GenAI is being adopted, followed by an exploration of the expressed concerns surrounding GenAI. Finally, we provide a summary of the enablers for the adoption of GenAI in policy work.

2.1 Current Adoptions of Generative AI

There were a range of views on the potential of GenAI to transform policy work, ranging from strong support for the possibilities of its application through to reluctance to engage with it. Most participants conveyed a shared perception of GenAI, seeing it principally as a means to augment the capabilities and capacity of humans, rather than as a replacement for public servants. Currently, GenAI is largely described as being focused on internal functions, rather than in client-facing functions. Participants described publicly available and widely accessible GenAI tools as being used in parallel with agency-led trials and applications of explicitly authorised tools. This section describes these perspectives within the context of the current adoption of GenAI across state, territory, and federal governments.

Views of GenAI's transformative potential spans hesitance to enthusiasm

Interviewees expressed several perspectives of the transformational potential of GenAI in policy work on a continuum. The continuum spans those who believe it will fundamentally change the ways this is undertaken to those who saw GenAI as a set of over-hyped tools. The latter considers there is the potential for significant risk that will lead to little change in how the work of public servants is undertaken.

A number of interviewees felt they were under sustained pressure to respond to increasingly complex policy issues at scale and also to do this in an efficient way. Within this context, they believed it was timely to re-examine how policy work is done and saw GenAI as a key part of a potential toolbox to do this. Some interviewees framed GenAI as part of a 'Fourth Revolution of Work' that provided public servants with an opportunity to revisit the fundamentals of their craft and transform how it is done. As Participant 2 comments: "Why improve the candle when you could use a light bulb?" This simple statement conveys several important implications. GenAI was seen as a qualitatively different tool that might significantly transform how policy work is currently done, including the pace and reliability with which outputs can be produced.

However, others were less positive about the potential of GenAI tools. Some interviewees recounted that senior public servants were reluctant to engage with GenAI, reflecting on how a lack of familiarity and understanding of these tools has implications for how they could ensure effective oversight for its implementation within their work units. Participant 23 explains, "So there's a bunch of public servants, including a decidedly large proportion

of senior executives, who are in the 'this just isn't happening to us and we're just going to pretend that it's all going to go away and it's just a whole lot of hype and there's nothing to see here' category." In the absence of clear direction on how to use GenAI or understandings as to how to use it within existing functions, there was a perception that the personal investment required of senior public servants to manage these changes outweighed the benefits.

Assessments of the potential transformative influence of GenAI were described by some as being tempered by the visibility and findings of the Royal Commission into the Robodebt Scheme, even though fundamental features of the scheme were reflective of flaws in human calculations and decision-making rather than automated technologies alone. Participant 31 describes the need for governments to rebuild credibility in relation to their technology oversight and implementation, highlighting, "The majority of the problems with Robodebt were at a human level... What's depressing about Robodebt is I feel that those obligations were already in place and not taken seriously... I do worry about people not taking their duties, both professionally and ethically, seriously, and interrogating what they get out of AI systems. It requires genuine, concentrated effort from each individual."

As shown in this quote, Robodebt shows how cultural conditions contributed to inadequate policy design and program governance, which reveal risks that emerge in environments that don't cultivate critical skills or foster conditions for critical discussion on technology use in support of aims. Despite cultures of technology adoption becoming more risk averse and less attuned to the necessity of ethically defensible policy and decision-making, Participant 2 cautions that the adoption of GenAI tools risks contributing to "learned helplessness" in the absence of adequate analytical skills to use GenAI in ways might transform policy work at face value yet ultimately compromise quality.

GenAI will augment human capabilities and not replace them

Regardless of individual views on the transformational potential of GenAI, the vast majority of interviewees expressed a belief that these tools would not lead to the replacement of public servants but would instead augment their capabilities. Participant 27 explains, "It is augmenting us. It is amplifying us." This viewpoint has a few important elements to it.

Firstly, interviewees believe that GenAI will not replace public servants because many aspects of policy work involve decision-making about ethical and moral issues, the responsibility of which lies with public servants. Many of the existing legislative frameworks require that decisions must be made by a specific accountable officer, which cannot be a technological application or agent. To illustrate, Participant 20 challenges a view that legislative powers and responsibilities could be enacted by something other than a human who was deemed to be competent to conduct functions to a satisfactory level, posing the question, "how will this mosaic of obligations be delegated away to generative AI?", and prompting consideration of the legitimacy of a scenario of an audit and risk committee that was enacted by a GenAI tool. Interviewees explain the necessity of ensuring 'a human in the loop' as a principles-based approach to formalise the importance of human oversight in policy decision-making.

Secondly, many interviewees saw the potential of GenAI to assist in completing activities that are dull, routine, and repetitive. In the process of doing so, they saw the potential of these tools to free up policy workers to focus on activities that require their higher-level skillset and expertise. This would create capacity in the workforce to address complex policy work that is dependent on human skills and expertise. Interviewees therefore saw a role in diverting low complexity and high-volume tasks to GenAI, creating capability to better meet the needs of underserved groups (e.g. creating more time for external facing services to engage with clients) and enable groups to engage with services in new ways that meet their preferences (e.g. those with verbal impairments that prefer to engage with chatbots and virtual assistants). Interviewees describe how their motivation to use GenAI tools is driven by a focus on improving outcomes for clients and end users of public services, rather than an efficiency measure or means to reduce headcount. As Participant 24 explains:

> *"The Copilot trial so far has been less about efficiency and more about delivering better customer experience... When, if, if in the process of doing that, we find that people are doing tasks quicker, that'd be great. We'd love to see a productivity gain... We've been really clear in the framing to our staff that this is not about an exercise of reducing staff. It's not about replacing staffing with machines."*

Yet some interviewees saw potential ramifications for staff in eliminating low value work, particularly the potential for greater threats to workforce psychological health and safety through sustained work complexity over time. As Participant 19 describes:

> *"I think there's probably some legitimate concerns in terms of like the type of work people would be expected to do, and you know, probably one of the, if we were successful in being able to remove some of the manual kind of drudge work out of the system, it potentially will create jobs that are at a different level and maybe much more psychologically demanding and things like that."*

As this quote illustrates, the use of GenAI tools is not a simple process of adoption; rather, getting the most out of these tools would mean significantly redesigning work processes and, as we come to below, also introducing new recruitment, training, and development processes.

Public servants prefer GenAI being kept away from external-facing activities

Interviewees expressed varying degrees of comfort in the application of GenAI tools to conduct policy work. There are certain activities perceived as potentially compromising citizen trust and confidence in governments. Senior public servants explain they are more comfortable in using GenAI to undertake internal functions that are predominately administrative in nature, rather than those that use client or end-user data or engages with these groups. Several interviewees explained it is preferable to initially focus on using GenAI with low-risk functions and datasets given the risks associated with its use.

Hence, agencies and branches seen as more progressive in their adoption of GenAI are those that are using it for activities that are predominantly administrative in nature or that are retaining low risk data sets. Central agencies responsible for internally facing functions were more likely to align with this view.

For agencies responsible for client-facing public service functions or for retaining sensitive client data, interviewees expressed greater concern with GenAI adoption. While there were some exceptions to this pattern, participants from agencies with client-facing public service functions are less satisfied that the risks of GenAI can be easily mitigated. Line agencies with service delivery responsibilities report there are features of their portfolios, linked to the nature of the clients served and the sensitivity of the data retained, prompting them to feel that some agencies were less comfortable with the use of GenAI.

Availability of GenAI tools can lead to unauthorised forms of use

Interviewees reported that some agencies have facilitated use of GenAI tools by granting access to various applications and providing training for their work-related use on official devices. Some agencies have also prohibited use of particular GenAI tools by blocking access to them from work devices. However, due to the availability of some GenAI tools, often at no cost, some agencies have seen not explicitly authorised use or even prohibited use of these for work purpose from personal devices. Interviewees explained the use of GenAI tools was becoming more prevalent because individuals saw them as a way to improve the effectiveness and efficiency of their work.

Participant 21 describes their experience of using not explicitly authorised GenAI tools: “Sometimes it’s not useful, sometimes it is so. It’s reasonably useful enough that I will do it. But I’ll do that on my personal iPad with a ChatGPT subscription that I got. Paid for myself.” As this quote illustrates, where individuals see GenAI tools as having utility for their work they may access these through personal devices on the basis that the benefits outweigh any particular risks.

For many individuals, their first experiences with GenAI are gained outside of work, meaning they are developing skills in the use of tools in a relatively unconstrained environment. As such, the notion of formal trials and explicitly authorised tools can be perceived by staff as being a constraint to the use of their existing skills. Agencies are therefore faced with a challenging context where, even though they have prohibited some tools, some staff are using them. Some interviewees did not think that preventing this required additional regulation, and, if detected, it could be addressed through existing code of conduct regulations.

Examples of current use of GenAI tools in policy work

Interviewees described a diverse suite of GenAI tools that are being used to support the design, implementation, and evaluation of policy. Senior public servants see the relevance of GenAI to different workforce types, including those at the executive level. Microsoft

products dominate in current uses, with many interviewees describing the use of Copilot and the Azure OpenAI suite. Most interviewees described that any current use of individual GenAI tools was being evaluated in some way to ensure they deliver what is expected and comply with any organisational requirement. In some cases, GenAI tools had lost their ability to be used in an authorised way as a result of these evaluations.

A summary of some current applications of GenAI tools described by interviewees include:

- Administrative applications including minute-taking, file management, and content search and retrieval using Microsoft Copilot.
- New content and idea generation using GenAI tools such as ChatGPT typically like a search engine, to more tailored policy-specific GenAI tools such as those developed by Dragonfly Thinking for analytical policy and problem-solving tasks. Individuals described using Copilot to develop a 'first draft' of a policy brief or correspondence, or Copilot or ChatGPT to establish a preliminary sense of policy issue considerations, risks, and solutions.
- Building an agency-specific knowledge base using bespoke Azure OpenAI platforms trained on agency data, policies, legislation, procedures, and other administrative data. These platforms are used in the form of virtual assistants that serve as an internal resource for functions such as human resources or call centre agents. Interviewees describe GenAI tools as being particularly useful in assessing coherence of policies and legislation, identifying potential inconsistencies over time. These platforms offer a means to structure content such as frequently asked questions and websites.
- Summarising large volumes of information or documentation. Some interviewees described using GenAI to summarise and retrieve information, for example, extracting historical ministerial and political party perspectives on issues to inform departmental approaches. GenAI has also been used alongside manual methods to summarise high volumes of submissions that are received as part of public consultation processes, with quality checks in place to assess similarity of summarised outputs.
- Service delivery applications include the use of Grammarly, Lyrebird AI, Heidi AI, or ChatGPT by practitioners to prepare client documentation such as client records or notes. Interviewees describe practices of health professionals recording clinical consultations to generate medical record entries, clinical correspondence, and discharge summaries.
- Quantitative and computer programming functions, such as integrating GenAI coding within econometric models as an extension of existing AI analytical approaches.

Interviewees described how many of these applications have been pursued through small-scale projects and pilots of explicitly authorised applications within the context of feedback cycles and evaluation. There are trials of Microsoft Copilot underway across many jurisdictions and at the federal level, with similarities in the use cases applied with

a focus on administrative functions, administrative data, and access by selected users. Activities enabled by Microsoft Copilot include content summarisation, search, and editorial review. However, while these trials offered participating workforces an explicitly authorised GenAI tool, the functionality offered by these applications often seemed limited in comparison with publicly available GenAI tools being used by staff in their personal lives. Participant 24 explains, "I think it's a really big challenge for the APS… Yes, we're giving people applications like Copilot. But we're not giving them ChatGPT. We're not giving them Claude." This quotation suggests the reliance on a narrow set of GenAI tools that were considered 'safer' within the context of obligations, yet didn't offer the same functionality as other publicly available tools to an extent that incentivised adoption.

The challenge of achieving transformational change via incremental projects

As described above, several interviewees were optimistic about the potential for GenAI to generate transformational change in the conduct of policy work. They conveyed a vision for the future of GenAI as amenable to 'blue sky' thinking that fosters creativity and innovation in the generation of ideas without constraints. Yet participants described approaches to technology implementation within the public sector through defined use cases, business cases, and trial projects which were seen as working against the conditions necessary for blue sky thinking. Interviewees explained these approaches to implementation as important mechanisms to be judicious with technology decision-making, investment, and risk management. However, they also describe a tension between the use of these methods as incremental projects and GenAI being able to realise its transformative potential. As Participant 18 explains the potential of GenAI in creating new understandings of the notion of policy specialisation and possibilities for the practice of policy work.

> *"What I'm personally interested in is let's explore the new conceptual spaces where we haven't even used tech before, that it might be able to create new possibilities…*
>
> *[When] you're in a business as usual environment, you want to do things in a sort of low risk, incremental way. And so that leads you to let's marginally improve the use cases that we know…*
>
> *But that approach means that you miss 'How can we do this thing in a way that we've never done things before? How can we apply computing power, not just to structured data, but how do we apply it to things that in the past we've thought of as [policy] craft or… an art that we do after 20 years of experience and judgement?""*

As with Participant 18, some senior public servants believe it is timely to reconceptualise policy work in ways that could challenge current understandings. While significant numbers of incremental gains could eventually realise transformational changes to the conduct of policy work, some participants did not think this was the most efficient or effective way to realise this form of change and nor is it guaranteed.

2.2 Risks and Concerns Surrounding Generative AI

Interviewees described a range of risks in the context of using GenAI tools in policy work. They include fundamental concerns that the outputs of GenAI tools, derived from their algorithms and training environments may mislead users through erroneous and biased outputs, a consequence that threatens the legitimacy of policy work. There are clear expectations that Australia should retain control over its data and risks to data sovereignty must be mitigated. Consideration of risks need to balance financial investment against the return to be realised. There are concerns that GenAI's environmental impact is at odds with government's broader environmental aims.

Data accuracy and risk of bias within the context of obligations

Senior public servants described the importance of the integrity of data used to inform policy work. Interviewees explained publicly available GenAI tools such as ChatGPT are trained on publicly accessible datasets, such as internet websites, and are therefore prone to generating inaccurate responses that may be outdated, inaccurate, or biased. These GenAI tools are reliant on data collected from a vast range of sources, which may include user uploaded content that when used in the conduct of policy work may contravene privacy and data protection obligations.

Senior public servants recognised the criticality of potential breaches of privacy and other obligations arising from fundamental features of GenAI tool algorithms, which are adaptive and prone to change. However, the details and nuances of these changes are not always transparent or understood by users. Participant 30 highlights that the algorithms used in publicly available GenAI tools are proprietary to vendors, often presenting 'a black box' where the details of internal workings are not known by governments in ways that are counter to expectations for transparency and accountability. Interviewees explained that GenAI tools can generate a response stated as fact despite the content being misleading or incorrect when assessed. Referred to as hallucinations, interviewees recognised that GenAI tools, particularly publicly available tools, are programmed to deliver a response irrespective of whether the model has access to the requisite data. This poses a clear risk in policy work when an inaccurate response is delivered and goes undetected by the user.

Interviewees explained the unique influence of the anthropomorphic features or human-like qualities retained by GenAI tools that risk confusing users. The use of GenAI tools necessitates that staff examine the integrity of responses, which requires familiarity with the content. A need for diligent application of adequate human oversight and specialised knowledge exists, alongside a critical lens to the assessment of GenAI outputs. As Participant 27 cautions:

> *"I think to anthropomorphise AI, I think it's actually really dangerous... The risk that… ChatGPT is kind of like turned into a buddy that you can have a conversation with. It sounds like a human, so maybe it behaves and thinks like a human. It's not a human, it's a machine."*

This quotation emphasises the risk of workforces establishing a connection with a GenAI tool akin to a work colleague. It may enhance trust in the tools such that users become unable to detect erroneous content.

Many interviewees reported a preference for the development of explicitly authorised GenAI tools hosted on an internal and secure government platform. These platforms enable agency oversight of the data training environment and underlying algorithms. These platforms trained on government data reflect the influence of historical policies and decisions which must be considered in modelling. However, training GenAI tools with historical government data also risks misrepresenting the nature of policy issues. Interviewees explain the potential for GenAI models to inadvertently reinforcing institutional policies and practices that have resulted in inequities for priority populations, such as First Nations, people with disabilities, and women. It is acknowledged humans are also biased and error prone. Past decision-making and policy processes have been susceptible to the influence of bias and errors made by the public service. Accordingly, capabilities to identify and eliminate the perpetuation of biases in all forms are required.

Threats to government's ability to retain sovereignty and control over its data

Interviewees described the data risks associated with the adoption of GenAI and the threats to Australian data sovereignty, including First Nations data sovereignty. There are risks that vendors gain control and influence of technology developments in ways that compromise Australian government entities' abilities to retain sovereignty or control their own data, including expectations for First Nations data sovereignty, where First Nations people retain agency regarding the collection, ownership, and application of their data. Interviewees explained that some GenAI models require overseas data storage and processing, which breaches Australian data sovereignty obligations. Many such tools have therefore not been explicitly authorised for use. Yet interviewees acknowledge that compared with other countries, Australia's federation and population means that the breadth of government entities, some with relatively small workforces, may be constrained in their individual influence when negotiating with vendors to assure future product developments continue to satisfy obligations for explicitly authorised GenAI tools that are already in use. Participant 30 states, "It's a small handful of companies predominantly out of Silicon Valley that are deciding what's an ethical way for this model to act." This quotation highlights the dominance of international technology vendors whose interests may ultimately not be aligned with Australian governments and its citizens.

Challenges in achieving return on financial investment in GenAI

Senior public servants explained the necessity to demonstrate a return on financial investments made in GenAI. The investment into GenAI costs can be considerable, including vendor licensing costs, which are prone to price escalation and budgeting risk over time. Interviewees explained that license costs can be based on a count of users. However, as workforces across state, territory, and federal government agencies may reach tens to hundreds to thousands of workers, these license costs can be significant when considered at scale. Other data and infrastructure costs can also be significant

including those to support the establishment of government data storage centres. There are additional costs of implementation incurred to support change management activities such as work process design, workforce training, and communications. Additionally, there is recognition of the risks of unplanned or planned GenAI outages, which mean the costs of retaining manual workforce capabilities must be considered as part of business continuity arrangements.

The risks associated with the magnitude of adoption costs are amplified by the challenges of calculating return on investment. Interviewees consider that in some use case scenarios there is no guarantee that investment in deploying GenAI in policy work will guarantee its use beyond minimal and short-term interactions. Reflections on the feedback and evaluation of the Microsoft Copilot trial demonstrates there are some users who had a few simple interactions with the tool, but they did not result in increased or regular patterns of use.

Senior public servants emphasised the need for metrics to monitor the costs of implementation and the value derived. Activities might include feedback and formal evaluations to apply lagging and leading indicators of benefits. These indicators of benefits should form the basis for a whole-of-government perspective of outcomes and costs derived from the use of GenAI. Yet an interviewee also highlights the challenges of budgeting for GenAI use with vendor consumption-based pricing not always aligning with government arrangements, whereby variation between projected and actual GenAI costs risks leaving work units prone to budget overspend or underspend. Consequently, there are considerations for how agencies quantify costs over time, where vendor narratives risk being prone to a short-term perspective; however, an assessment of the costs and value from GenAI deployment necessitates a longer-term view.

There are concerns over GenAI's environmental impact given sustainability aims

When considering value, governments need to take into consideration a broad range of impacts. Governments retain stewardship responsibilities, which include aspects pertaining to the environment that must be considered alongside the development of GenAI strategies and assessment of value. Interviewees explained the opportunities and benefits that GenAI presents must be balanced against its considerable environmental impact. The risks associated with the scale and pace of GenAI adoption in policy work are linked to increased energy and water consumption in the training, and use of GenAI models. Interviewees explained that data centres have significant cooling requirements, which mean there are locations in Australia where water requirements may be challenging, especially given the dry and changing climate conditions across much of the Australian landscape. There are risks that a growth in GenAI adoption across government functions may shift water and energy needs in ways that will elevate its environmental impact in locations where data centres are hosted. As Participant 33 explains, GenAI's environmental impact is a threat to the realisation of its environmental sustainability and stewardship goals.

> *"We know that these models, specifically GenAI, has an absolutely monstrous carbon footprint, and the water consumption is also massive. So aside from the lens of, you know, environmental ethics, where, you know, we should consider ourselves stewards of our resources."*

This quotation emphasises the need for alignment between the government's GenAI strategies and its environmental goals.

2.3 Enablers for the Adoption of Generative AI in Policy Work

This section describes some of the enablers for GenAI adoption in policy work. Interviewees suggested the need for strategies that outline the intended role and purpose of GenAI. Clarity regarding strategy is key to building understanding and confidence in the approach across the workforce. Senior executive capability to oversee the adoption and use of GenAI is considered a critical enabler, alongside enhancement of workforce training, development, and supervision models. The notion of GenAI augmented policy craft must be well-understood as a basis for building broader workforce capability. Alongside these considerations, alignment with regulatory, governance, and data obligations remain important to assure the quality of government data as an asset.

Strategies are needed to build clarity, trust, and confidence in GenAI

Senior public servants described the necessity of clarity regarding the role and purpose of GenAI within government. They perceived shared understanding as being essential to the development of a narrative for citizen engagement to build trust and confidence. Interviewees explained that a focus in their agencies on incremental projects and trials that are subject to a structured approach to feedback, review, and formal evaluation seeks to build transparency about the risks and costs of GenAI adoption. Yet these efforts occurred in accordance with a 'wait and see' approach to new technology adoption, lacking a strategic commitment to transformative change and investment in GenAI.

Interviewees described that governments are prone to external pressures for them to be market leaders in GenAI adoption. As Participant 20 explains, these influences are leading to GenAI adoption becoming a symbol of public sector modernisation and progress.

> *"It seems that if you don't have AI use even for the sake of it, you're seen as someone who's not future proofed. Archaic. Not with the times. Potentially wasting money. Missing opportunities."*

Interviewees expressed a belief that government GenAI strategies should be grounded in aims to improve the quality of policy work, including responsiveness to citizen and client expectations. Accordingly, the development of GenAI strategies necessitates a reassessment of government's moral, ethical, and democratic responsibilities. Within the context of the potential of GenAI, many felt it is necessary to reconsider the unique human contribution of public servants in executing these responsibilities. Interviewees emphasised the necessity of ensuring 'a human in the loop', and, as Participant 11

explains, “If you’re making a decision that affects people, you need to have a level of accountability and transparency.”

Interviewees framed the effectiveness of government strategies for GenAI adoption as central to building citizen trust and confidence in its application to policy work. In the absence of a shared understanding of, and commitment to, GenAI strategies, adoption efforts risk being vendor or industry-led in ways that might work against public interest. Participant 3 highlighted the necessity to be transparent about the risks of maintaining current arrangements and a reliance on human-led practices alongside the opportunities that GenAI presents:

> *“And we always talk about the risks of the new, but not about the default... At least you can generally sort of mitigate some of the risks of the machine... It’s much harder to do so with us as humans…*
>
> *There are going to be things that we do as in the public sector where you want a machine to help you because it’s hard and complex and humans aren’t going to do a very good job.”*

The quotation explains the need to be clear about the rationale for changes to arrangements that aim to foster improvements in the quality of policy work.

Senior executive capability is needed to deliver GenAI oversight

Senior public servants considered contemporary GenAI knowledge, skills, and experience as necessary capabilities to support the adoption of GenAI. Yet interviewees reported there are some areas within the executive that have minimal capability to do so. The generalist nature of some senior public servant roles does not necessarily encompass digital technologies as part of the core skillset. The scale and pace of GenAI development likely mean that senior public servants will need to improve their critical AI capabilities to oversee its application. Interviewees explained that limitations in executive capability to do this can manifest as enculturating resistance to GenAI within branches or work units. Accordingly, all senior public servants were seen to hold a responsibility to build GenAI knowledge, skills, and experience to manage its risks and realise its potential.

Interviewees explained that the Australian technology workforce is relatively small, which means that governments must compete with vendors and industry for skilled talent. This also poses challenges for public sector recruitment and retention of GenAI expertise. Senior public servants described that technology portfolios and responsibilities have been assigned to those with a generalist skillset given difficulties in hiring and retraining, but with a view that these functions can be delivered at a senior level by a generalist skillset. Participant 11, however, reflects that chief finance and legal roles require a high level of qualification and experience, but in some instances chief technology roles may not. Participant 19 shares the perspective of becoming responsible for the GenAI portfolio without a specialist skillset, acknowledging the importance of decisions made.

> *"You get people like me who find themselves all of a sudden responsible for AI. And I'm probably… equipped as anyone, but still like we're not AI specialists. And as you go up the line, people are going to be in those jobs not because they're AI specialists, but they're going to be expected to make some pretty gnarly decisions about how we're using AI."*

This quotation explains that improvements in GenAI capabilities at senior levels of government are critical to mitigating the risks and realising the benefits to policy work.

Senior public servants engage with new technologies in diverse ways. They are often not able to participate as members of communities of practice, creating a reliance on formal training or being self-taught. There are also unique features of time-pressured senior public sector roles meaning work pressures are not conducive to in-depth or sustained use of GenAI, with higher proportions of time spent with non-desktop-based activities such as meetings. Some interviewees concede their experience of limited regular use of GenAI despite oversight for an agency's technology portfolio. Interviewees spoke of relying on 'public sector gossip' to garner an understanding of GenAI developments, including the challenges experienced by peers. Executive coaching and non-traditional strategies such as reverse mentoring involving the delivery of training supports by younger workforces are profiled as another means to build understanding and readiness to GenAI across mature aged workforces.

Effective use of GenAI needs workforce training, development, and support.

The advent of GenAI tools to augment human capabilities has prompted a reassessment of the fundamentals of policy craft. Senior public servants recognise the quality of policy work is a product of the expertise and experience of its workforce. Therefore, the task is to maintain or enhance the quality of public service provision where features of expertise and experience are vested into GenAI tools. Many interviewees anticipated shifting reliance from humans as a source of institutional memory and analytical capability over time. Irrespective of this shift, many indicated that an understanding of the technical features of GenAI tools as a contemporary policy capability that supports effective validation of outputs for use in policy work. For instance, Participant 33 explained the risk that individuals defer to the outputs of GenAI tools on the basis that "maybe it knows something that I don't" - in essence, letting the tool override their expertise. Hence, the practice of expertise-informed policy work is expected continue to evolve alongside the adoption of GenAI, creating a need for building workforce capability in new forms of GenAI-augmented policy craft.

There are several workforce training and development needs to be addressed to enable the adoption of GenAI in policy work. To address the implications of attributing human characteristics of expertise and skill to GenAI tools, which may inadvertently lead to negating the need for human oversight, the workforce must be supported to build a basic understanding of GenAI's underlying algorithmic approach. Workforce training and development should encompass expectations for validation processes, document requirements, and citation methods. Building workforce capability in these expectations

is also considered essential to meeting documentation requirements for potential increased Freedom of Information requests using GenAI tools. Aside from these practical skills, interviewees emphasised the relevance of workforce capability in responsible and ethical use of GenAI, including the risks of longstanding issues of error and bias being amplified by GenAI.

There is also a need to retain a level of skill and capability in functions that are undertaken by GenAI tools such that these could be manually conducted in the event of GenAI outages. Senior public servants explained the introduction of GenAI may expedite a shift from team-based to individualised ways of working. It is anticipated that individuals may pair with GenAI tools akin to a team member or work colleague. The benefits of this paired arrangement with GenAI tools may strengthen individual policy craft capabilities; however, additional strategies are required to retain the benefits of team-based work with other people. These include the merits of collective deliberation, brainstorming, and development of a shared understanding to advance policy work. Interviewees describe the necessity for new ideas to address complex problems, mindful that GenAI tools fundamentally cannot create any new ideas as algorithms draw from patterns in historical data to generate words and content.

New methods of working may require revisiting accepted forms of supervision and support. This may be relevant to more established members of the workforce but is particularly relevant to junior workforces keen to advance their careers in policy work. Yet alongside the adoption of GenAI, responsibilities such as minute-taking or preparation of first drafts of briefs are increasingly being conducted by more experienced employees with the assistance of GenAI. These functions, often undertaken by junior staff, are opportunities for observation and skills development. There exists a need for workforce planning efforts to consider training pathways and career development to assure the pipeline of properly prepared future senior public servants.

There needs to be alignment across regulatory and governance arrangements

Many interviewees perceived regulatory and governance frameworks that underpin GenAI implementation as multi-layered, involving agency efforts across state, territory, and federal governments. The adoption of GenAI in government occurs in the context of obligations, including, but not limited to, privacy, records management, freedom of information, public sector code of conduct, the human right to privacy and reputation, and other government solicitor advice. Compliance with current regulatory and governance frameworks, whilst retaining readiness for future developments, are seen as potentially an enabler for GenAI adoption. However, current frameworks are considered prone to ambiguity and fragmentation. As participant 10 explained, “We’re not just dealing with one other government agency. We’re dealing with multiple agencies that have their finger on the pulse and are doing this. You know, [the Digital Transformation Agency] develops the policy, but [the Department of Industry, Science and Resources] are developing something else… We just need to make sure that we’re across all of that. So, we are doing the right thing.” It is important to ensure that the evolving regulatory and governance arrangements across levels remain aligned across the breadth of agencies involved in these functions.

3. What Can We Take from This Research?

This was exploratory research that sought to examine GenAI as an emerging technology. Perhaps unsurprisingly, we did not find concrete agreement around how a range of senior public servants see the future of GenAI use in policy work. Instead, we found interviewees discussing a series of dilemmas they had encountered and were dealing with in their everyday practice.

Some of these dilemmas related specifically to the technologies of GenAI, such as how to use these technologies safely and in a cost-effective way in a context where vendors of these technologies have significant power within the market. Or, how can policymakers deal with issues of GenAI tools drawing on data sets that might contain significant biases?

Other dilemmas relate more specifically to the practice of policy work and the implications of using new tools. This includes issues like: if GenAI tools become more widespread in applications in policy work, then what does this mean for how we train and develop policy workers? If GenAI tools take over aspects of activities that might be routine and repetitive, such as minute taking in meetings or compiling research about a policy issue, how do we provide alternative pathways for workforce development? Or, if the use of GenAI tools is to be transformational, are we able to achieve such change through multiple iterations of relatively constrained pilot projects?

These dilemmas illustrate an important over-arching challenge: that utilising technologies to their full extent is about more than a simple process of adoption. Adopted technologies rarely deal with a simple and straightforward problem. As such, they require careful consideration in terms of how they integrate with organisational systems and how they impact on professionals and users of public services. Indeed, this has been one of the mistakes made in a number of attempts to introduce large-scale IT projects (7). If GenAI technologies are to have a substantial impact on the way that policy work is done, we need to not just consider the strengths and limitations of these tools but also address the ways that these features might fundamentally change how policy work is conceptualised, developed, and supported within organisations.

3.1 Implications for Policy and Practice

Findings from this exploratory study have several implications for policy and practice. These include, but are not limited to, the following:

- Ensure that strategies for GenAI adoption are aligned with governance arrangements as enablers for the realisation of aims. Strategies should be consistent with public service values, including stewardship, impartiality, commitment to service, accountability, respectfulness, and ethical practice. This includes government's environmental sustainability interests.

- Develop plans and resources to build confidence in the strategy for GenAI adoption, including reports of measures of value and methods for iterative review. This should include a focus on the legitimacy of decision-making processes, underpinned by human oversight, and the formality with which GenAI contributions to policy work are recorded to meet expectations for accountability (e.g. Freedom of Information requirements).
- Re-examine workforce training, development, and support across all levels in GenAI augmented policy craft, including senior executives.

3.2 Areas for Future Research

As outlined in the introduction, to date limited research has been undertaken into the use of GenAI in policy work. As such there are a range of gaps within our existing knowledge base. This is an evolving area that research could make a significant contribution to supporting practice, especially as these technologies are more substantively engaged with across governments. This has implications for not only what is researched but how this research is undertaken.

In terms of what is researched, there are clear gaps in terms of topics, such as the value of GenAI in policy work, how GenAI disrupts and shapes policy craft, how organisations and regulatory cultures and practices shape the use of GenAI tools, how public organisations can engage with vendors of GenAI technologies to ensure safety and meet broader stewardship commitments (e.g. environmental goals and data sovereignty obligations), and how public service organisations can build and maintain community trust in activities undertaken using GenAI tools.

Longitudinal research with different areas and types of policy work may be helpful to capture both the evolution of GenAI tools over time and the full range of implications they have for policy work at different points in their lifecycle. Technology adoption is not one occurrence in time; professions and organisations need time to respond to new tools and fully embed them within their organisational practices. Given resource constraints, all too often research projects – such as this one – provide a snapshot in time rather than a more complete view on the process of implementation.

Embedding researchers within teams as they seek to implement GenAI tools may also offer a helpful way for teams to reflexively consider how these technologies are being rolled out in organisations and the impacts that they have for broader systems. It is also important that multi-disciplinary research teams are assembled for this process bringing together insights from disciplines such as sociology, public administration, psychology, economics, law and organisational theory to ensure that the full extent of the implementation of these tools are captured.

3.3 Limitations

We acknowledge some limitations with this research. While all participants met the eligibility criteria for this study, it was identified that some senior public servants retained varied depth in relation to both technological knowledge and programmatic adoption of GenAI. However, this limitation is not unexpected given the nature of GenAI adoption across and within public sector agencies, varied approaches to knowledge management, and prevalence of use of GenAI tools.

While we retained good spread of representation across the country from various state, territory, and federal governments, this report is a snapshot of considerations related to GenAI tools at a particular point in time. As GenAI tools continue to develop and evolve the types of applications, issues, and debates will also continue to change.

4. Summary

This research identifies a high degree of divergence in understandings and expectations of GenAI in policy work. This is not unsurprising given the diversity of agency engagement with an emerging technology. Senior public servants who form the executive structures across state, territory, and federal government agencies have varied perspectives of the goals and outcomes that the introduction of GenAI could achieve. The emphasis on GenAI-related risks and the absence of explicit objectives about the improvements that GenAI is expected to realise creates conditions that may not enable the realisation of AI-related benefits. As such it is important to emphasise the importance of attending to enablers for GenAI adoption.

A recent report highlighted that “technology affects the evolution of jobs” (1, p. 41). Accordingly, agencies are facing fundamental questions as to the features of its roles that rely on qualities, judgements, and decisions that are distinctly human. The task ahead is to be proactive in building trust in GenAI and confidence in its use to maintain high-quality policy craft. There remain considerations for realising value with attention to the costs and benefits of GenAI adoption. The strength of the GenAI strategy and the implementation of plans is essential to ensuring policy work remains primed to improve to meet the needs of Australian citizens.

5. References

1. House of Representatives Standing Committee on Employment EaT. *The future of work: inquiry into the digital transformation of workplaces*. In: Australia Po, editor. Canberra Commonwealth of Australia; 2025.

2. Feuerriegel S, Hartmann J, Janiesch C, Zschech *P. Generative AI. Business & Information Systems Engineering*. 2024;66(1):111-26.

3. Al Naqbi H, Bahroun Z, Ahmed V. *Enhancing Work Productivity through Generative Artificial Intelligence: A Comprehensive Literature Review*. Sustainability. 2024;16(3):1166.

4. Commonwealth of Australia Digital Transformation Agency. *Policy for responsible use of AI in government In: Commonwealth of Australia Digital Transformation Agency, editor*. Commonwealth of Australia Digital Transformation Agency,: Canberra 2024.

5. Lorenz P, Perset K, Berryhill J. *Initial policy considerations for generative artificial intelligence*. OECD Publishing; 2023.

6. Pfohl SR, Cole-Lewis H, Sayres R, Neal D, Asiedu M, Dieng A, et al. *A toolbox for surfacing health equity harms and biases in large language models*. Nature Medicine. 2024;30(12):3590-600.

FOR MORE INFORMATION

psrg@unsw.edu.au

unsw.adfa.edu.au/public-service-research-group/

@PSResearchG

UNSW CANBERRA | Public Service Research Group

Australian National University