



Australian Government




RBA Policies + AI

Information Retrieval Tool for RBA Corporate Policies


Reserve Bank of Australia





Executive Summary

Problem Faced	Solution Overview	Benefits and Impact
RBA staff spend a significant amount of time going through lengthy (often difficult to read) documents to answer questions related to the Bank’s corporate policies, procedures and guidelines (hereafter referred to as “Corporate Policies”). This process can be inefficient, particularly when staff are seeking quick clarification or need to locate specific information across multiple documents. Traditional document search methods are not well suited to this task. They are limited in their ability to interpret natural language queries or return contextually relevant results, which can slow down workflows and reduce the ease with which staff engage with policy content.	“Policies + AI” is a secure, in-house AI-powered search tool that helps staff efficiently retrieve information from the RBA’s corporate policies, procedures and guidelines. It uses Retrieval-Augmented Generation (RAG) to enhance the output of a Large Language Model (LLM) by referencing an authoritative knowledge base before generating a response. Users submit a query and receive: • Links to relevant policy and procedure documents • The top three relevant paragraphs • An LLM-generated answer summary The tool supports the RBA’s strategic objective to work <i>Smarter, Simpler and Faster</i> by making policy content easier to access and interpret.	The tool improved how staff engage with policy content, delivering measurable efficiency gains and broader organisational benefits: • Faster access to relevant, policy-aligned information, with an estimated time savings of 7½ minutes per user per week, equivalent to approximately 5¼ full-time staff annually. • Improved consistency in interpreting and applying policy • Strong early adoption, with around 25 unique logins per day following launch These outcomes support more responsive internal services and contribute to a more efficient and informed workplace.



Target Audience and Stakeholders

The tool was developed to support all RBA staff, regardless of business area or technical expertise.

Key stakeholders include:

- The RBA's Data Science Hub and Risk Management Department, who led development and oversight.
- Staff from both policy and operational areas, who participated in user testing and feedback during the proof-of-concept phase.

Stakeholder engagement was central to the tool's development, with feedback used to refine performance and ensure alignment with user needs.

Risks and Mitigation Overview

A formal risk assessment conducted following the Proof of Concept identified key risks associated with the tool. These included:


- **Performance quality:** Managed through human-in-the-loop testing, prompt refinement, and grounded responses.
- **Over-reliance on AI outputs:** Mitigated by interface warnings, mandatory user acknowledgement, and clear guidance on appropriate use.
- **Compliance and legal risk:** Addressed by restricting the tool's knowledge base to approved internal policy documents.
- **Privacy and data protection:** Managed through secure internal deployment and by ensuring the tool did not process personal or sensitive data.

Use Case Status

Implemented

Use case timeline

- **Proof of concept (4 months):** Preparation, testing, optimisation, user feedback, and decision phase
- **Bank wide rollout (12 months):** Scaling infrastructure, automating data ingestion, preparing risk assessment and sign-off, and delivering a Bank wide communications campaign.



Additional Information

The tool was designed to integrate with the Bank's internal systems and workflows. It includes automated ingestion of policy documents to ensure content remains current and aligned with official sources. The Retrieval-Augmented Generation (RAG) architecture allows the tool to deliver grounded responses while maintaining flexibility for future enhancements. The solution was developed entirely in-house, allowing for greater control over data handling, deployment, and risk management.

Lessons Learned

- **Start with real user questions:**
Frequently asked questions helped focus development on practical needs and supported meaningful evaluation.
- **Expert review adds value:**
Subject matter expert review of AI outputs against known answers improved quality and trust in the tool.
- **Keep expectations clear:**
Interface warnings, user guidelines, and internal communications helped reinforce that the tool is a support resource, not a source of official advice.

Contact information

Responsible Entity Name

Reserve Bank of Australia

Open for Collaboration?

Yes

Area of Entity

Data Science Hub
Economic Research Department

Use Case Contact


s 22
s 22 [@rba.gov.au](mailto:s 22@rba.gov.au)

Use Case Website/s

N/A

Use Case Owner

s 22
s 22 [@rba.gov.au](mailto:s 22@rba.gov.au)



Screenshot/s

Policies + AI

This is an experimental tool to be used as a resource to help you extract information from the Bank's policies, procedures and guidelines as listed on [Policies+](#)

Open for  Tips: 


Show Terms of Use

☒ I have read and I agree to the Terms of Use.

Question

Ask your policy question here 

Search Policies +



Detailed Overview

Version Control

Version	Date	Author	Description of Changes
1.0	3 Feb 2025	GovAI	Version 1 created
1.1	17 Mar 2025	GovAI	Modified based on feedback

Index

Responsible Organisation Category.....	5
Scope of the Use Case	5
Ethical Considerations	6
Value of the Use Case.....	7
AI Process Type.....	8
AI Technologies Utilised	9
Technical Elements.....	10


Note: For details about category items in the detailed overview, see *APS AI Use Case Repository Guidance- Guidance for Use Case Owners and Editors*.

Responsible Organisation Category

Select the Classification of the Functions of Government - Australia (COFOG-A) 3-digit category that best identifies the functional area associated with your AI use case.

☒ 01 - General Public Services	014 - Planning and statistical services
☐ 02 - Defence	Choose an item.
☐ 03 - Public Order and Safety	Choose an item.
☐ 04 - Economic Affairs	Choose an item.
☐ 05 - Environmental Protection	Choose an item.
☐ 06 - Housing and Community Amenities	Choose an item.
☐ 07 - Health	Choose an item.
☐ 08 - Recreation, Culture, and Religion	Choose an item.
☐ 09 - Education	Choose an item.
☐ 10 - Social Protection	Choose an item.
☐ 11 - Transport	Choose an item.

Scope of the Use Case




Use the dropdown menus below to identify the scope of your use case.

Geographical focus Choose the region for implementation from the dropdown list	National
Primary type of government interaction Choose the type of government interaction from the dropdown list	Government-to-government (G2G)
Cross-features - Sector Indicate if the use case describes a solution that can be used across sectors or in cross-sector scenarios (Yes/No).	Yes
Cross-features - Jurisdiction Indicate if the use case describes a solution that can be used across State/Federal borders or in cross-border scenarios (Yes/No)	Yes

Ethical Considerations

Accuracy, Fairness, Accessibility, Bias and Discrimination	• Accuracy: The tool's responses are grounded in the Bank's official policy documents. Outputs were validated against a structured set of FAQs and "golden" answers during the proof of concept. • Fairness: The tool is available to all staff and does not personalise responses. It treats all queries equally, regardless of user role or background. • Accessibility: A natural language interface allows users to ask questions in plain English, making policy content easier to access and interpret. • Bias mitigation: The tool does not use personal data or external sources. Human reviewers assessed outputs to ensure alignment with policy intent and to identify any unintended bias. • Transparency: Users are shown the source paragraphs and links alongside
---	---




	the AI-generated response, supporting interpretability and trust.
Privacy	• The tool does not use or process personal or sensitive data. • All responses are generated from a curated internal knowledge base of corporate policies and procedures. • The system is deployed securely within the Bank's secure cloud-based environment, with access restricted to staff. • No identifiable information is stored or processed, and the tool was developed in line with privacy-by-design principles.
Rights of Users	• Users are informed of the tool's purpose and limitations through: • A mandatory user acknowledgement at first use • Clear interface warnings that reinforce appropriate use • The tool does not make decisions or provide official advice. • Users are encouraged to verify responses against source documents. • Feedback was collected during the proof-of-concept and user testing phases, and ongoing engagement supports continuous improvement.

Value of the Use Case

Identify the public value that the solution provides or is expected to provide. Select from the multi-select options.

Improved public service This category refers to solutions that enhance the services provided to end users, whether they are citizens or businesses.	☐ Personalised services ☐ Public (citizen)-centred services ☒ Increased quality of public information and services ☒ More responsive, efficient and cost-effective public services ☐ New services or channels
---	--



Improved administrative efficiency This category refers to solutions that increase efficiency, effectiveness, and quality while reducing costs within administrative processes, systems, and services.	☒ Cost reduction ☐ Responsiveness of government operation ☒ Improved management of public resources ☒ Increased quality of processes and systems ☐ Better collaboration and better communication ☐ Reduced risk of corruption and abuse of the law by public servants ☐ Greater fairness, honesty and equality enabled
Open government capabilities This category refers to solutions that enhance the level of openness, transparency, engagement, and communication within public organisations.	☒ Increased transparency of public sector operations ☐ Increased public participation in government actions and policymaking ☐ Improved public control of and influence on government actions and policies

AI Process Type

Select the types of tasks within government operations that the AI solution is performing or expected to perform

Supporting Decision Making- Tasks that support formal or informal agency decision-making on benefits or rights.	☐ Taking decisions on benefits ☐ Managing copyright and intellectual property rights
Analysis, monitoring and regulatory research - Tasks that collect or analyse information that shapes agency policymaking.	☒ Information analysis processes ☐ Monitoring policy implementation ☐ Innovating public policy ☐ Prediction and planning
Enforcement - Tasks that identify or prioritise targets of agency enforcement action.	☐ Smart recognition processes ☐ Management of auditing and logging ☐ Predictive enforcement processes ☐ Supporting inspection processes ☐ Improving cybersecurity ☐ Registration and data notarisation processes ☐ Certification and validation processes




Internal management - Tasks that support agency management of resources, including employee management, procurement, and maintenance of technology systems.	☐ Internal primary processes ☒ Internal support processes ☒ Internal management processes ☐ Procurement management ☐ Financial management and support
Public services and engagement - Tasks that support the direct provision of services to the public or facilitate communication with the public for regulatory or other purposes.	☐ Engagement management ☐ Data-sharing management ☐ Governance and voting ☐ Payments and international transactions ☐ Supporting disintermediation ☐ Authentication of self-sovereign digital ID services ☐ Service integration ☐ Service personalisation ☐ Tracking of goods and assets along the supply chain

AI Technologies Utilised

Select the types of AI technologies proposed / utilised to deliver the use case.


Reasoning or Knowledge Representation AI systems that store, structure, and process knowledge to make inferences, derive conclusions, or support decision-making.	☒ Knowledge Representation ☐ Automated Reasoning ☐ Commonsense Reasoning
Planning and Optimisation AI techniques that generate, refine, and optimise action sequences or resource allocation to achieve specific goals efficiently.	☐ Planning and Scheduling ☐ Searching ☐ Optimisation
Learning and Adaptation AI systems that identify patterns, extract insights, and improve performance over time based on data.	☐ Machine Learning ☐ Deep Learning ☒ Generative AI
Communication and Natural Language Processing	☒ Natural Language Processing (NLP)




AI systems that process, interpret, and generate human language for interaction, comprehension, and automation.	☒ Text Generation ☐ Text Mining ☐ Machine Translation
Perception through the Senses AI systems that process and interpret sensory data, such as visual, auditory, or tactile inputs, to understand and respond to their environment.	☐ Computer Vision ☐ Audio Processing
Integration and Interaction with the Environment AI systems that interact with physical or digital environments, including autonomous agents, robotics, and interconnected systems.	☐ Multi-agent Systems ☐ Robotics and Automation ☐ Connected and Automated Vehicles (CAVs)
AI as a Service AI capabilities delivered through cloud-based platforms, offering tools, models, and infrastructure for AI-powered applications.	☐ AI Services (e.g., cognitive computing, machine learning frameworks, bots) ☐ Infrastructure as a Service (IaaS) ☐ Platform as a Service (PaaS) ☐ Software as a Service (SaaS)
Additional Comments or Explanation:	If you have selected any of the subcategories above, feel free to provide more detailed comments or a description of how these elements apply to your specific use case.

Technical Elements

Platform implementation	The Policies + AI tool is hosted securely within the RBA's secure cloud-based infrastructure. It was developed in-house and integrates with the Bank's internal systems to ensure alignment with existing workflows and security protocols. Key technical considerations include: • Hosting: Secure cloud-based RBA infrastructure • Security: Access restricted to staff; no personal or sensitive data is processed • Integration: Connects to a curated internal knowledge base of corporate policies and procedures.
--------------------------------	---




	• User interface: A browser-based interface allows users to submit natural language queries and receive grounded responses. The interface includes a mandatory Terms of Use acknowledgement and guidance features such as "Open for Tips". • Identity and access: Access is managed through internal authentication systems	
Model / Algorithm used	The tool uses a Retrieval-Augmented Generation (RAG) architecture. This combines: • A Large Language Model (LLM) to generate natural language responses • A curated internal knowledge base of corporate policies to ground those responses The LLM is prompted using structured templates and constrained to generate answers only from retrieved, authoritative content. Prompt design and model tuning were informed by human-in-the-loop evaluation during the proof-of-concept phase.	
Data Sources Select the types of data sources used and provide relevant details.	☒ Internal	☐ Third-party
	☐ Public	☐ Synthetic
Risk Assessment and Mitigation Details	Details: Curated policy documents and procedures from the Bank's internal repositories.	
	The risk assessment process was embedded throughout the development lifecycle, beginning with the proof-of-concept phase and continuing through to rollout. Approach: • A human-in-the-loop evaluation was used to assess the tool's accuracy. A set of 59 frequently asked questions (FAQs) and corresponding "golden" answers were compiled. Responses were evaluated for fluency, coherence, correctness and completeness.	




- Feedback from this evaluation informed prompt engineering and model tuning.
- A two-week user testing phase with 10 staff from policy and operational areas provided additional insights into usability and performance.

Key risks and mitigation strategies:

- **Accuracy and hallucination:**
 - Responses are grounded in a curated internal knowledge base of official policy documents.
 - The tool does not access external sources, reducing the risk of hallucinated content.
 - Prompt templates are version-controlled and locked to prevent unintended changes.
- **Performance drift:**
 - Ongoing monitoring and feedback mechanisms are in place to detect and address any degradation in output quality over time.
- **Over-reliance on AI-generated responses:**
 - The interface includes clear warnings that the tool is a support resource, not a source of official advice.
 - A mandatory user acknowledgement reinforces this message.
- **Compliance and legal exposure:**
 - The tool only references approved internal documents, ensuring alignment with official policy positions.
 - No external or unofficial content is included in the knowledge base.
- **Privacy and data protection:**
 - The tool does not process or store personal or sensitive data.



	- It is deployed securely within the Bank's secure cloud-based environment, with access restricted to staff. Cybersecurity: - The system is hosted on the RBA's secure cloud-based infrastructure. - Access is controlled, and usage is monitored in line with the Bank's information security protocols. User expectations and safe use: - A Bank wide communications campaign supported rollout, helping staff understand the tool's purpose, limitations and appropriate use. These measures ensured that the tool's deployment remained within the Bank's risk tolerance and aligned with responsible AI principles.	
Security and Compliance Frameworks Select the security and compliance frameworks and measures implemented. Provide details or additional artifacts if relevant.	☐ Authority to Operate (ATO) ☐ System Security Plan (SSP) ☐ Security Risk Management Plan (SRMP)	☐ Information Security Registered Assessors Program (IRAP) ☐ Penetration Testing
	Details: No formal frameworks (e.g. ATO, IRAP) were conducted. The tool is deployed securely within the Bank's secure cloud-based environment, with access restricted to staff and no processing of personal or sensitive data.	
Assurance and Government Frameworks	No formal assessment against frameworks such as the DTA AI Assurance Framework has occurred at this stage.	
Record maintenance	Documentation was developed during the proof-of-concept and rollout phases. This includes: Evaluation results from human-in-the-loop testing	



	• Prompt templates and version control • Feedback from user testing • Internal communications and usage guidance
Disengagement	No formal disengagement plan is documented. However, as the tool is hosted on the RBA's secure cloud-based infrastructure and does not process personal data, the risk profile is low. The system could be decommissioned by disabling access and archiving the knowledge base if required.
Performance Metrics and Results	Performance was evaluated through: • Time savings: Estimated at 7½ minutes per user per week, equivalent to approximately 5¼ FTEs annually • Adoption: Around 25 unique logins per day following launch • Accuracy: Validated against 59 "golden" answers during proof-of-concept testing • User feedback: Collected during a two-week trial and used to refine the tool These metrics support the tool's value in improving internal efficiency and access to policy information.