

EVIDENT RESPONSIBLE AI BANKS

Emerging best practices for Responsible AI deployment in banking

RESPONSIBLE AI
REPORT

2025/03

Disclaimer

Evident Responsible AI Report, March 2025, is published by Evident Insights.

The information in this report has been obtained from sources we believe to be reliable and accurate. However, it has not been independently verified. We make no representation or warranty, express or implied, as to the accuracy or completeness of this report. We undertake no obligation to publicly revise or update this report and shall in no event be liable for any loss or damage arising in connection with the use of the information in this report.

© Evident Insights Ltd 2025

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, or otherwise for commercial gain. You may use content from the report for non-commercial use, provided that you attribute the content to “Evident Responsible AI Report”, March 2025.

1. Introduction	1
2. Chapter 1: People	5
3. Chapter 2: Policy	10
4. Chapter 3: Process	13
5. Appendix A: Draft Maturity Model	16
6. Appendix B: Glossary of Terms	17

Introduction

Banks are traditionally risk-averse organizations. They operate in a highly regulated and competitive industry, where establishing (and maintaining) consumer trust is paramount.

As such, they maintain robust technology control, risk management, and governance frameworks – including model risk management, operational resilience programs, and regulatory compliance structures. These guardrails were built over many decades to help financial institutions adapt to the latest technological innovations and regulatory expectations.

However, AI poses new risks. While these past investments provide a starting foundation for banks to manage emerging risks specific to AI, their ability to address the dynamic, black-box, and/or data-intensive nature of AI systems has been tested. Operating in a constantly evolving regulatory environment only adds complexity (and urgency).

No doubt the advent of Agentic AI will further exacerbate these challenges. Even if parameters are well defined, agents have the ability to self-optimize and engage with external sources through APIs and even customer interactions. This makes them hard to predict, control and audit using more static governance protocols.

As banks look to expand and scale-up their Generative AI capabilities across the enterprise in the coming years, balancing robust existing governance processes with more dynamic controls that account for evolving AI complexity will be critical. Those who invest early and often in these domains will see Responsible AI become a source of competitive advantage.

In this report we'll look at where the banking sector is now in terms of Responsible AI adoption. We'll explore how they're actively adapting their people, policy, and processes; what lessons they've learned; and what best practices they've adopted to accelerate innovation while ensuring they deploy AI responsibly.

In the end, Responsible AI is not a barrier to faster innovation – it is a key enabler.

ACKNOWLEDGEMENTS

Our special thanks to several members of the Evident community who informed this report, including:

Sebastien Devoivre (Société Générale), **Dr. Paul Dongha** (NatWest), **David Heike** (Strategic Consultant), **Dan Jermyn** (CommBank), **Kirsten Mycroft** (BNY), **Raquel Ettrick Thompson** (BNY), **Luke Vilain** (UBS), and **Ozge Yeloglu** (CIBC).

Alexandra Mousavizadeh

Co-CEO & Co-founder

alexandra.mousavizadeh@evidentinsights.com

Annabel Ayles

Co-CEO & Co-founder

annabel.ayles@evidentinsights.com

WHY IS RISK FUNDAMENTALLY DIFFERENT IN THE CONTEXT OF AI?

This preamble begs a simple question: *If risk management is constantly evolving alongside technology to identify, evaluate, and address potential risks to the organization, what sets AI apart?*

The panel of subject matters experts consulted for this report identified four specific points of differentiation.

First, Generative AI has a “black box” problem. At the end of the day, these models produce outputs in a way that makes it difficult to pinpoint the exact reasoning that led to a given result. While banks are pouring resources into Explainable AI (XAI), many feel like we’re only scratching the surface – especially in an operating environment that demands strict controls and auditability.



Dr. Paul Dongha
*Head of Responsible AI
and AI Strategy*
NatWest

‘Predictive AI poses challenges for model validation teams, who pride themselves on making sure their models are resilient, transparent, and explainable. Controls required to ensure these concepts are upheld are now commonplace in enterprise risk management and model development platforms. The onset of Generative AI has posed heightened and new risks which require a greater investment in governance, risk, and compliance activities.’

Second, the risks associated with AI are far broader than just model risk. They extend to individual systems, algorithms, and organization practices. Based on emerging standards (see ISO/IEC JTC 1/SC 42), stakeholders can already anticipate comprehensive audits that require access to training data, development and deployment processes, and systems engineering documentation.



Luke Vilain
AI Governance Lead
UBS

“There’s no reason why the people who have come from model risk backgrounds cannot be cross-trained, but the pure fact of the matter is that the risks to do with AI are much broader than model risk. You have data risk. You have ethical risk. And you have architectural and technology stability production risk with the increased use of computing demand. You have cyber risks. You have risks associated with Third and Fourth parties. You also have risks around sustainability, climate, and HR. So there are a variety of risks that are far broader than the remit we had before.’

Third, many Generative AI applications involve general purpose platforms (e.g. RAG chatbots) that will ultimately be employed by different users for different purposes across different lines of business. As a result, established risk vectors shift from models (highly constrained, built for purposes) to users (diverse range of use cases, general purpose).



Luke Vilain
AI Governance Lead
UBS

‘When it comes to Generative AI, that dynamic of having a system that can be constrained in terms of purpose and scope is completely shifted – because it’s the AI user who’s deciding what is within the boundaries of the system. So how do we actually end up defining a use case?’

Fourth, the increased use of Generative AI has raised sustainability concerns due to high energy demands, carbon emissions, water consumption, and hardware waste. While AI continues to revolutionize industries, balancing innovation with environmental responsibility will prove crucial in ensuring a sustainable future for AI deployment.



Dr. Paul Dongha
*Head of Responsible AI
and AI Strategy
NatWest*

‘The rapid development and deployment of powerful Generative AI models significantly increases the demand for electricity and water due to the massive computing power required to train and run these models. This can exacerbate environmental risks including climate emissions. Banks need to pre-emptively address this growing externality by extending their risk management and reporting practices. Existing FinOps and GreenOps practices can help with this.’

WHAT ARE BANKS DOING TO ADDRESS THESE RISKS?

In light of the current pace of change driven by AI, banks need to evolve and adapt their existing governance, risk, and compliance frameworks to respond to shifting risks as they evolve. However, no institution can update these guardrails in a vacuum.



Raquel Ettrick Thompson
*Global Head of AI Governance
& Control, BNY*

‘Do you have the right guardrails in place to ensure the solutions are being delivered both internally and externally are appropriately governed? Do you have policies and standards in place that help your employees to understand how to use AI responsibly? Are you measuring the effectiveness of those solutions – encompassing responsibility, fairness, ethics, accuracy, and transparency? And do you have your inventory of solutions mapped? These are the first things regulators are going to come and ask you about...’

Industry practitioners consistently reinforce this process of adaptation follows a clear, logical sequence – working back from first principles.

But what are these first principles? Increasingly, the banking industry is achieving a consensus definition of what falls under the banner of Responsible AI, including:

- Establishing **accountability** for the outcomes and associated risks of AI
- Creating **transparency** of practices and processes in AI development to support ongoing efforts to make AI models and decision making explainable
- Anticipating evolving **regulatory requirements**
- Upholding company **ethical commitments** and operational standards by ensuring fair, unbiased and human centered AI

By mapping responsible principles to responsible practices, leading banks are baking Responsible AI into every step of the production lifecycle – from design principles, to rigorous testing, to post-deployment monitoring, to auditability of each and every risk vector (including data, model, systems/infrastructure, use/process, and legal/compliance).

Ultimately, the more banks enhance these processes, the more streamlined deploying AI use cases will become – representing a significant competitive advantage in terms of reducing current time to production.

WHAT ARE BANKS DOING TO MOVE RAI EFFORTS FORWARD?

By interviewing leading subject matter experts and industry practitioners, Evident identified seven key learnings:

1. **TALENT:** Establish RAI-specific expertise to promote cross-functional collaboration and establish an effective AI Governance Committee focused on data handling, model explainability, and bias mitigation
2. **LEADERS:** Hire a designated Responsible AI leader to complement legacy Governance, Risk, and Compliance functions
3. **TRAINING & EDUCATION:** Invest early in training and education to promote awareness of evolving AI governance, ethical consideration, and trustworthy development processes
4. **CONTROL MAPPING:** Complete a detailed mapping of RAI Principles to AI Controls to ensure clear mechanisms for continuing oversight and accountability
5. **EXPLAINABILITY:** Double down on Explainable AI (XAI) research expertise to provide strong foundation for technical guardrails
6. **SAFETY CHECKS:** Progressively, equip first-line operators with “self-service” diagnostic resources to ensure appropriate use of emerging Generative AI tools, thereby accelerating approval of low-risk use cases
7. **ASSURANCE:** Work towards implementing a true AI Assurance Platform, which can keep pace with a rapidly evolving regulatory environment

Collectively, these seven learnings underscore where and how leading banks are investing – establishing the aspirational criteria by which other organizations (and industries) can assess their relative stage of maturity (see Appendix A: Draft Maturity Model for Responsible AI).

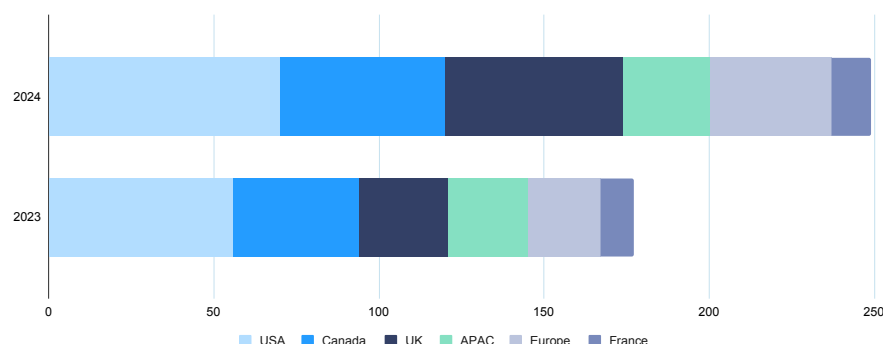
1. Hire AI-specific expertise within governance, risk and compliance functions

When we update the [Evident AI Index](#) each year, we examine the total number of individuals working in roles related to “Responsible AI” across the 50 banks we track – extending to specialization in AI governance, AI risk, and AI ethics. Collectively, we refer to this capability area as “RAI Talent.”

In aggregate, RAI Talent represents a relatively small pool within the wider AI talent stack – totaling 249 individuals across the 50 banks, as of Q4 2024. But bank employees with remits related to Responsible AI are growing at a brisk clip (41% year-on-year). In addition, this talent capability is now found at more banks than ever before. In 2024, we found evidence of these roles at 41 out of the 50 banks in 2024, up from 31 banks in 2023.

This is not only a reflection of the increasing focus on Responsible AI across the world’s largest banks, but also a growing appreciation of the need for individuals with more specific expertise of AI specific risks within these functions.

RESPONSIBLE AI: VOLUME OF RAI TALENT, BY REGION



October 2024 vs.
November 2023, n=41
banks

However, not all banks are equal. Generally speaking, the US banks lead the way with the greatest overall number of these AI specialist positions. As an early mover in this space, we found over 20 matching roles at JPMorganChase – over 3x that found at the average bank.

But this regional view is also changing quickly. Across each geography where banks are headquartered, we observe universal growth in RAI Talent. However, the UK banks and the European banks are accelerating talent acquisition the fastest, growing their share of the overall pie.

This makes sense. While all regions are investing in talent in anticipation of changing regulations and standards specific to AI models and systems – these requirements are not aligned in terms of scope or sequence. As of February 2nd, obligations stemming from the EU AI Act are already hitting regional players (as well as those with operations in the EU).

Not all roles are equal

A double-click into RAI Talent reveals significant variation in the types of roles – and remits – depending on the organization under review. For example, over 70% of specialists in AI ethics found during our sweep work at either European or UK institutions.

While every bank has its own idiosyncrasies with regards to job titles, talent working in Responsible AI positions can generally be grouped into five key buckets (distinct from data scientists, engineers, product managers, etc. in other capability areas):

TEAM	ROLES	RESPONSIBILITIES
1. Responsible AI Risk & Compliance Team	<ul style="list-style-type: none"> → AI Risk Lead → AI Compliance Manager → Model Validation Specialist 	<ul style="list-style-type: none"> → Conduct risk assessments and implement scoring frameworks → Ensure AI models align with global regulations (GDPR, Basel III & IV, SR 11-7, MAS-FEAT, etc.) → Develop AI risk mitigation strategies
2. AI Governance & Policy Team	<ul style="list-style-type: none"> → AI Policy & Ethics Lead → AI Governance Specialist → Regulatory Affairs Liaison 	<ul style="list-style-type: none"> → Ensure risk policies align with compliance requirements → Develop internal AI risk policies to integrate with or complement broader risk management → Provide guidance to business units on RAI adoption and training
3. AI Model Oversight & Assurance Team	<ul style="list-style-type: none"> → Model Risk Manager → Bias & Fairness Auditor → Explainability & Interpretability Lead 	<ul style="list-style-type: none"> → Oversee AI model validation process → Develop model validation guidances → Enforce human-in-the-loop governance → Document and monitor AI model lifecycle
4. AI Ethics & Responsible Innovation Team	<ul style="list-style-type: none"> → AI Ethicist → Human-AI Interaction Researcher → Fairness & Inclusivity Officer 	<ul style="list-style-type: none"> → Ensure AI models adhere to the bank's RAI principles → Implement AI accountability mechanisms (fairness audits, bias mitigation, and algorithmic transparency) → Develop guidelines for high-risk AI applications
5. Technology & AI Security Team	<ul style="list-style-type: none"> → AI Security Lead → Data Privacy Engineer → Adversarial ML Researcher 	<ul style="list-style-type: none"> → Establish an AI incident response framework → Maintain an AI risk register that logs all identified issues and remediations → Provide quarterly reports on AI risk trends and emerging standards (ISO-42001:2023)

As banks build out AI expertise within these distinct teams, collaboration becomes more critical than ever.

To that end, 18 of the 50 banks we track publicly refer to the formation of a cross-functional AI Risk or Governance Committee as a necessary first step towards getting organized around Responsible AI — and an additional 11 banks have expanded the remit of existing working groups to address issues of AI risks and/or Data Ethics.

Beyond that, we've already seen the emergence of a wider range of Working Groups that leverage this diverse range of subject matter expertise, including:

WORKING GROUP	MANDATE
AI Governance	Sets policies, guidelines, and standard for Responsible AI development and usage, including data handling, model explainability, and bias mitigation
AI Ethics	Reviews the ethical implications of new AI projects, ensuring alignment with organizational values and norms
Data Privacy	Focused on protecting user data and ensuring compliance with relevant privacy regulations
Model Review	Assess the performance of and potential biases of AI models before deployment
AI Impact	Analyzes the potential societal and business impact of AI projects, identifying potential risk and mitigation strategies

All of this signals a growing focus on – and organization of – Responsible AI practices within the world's largest financial institutions.


2. Hire a designated Responsible AI leader to complement legacy Governance, Risk, and Compliance functions

With more people comes more complexity, and as this RAI talent pool grows, we're also seeing an increasing need for RAI leadership positions.

In 2024, we found evidence of a senior RAI leader (defined as a VP-level or more senior position whose primary responsibility is to oversee RAI activities, drive awareness, or build expertise within the institution) at 33 of the 50 banks we track. This represents a significant increase from just 15 banks the year prior, indicating a growing centralization and prominence of these roles across the industry.

RESPONSIBLE AI: EVIDENCE OF RAI LEADERS
FOUND ACROSS INDEX BANKS, BY YEAR

October 2024 vs.
November 2023,
n=50 banks



While RAI leaders are becoming more common, their relative seniority, titles, and remit (often extending to Data, Governance, Risk, or Trust & Safety) varied widely across the sample. At this juncture, “Heads of RAI” are still rare – most commonly found across UK banks. This makes sense: UK banks grew RAI roles the fastest last year (2.5x faster than the Index average). Consequently, this is where we see the greatest need for RAI leaders that can manage and direct growing headcount.

BANK	RAI LEADER	CURRENT TITLE	LATEST ROLE CHANGE	REGION
Lloyds Banking	Magdalena Lis	Group Head of Responsible AI	February 2025	UK
NatWest	Dr. Paul Dongha	Head of Responsible AI & AI Strategy	July 2024	UK
CommBank	Nella Luan	RAI Centre of Excellence Lead	May 2024	APAC
UBS	Luke Vilan	AI Governance Lead	March 2024	Europe
Standard Chartered	Emma Johnson	Executive Director, Head of Responsible AI & Data Ethics	February 2022	UK

So given their diversity in profiles and backgrounds, what do RAI leaders actually do?

The primary task of RAI leaders is to establish RAI principles and translate them into guidelines for AI development and deployment. As the RAI leader makes progress in his or her primary task, responsibilities will often expand to include:

- **Risk Assessment:** Reviewing ethical risks associated with new AI projects and strategies, often as a test of enhanced monitoring and/or auditing procedures
- **Education & Training:** Designing RAI training and awareness curriculum specific to RAI principles, ethical considerations, and best practices
- **Stakeholder Engagement:** Communicating RAI initiatives to both internal and external stakeholders, including customers, partners, and regulators

- **Policy Development:** Contributing to the development of company policies related to Responsible AI usage and governance (in tandem with legal, compliance, and reporting/audit functions)
- **Advocacy and Thought Leadership:** Representing the organization as a leader in Responsible AI discussion within the industry

So, what happens next?

As with any other digital transformation effort, centralized roles such as RAI leaders are often stood up to review existing governance and risk frameworks. As modified frameworks and resulting policy and processes evolve into standard operating procedure, functions of the central team become mirrored across applicable areas of the business.

Over time, embedded expertise shifts accountability back to the broader enterprise – as the centralized team retains a monitoring and reporting function (to the degree it's necessary). Given that many banks are in the early stages of their AI transformation, many stakeholders expressed support for a centralized approach:



Luke Vilain
AI Governance Lead
UBS

‘The great value of having a responsible AI function, of having a data ethics function, is we can look over all of these risks and consider them in a holistic way and put the end impact to the customer, to our clients, to our employees, to our other stakeholders, at the real heart of what we’re doing.’

3. Invest early in training and education to promote awareness of evolving AI governance, ethical operations, and trustworthy development processes

While centralized RAI leadership is necessary to “jump start” several critical workstreams, subject matter experts are quick to warn against the potential pitfalls of assigning accountability without a plan for bringing the wider enterprise along on the journey:



Dan Jermyn
Chief Decision Scientist,
CommBank

‘I do feel there’s a fine balance to be had between bringing in specialists that are necessary to really unpick the complexities that AI introduces into the system, really understand what they are, what risks that exacerbates or creates, and also mythbusting to some extent too... But it’s also incredibly important that this work doesn’t become siloed—and that absolutely everybody in the entire bank has the accountability for Responsible AI. We’re very clear on that from the outset...’

This is where training and education become critical. At present, 31 of the 50 banks tracked in the Evident AI Index provide some form of AI training to bank employees. Of these, only 16 banks focus on specific areas of AI, extending to internal tools and the ethics of AI.

In May 2024, CommBank launched “AI for All” – a microlearning series designed to upskill 43,000 employees on Generative AI. In designing the curriculum, Responsible AI was deemed a core topic.

In August 2024, RBC also launched a “Leading in AI” pilot targeting senior leaders and Board members. The one-day crash course is designed not only to impart a deeper understanding of Generative AI, but its potential risks – creating a solid foundation for responsibly integrating AI into operations.

While these two examples provide evidence of how Responsible AI education is being adapted for different stakeholder groups, subject matter experts are quick to stress the importance of training aligned to the lifecycle of AI product development and employment, including:

A. “Safety-First” Approach to AI Systems and Models



Dr. Paul Dongha
*Head of Responsible AI
and AI Strategy
NatWest*

‘Embedding risk management and ethics into culture and bringing education into the early parts of development life cycle is really crucial to avoid the loss of invested cost or realising there’s an issue later.’

B. Evolving Standards & Regulation Impacting Governance, Risk, and Compliance Efforts



Luke Vilain
*AI Governance Lead
UBS*

‘Educating our AI-focused teams on AI governance topics allows us to get the authority to go through meaningful transformation for how control and risk frameworks work, so that the company has a culture of proportional risk taking as we adapt to what’s coming over the horizon. This type of education is complementary to our broader education agenda (e.g. Mandatory learning on Responsible use of Generative AI for all employees in the group).’

C. Accountability for End Users Accessing Generative AI Tool & Applications



Ozge Yeloglu
*VP, Advanced Analytics & AI
CIBC*

‘Education is crucial. At CIBC, access to our Generative AI tool is contingent upon team members completing required training and securing the appropriate permissions for its intended use.’

In terms of people that sit at the nexus of Responsible AI, banks are clearly investing in specialization, leadership, and training. But people are only part of the solution. For Responsible AI efforts to scale effectively across the wider enterprise, banks need to pursue policies that define how risks are identified, assessed, and managed. This includes mapping RAI Principles to discrete AI Controls (short-term), as well as investing to address known gaps in explainability (long-term).

4. Complete a detailed mapping of RAI Principles to AI Controls to ensure clear mechanisms for oversight and accountability

While education and training is often reactive to a rapidly changing operating environment, the painstaking process of mapping of RAI principles to specific AI controls provides a proactive means of getting ahead of future regulatory requirements and emerging standards.

Current documentation of RAI principles articulate the core values that guide the ethical development and implementation of AI. Over the last two years, these frameworks have become more commonplace across the industry. In 2024, 18 banks published their existing set of RAI principles, up from only 6 banks the year prior. In choosing to disclose those statements publicly, banks are capitalizing on an opportunity to reinforce trust with a wider set of stakeholders – including their customers, partners, investors, and regulators.

Yet a critical challenge for banks has been translating those lofty principles (sometimes derided as “platitudes”) into pragmatic AI controls that provide structured, actionable policies and safeguards.

A typical principles-to-controls mapping exercise is often informed by technical regulatory guidance on what a given principle means, how that can be measured, and how it can be matched to the relevant application. For example, the [OCC Comptroller's Handbook](#) explainability tiering. Here are some public examples of control mapping at a high-level:

BANK	PRINCIPLE	AI CONTROL	IMPACT	EVIDENCE
JPMorganChase	Fairness	Bias Detection & Mitigation in Credit Scoring	Ensures equitable access to financial services by preventing discriminatory lending practices	JPMorgan has filed patent applications for algorithmic bias evaluation of risk assessment models
Deutsche Bank	Transparency	Explainable AI in Fraud Detection	Enhances customer trust and regulatory compliance by making AI-driven fraud alerts interpretable and auditable	Deutsche Bank has a long-standing partnership with Google Cloud and NVIDIA to expand and document the model predictions
Capital One	Accountability	Human-in-the-Loop Oversight of AI applications	Reduces risks, ensures accuracy, and enables AI models to learn in real-time from human contextual intelligence	Capital One's prioritizes human-in-the-loop AI applications for less routine, more critical decision-making processes

Each mapping satisfies a very specific requirement. Collectively, this mapping can be used to demonstrate clear mechanisms for continuing oversight and accountability – before customers, partners, and regulators come asking.

Subject matter experts stressed two aspects of control mapping that can make the process more efficient and effective.

First, they emphasized that no one is starting from a blank sheet of paper. Banks maintain a multitude of administrative, physical, and technical controls. In addition, regulations and standards that may become applicable in the near future maintain a highly transparent development and review process. As a result, banks need to focus their efforts on gap analysis.



Raquel Ettrick Thompson
Global Head of AI Governance
& Control, BNY

‘How we approached this was looking at the next [NIST] Risk Management Framework as well as components of the EU AI Act – and map them to our existing technology control framework. Gen AI applications were typically where we had to consciously look at enhancing our existing controls, and in some cases, adding new controls to our Risk & Controls Self-Assessment (RCSA) so that we could account for those risks within our taxonomy.’

Second, they emphasized that control mapping is often progressive - allowing banks to gauge whether use cases satisfy or fail existing controls across a spectrum of risk. Conveniently, this is how banks prefer to implement AI use cases, starting with low-risk applications and slowly, iteratively laddering up the risk spectrum as the countervailing controls allow.



Luke Vilain
AI Governance Lead
UBS

‘This work provides a means to get the authority to go through meaningful transformation for how control and risk frameworks work, so that the company has a culture of proportional risk taking as we adapt to what’s coming over the horizon.’


5. Double down on Explainable AI (XAI) research expertise to provide strong foundation for technical guardrails

With Generative AI, the problem of explainability – how models arrive at their outputs and how that affects AI-driven decision making – has become orders of magnitude harder to solve due to the more dynamic, complex, and “non-deterministic” nature of large foundation models underpinning it.

This not only presents a challenge for all relevant GRC teams, who need to understand how AI models behave and how to govern them, but it is also a largely unsolved problem for banks prioritising Generative AI-driven, client-facing use cases – where the need to provide a breadcrumb trail on how AI impacts decision making is seen to be critical to minimise product liability and maintain user trust.

Many banks have laid the groundwork to address this challenge by deepening explainability expertise within the remit of RAI-specific research, which continues to be a priority topic among a select group of Index companies. In fact, RAI research output has seen a cumulative growth of 136% over the last 3 years – accounting for almost 15% of all research during this period.

RESPONSIBLE AI: TOP-10 BANKS PUBLISHING RAI RESEARCH PAPERS



2022-2024



Dan Jermyn
Chief Decision Scientist,
CommBank

‘Our research work on explainability was the foundation for some of the control processes that we put in place for our model developers and it allowed us to move nimbly and flexibly - so it's actively being used, and we continue to evolve those things.’

Some of the leading RAI research banks, such as JPMorgan Chase, have leveraged [dedicated explainability expertise](#) to develop and publish papers frontier explainability tooling. Whilst others, such as Wells Fargo, choose to [open source that tooling](#) developed off the back of in-house pure research. The combination of developing state of the art capabilities internally and sharing that in a broader research ecosystem reflects the pressingness of the challenge in the scientific community.

While strict explainability remains aspirational (for now), select banks are starting to see dividends from early investments in this space as that capability is translated into technical guardrails for model development.

Control Mapping and Explainable AI both hold great promise for effective policies by which to audit and review AI models and systems, ensuring banks adhere to principles of fairness, transparency, privacy and accountability. But these investments, in and of themselves, do not tell us how policy will be enforced - especially as more AI use cases impact wider swaths of employees spread across numerous lines of business. For that, we need to examine specific processes that are percolating across leading banks.

6. Progressively, equip first-line operators with “self-service” diagnostic resources to ensure appropriate use of emerging Generative AI tools, thereby accelerating approval of low-risk use cases

Meanwhile, as XAI investments come to fruition, banks are getting practical by engineering processes that ensure proper use of Generative AI tools – by isolating and escalating more high-risk applications back to 2nd line operators in charge of monitoring operational risk.

While this process is not new, it is becoming radically more efficient. Early Generative AI use cases historically were accompanied by comprehensive questionnaires that covered every aspect of:

AREA OF INQUIRY	SAMPLE QUESTIONS
Data Governance & Security	<ul style="list-style-type: none"> → What type of data is being used? → Does it involve PII? → Is it being anonymized?
Model Risk & Explainability	<ul style="list-style-type: none"> → Is the model capable of explaining outputs? → How frequently is the model monitoring for accuracy, drift, bias? → Are human oversight mechanisms in place to review AI-generated outputs?
Ethical AI & Responsible Use	<ul style="list-style-type: none"> → Has the AI tool been tested for bias against different demographic groups? → Is there a process for handling misinformation, hallucinations, or bias?
Regulatory & Legal Compliance	<ul style="list-style-type: none"> → Does the AI tool comply with relevant banking regulations? → Are records of AI model training, decision making, and data usage documented? → Has legal counsel reviewed the AI tool's intended use case against compliance requirements?
User Accountability and Risk Controls	<ul style="list-style-type: none"> → Who is responsible for AI-generated decisions and their impact? → Have intended users been trained on appropriate and responsible use? → Are AI-generated outputs validated before being used in other business processes?

While the original versions of these questionnaires were so extensive as to risk discouraging use of Generative AI tools, newer versions have been streamlined based on a logic tree that quickly reconfigures and narrows the question set based on some initial screening prompts clarifying user intent.

Front loading with these self-service tools establishes an “AI risk management by design” workflow that is both proactive and self-improving. In effect, diagnostic questionnaires make the job of 2nd line validators more efficient, while providing clear escalation tracks and remediation steps to corresponding low-, medium- and high-risk applications.

By eliminating arduous, manual review of each and every variation of closely related use cases with marginally different parameters, banks can accelerate use case deployments and reduce average time to production. Without these automated procedures in place, banks never reach the “trust but verify” stage of AI transformation. Ultimately, they never achieve underlying confidence in the AI system’s capabilities – blunting how quickly systems that demonstrate promising AI can be scaled across the enterprise.

For banks that have prioritised processes like these, there is a clear pathway towards a rigorous, scalable and automated control environment to manage this evolving risk landscape:

BANK	WHAT?	HOW?	IMPACT?
UBS	Crafted a questionnaire diagnostic tool based on AI principles documentation	<p>Starts by asking developers to flag key risks associated with a use case, then outlining steps taken to mitigate those</p> <p>An independent second line team, often composed of data ethicists, reviews responses, often asking to show how they’re addressing RAI principles and what they’re implementing to address those.</p> <p>Above a certain risk threshold, that team then issues a formal impact assessment to approve or deny the use case to production</p>	Provides a frictionless bridge between first and second line teams, with AI principles becoming more embedded into the design stage.
BNY	Built a decision tree matrix into user questionnaires to help isolate and elevate high risk use cases that merit additional scrutiny	<p>Review of AI use cases goes through multiple pathways, starting with the Review Board on Data Usage.</p> <p>While the original questionnaire design was 60+ questions, the latest version starts with a set of only 14 “triage” questions – which trigger more or less questions depending on the answers provided.</p>	By “front-loading” questions in terms of criticality and risk, the bank is effectively educating 1st line operators about where a given AI application falls on the risk spectrum – and providing escalation points that help isolate high-risk instances against the backdrop of expanding use cases.

7. Work towards implementing a true AI Assurance Platform, which can keep pace with a rapidly evolving regulatory environment

While diagnostic questionnaires provide a critical safety check as deployment of AI use cases expands, they do not allow organizations to see the wider picture – encompassing automated assessments, audit capabilities, and governance controls. To accomplish this, banks are developing AI Assurance Platforms – a term that is currently ascending to a place of prominence in the hype cycle.

At the most basic level, an AI Assurance Platform is the amalgamation of systems, frameworks, and toolkits designed to evaluate, monitor, and ensure AI models operate in compliance with a rapidly changing regulatory landscape. As Assurance Platforms become more automated, they will help reduce risks, enhance trust, and ensure compliance – making them an essential part of RAI strategy not only for banks, but other highly regulated industries.

Consequently, one of the first areas where industry practitioners expect to see dividends from Assurance Platforms is in regulatory compliance by providing automated compliance checks and audit trails aligned against global AI regulations (e.g. EU AI Act, NIST AI Risk Management Framework, Federal Reserve Board SR 11-7).

By providing a “one-stop shop” or “clearing house” to validate organization-, system-, and model-level information, banks will significantly reduce the legal and financial risk related to (unintentional) non-compliance in a world where fast-moving technology competes with slow moving policy guidance.

Today, banks are already demonstrating several building blocks that demonstrate the potential of models and systems designs to evolve alongside changing requirements - both internal and external:

BANK	WHAT? (STATUS)	HOW?	WHY?	IMPACT?
CommBank	Combining model check automation with human oversight (implemented)	Technical control measures, such as assessing the stability of models, run by automated tools.	<p>'Increased ability to automate the control environment of the AI you deploy is not only going to make you safer, it's going to make it quicker and more efficient as well.'</p>  <p>Dan Jermyn <i>Chief Decision Scientist, CommBank</i></p>	Creating greater auditability and volume of model checks
CIBC	Operationalizing bank's AI Risk Assessment Process through a tool that facilitates a thorough identification, assessment, and mitigation of AI-related risks for all their AI use cases (implemented)	The tool leverages the best practices in risk mitigation methodologies to provide teams with the necessary guidance to identify applicable AI-related risks for a particular AI use case, assess those risks collaboratively with various stakeholders, and then track proposed mitigations to completion with appropriate accountability.	<p>'We believe effective AI Governance and Risk Management requires that teams adopt a lifecycle approach, i.e., Governed AI by Design, through which AI-related risks are considered at every stage of the AI lifecycle. Our AI Risk Assessment Tool helps us achieve precisely that.'</p>  <p>Aditya Anne <i>Senior Director, Enterprise AI Governance, CIBC</i></p>	Ensures a robust and standardized treatment of AI-related risks for every AI use case which is essential for effective AI Governance
BNY	Enhancing monitoring dashboard for universal platforms, focused on anomaly detection and critical alerts (iterative)	In addition to monitoring and alerting capabilities, the bank's primary Gen AI platform features technical guardrails designed to address evolving LLM safety risks.	<p>'Being able to articulate how we operationalize Responsible AI and data commitments is a key driver of trust. Our Eliza platform is helping to enable competitive advantage, because it combines the facets of AI innovation and governance into a single platform. We're combining our principle-based control framework with democratized access to AI.'</p>  <p>Kirsten Mycroft <i>Chief Privacy & Data Ethics Officer, BNY</i></p>	When requirements change, platform development teams can quickly flag any instances that merit review in response to newly published standards or guidelines
NatWest	Building towards a company-wide, distributed Assurance Platform for end-to-end management of AI risk (pilot)	<p>One-stop-shop platform that aligns policies, compliance criteria and technical measures to manage AI risk from development to production.</p> <p>Ensures only compliant models progress while monitoring post-deployment risks like model drift.</p>	<p>'The idea with the Assurance Platform is that it ties everything together so you have a consistent and efficient single source where model validation teams can provide assurance across all of the controls stemming from your internal and external principles and policies, including the EU AI Act, ISO, OECD.'</p>  <p>Dr. Paul Dongha <i>Head of Responsible AI & AI Strategy, NatWest</i></p>	Improves efficiency of AI lifecycle risk management




NEXT STEPS TOWARDS A MATURITY MODEL

Following our review of People, Policy, and Process, we were left with an open-ended question. Given what we know now about the current state of play versus where leading banks are headed next, what are the aspirational criteria by which other organizations (and industries) can assess their current stage of maturity with respect to Responsible AI.

To that end, we've assembled a Draft Maturity Model that summarizes the progressive steps along the journey from "emerging" capability to "pioneering" features of a leading bank in this domain.

March 2025

In this operating environment, it remains critical to define where institutions sit on a comprehensive **Maturity Model**. At present, we observe many leading banks profiled in the [Evident AI Index](#) have reached Level 4 on our emerging set of criteria outlined below. Please reach out with your thoughts and feedback on this construct as we continue to explore the building blocks for the “bank of tomorrow.”

	LEVEL 1 EMERGING	LEVEL 2 DEVELOPING	LEVEL 3 OPERATING	LEVEL 4 SCALING	LEVEL 5 LEADING	LEVEL 6 PIONEERING
 PEOPLE	RAI roles exist, but remain highly dispersed	Dedicated RAI leadership roles emerge, but lack line of sight to strategy and executive sponsorship	RAI staff drive agenda across newly formed AI committees and working groups	RAI education and training expands across organization in tandem with adaptation of existing policies and processes	RAI staff expands across relevant business units, evolving from a centralized to “embedded” function	Previously scarce RAI expertise becomes widely available and accessible across the larger enterprise
 POLICY	Early conversations evidenced by posts referencing AI ethics and risk	RAI Principles circulated internally for testing and feedback	RAI Principles shared proactively and publicly	AI risk management and governance teams engage in comprehensive mapping principles to specific controls	Bank actively incorporate outputs from long-term investment in Explainable AI (XAI)	RAI integrated in bank’s overall AI strategy
 PROCESS	Responsibility for managing AI risk mapped onto existing governance structures	Cross-functional teams assembled to explore adaptation of existing risk management frameworks	Initial focus on AI Governance quickly expands to AI Ethics, Data Privacy, Model Review, and AI Impact	AI risk assessments and compliance checks become more self-service and scalable	Human-in-the-loop verification and active monitoring of enterprise AI platforms demonstrate promise of risk management automation	Automated assurance platforms help organizations quickly respond to organization-, system-, and model-level audit requests

Appendix B

Glossary of Terms

TERM	DEFINITION
Responsible AI	<p>In banking, refers to the development and deployment of AI in a way that mitigates risks, whilst complying with existing and emerging regulations and ethical standards.</p> <p>It integrates robust and established governance, risk, and compliance processes with the need to adapt to emerging AI risks; serving as a specialist function that guides and validates the development of AI in the business so that it aligns with legal and societal expectations.</p>
AI Governance	The framework of policies, practices, and oversight mechanisms that ensure AI systems are developed, deployed, managed and monitored in a way that aligns with the governance values of the organisation
AI / Data Ethics	The deliberation and associated guidance on how the impacts of AI development and usage can uphold company and societal values such as fairness, transparency, privacy and accountability
RAI Principles	Core ethical and operational guidelines that underpin Responsible AI practices, typically including transparency, accountability, fairness, privacy, security, and human oversight. Shared with multiple stakeholders as a north-star to enable trust
RAI Policy	An expansion of RAI principles that are used internally to offer more practical guidance on how those should be implemented for the development and use of AI
RAI Talent	Professionals with the skills and remit to address Responsible AI: this includes ethics and policy specialists, but also adapted roles in governance, risk management, legal and compliance
RAI Leader	An individual responsible for orchestrating Responsible AI initiatives within an organization: ensuring alignment across the controls teams that manage and govern AI risk, as well as with AI development and implementation staff
AI Risks	Exposure a company faces to negative consequences of AI systems. This includes risks that fall under pre-existing types - such as operational or model risks - but the list is evolving to also include: Data, Ethical & Sustainability
AI Risk Committee	A dedicated group within an organization that oversees AI-related risks, ensuring AI systems are aligned with regulatory, ethical, and governance standards.
Explainable AI / Explainability	Ensuring that an AI system provides understandable and interpretable outputs, and traceability for the decisions an AI system makes. This is relevant for technical stakeholders to refine the robustness and auditability of AI models, but also for end users to have visibility of if and how AI effects decisions made about them
AI Bias	Systematic errors or unfair outcomes in AI models that result from skewed data, flawed algorithms, or unaccounted for human biases - with the risk of discriminatory or unfair decisions being made
AI Audit	The systemic review and assessment of AI use cases to evaluate how they comply with existing and emerging regulations and internal policies, if they introduce unacceptable risks, and whether they meet operational and ethical standards.
Assurance Platform	<p>A system or tool that supports the oversight, monitoring, and validation of AI use cases by demonstrating it has been aligned with regulatory requirements, internal policies, technical controls, operational and ethical standards</p> <p>These tools are often integrated with centralised platforms or inventories that manage the production lifecycle for AI use cases</p>
Diagnostic Questionnaires	A list of questions, typically based on Responsible AI policies, that are integrated with 1st line AI development workflows to identify risks and explain approaches to mitigating those risks, so 2nd line operational staff engage more easily manage and approve use case production.
1st Line Staff	First line is generally embedded in individual business areas, with risk managers serving as both subject matter experts immersed in the business while also attempting to remain impartial enough to weigh potential exposures against business strategy.
2nd Line Staff	Second line is fully removed from day-to-day decisions that happen in individual lines of business, focusing on overarching strategic imperatives and broader areas of operational risk (training and education, policy enforcement, technology compliance, resiliency planning for internal, external, or human trigger events).

Evident provides the most in-depth analysis and tracking of how banks are adopting AI, and publishes the annual Evident AI Index (Expanding to Insurance in June) that benchmarks the Talent, Innovation, Leadership and Responsible AI capabilities of 50 major banks.

Between Index updates, Evident members have access to in-depth analysis of each of the four Index pillars—fuelled by live data trackers and in-depth interviews with senior leaders across the sector—as well as high-profile industry events.

Evident membership includes:



INSIGHTS

Keep up to date with the latest sector-wide AI activity with access to our latest long and short-form research:


Member Hub: Evident members can access our latest insights each month, spanning 10 sub-pillars of the Evident AI Index, including Talent Development, Talent Capability, Research, Patents, Ventures, and more...

Insights Reports: flagship quarterly reports dive deeper into the latest sector-wide trends, depicting the current state of play of 50 global banks across each pillar of the Index.

Coverage includes: Talent, Innovation, Leadership, Responsible AI, and Outcomes.




Available Now



Available Now



Available Now




Next Update: May 2025



BENCHMARKING DATA


Evident members receive access to additional assets that allow them to identify opportunities, accelerate initiatives, review progress, and prepare for the next Index update. These assets include the Index Dashboard and Index Diagnostic Report. The Dashboard affords access to raw data and scores for all 50 banks across all 90 indicators underpinning the Evident AI Index. The Diagnostic Report provides a bespoke analysis of the bank's Index performance versus select peers, clearly identifying strengths and opportunities specific to individual banks.

If you are interested in learning more about your bank's performance profile, please [Contact Us](#) to learn more about Membership product offering and pricing options.



● Among its peers, BANK is the most improved overall, with the largest ranking increase in Leadership and smaller gains in more heavily weighted pillars, Talent and Innovation

INDEX		Talent		Innovation		Leadership		Responsible AI	
Bank	Score	Bank	Score	Bank	Score	Bank	Score	Bank	Score
Bank A	7	Bank B	9	Bank C	7	Bank D	5	Bank E	3
Bank F	10	Bank G	11	Bank H	9	Bank I	12	Bank J	10
Bank K	15	Bank L	15	Bank M	11	Bank N	10	Bank O	19
Bank P	27	Bank Q	23	Bank R	17	Bank S	33	Bank T	50
Bank U	26	Bank V	28	Bank W	21	Bank X	34	Bank Y	45



The [Evident AI Index](#) assesses the various approaches banks are taking towards AI readiness. The October 2024 Index covers 50 of the largest banks in North America, Europe, and Asia. Each bank is assessed on 100+ individual indicators drawn from millions of publicly available data points specific to four pillars: Talent, Innovation, Leadership, and Transparency.

COMPANY	RANK +/- YoY change	Talent Capability & development		Innovation Research, patents, ventures, ecosystem		Leadership in public comms and strategy		Transparency of responsible AI activities	
JPMorganChase	1		2		1		3	↓1	1
Capital One	2		1		3	↓1	14	↓6	14
Royal Bank of Canada	3		6	↑1	2	↑1	8	↓3	10
Wells Fargo	4		4		5		36	↓4	8
CommBank	5		↑1 7	↑9	12	↑8	2	↑1	5
UBS	6	↓1	3		18	↑1	15	↑2	15
HSBC	7	↑6	15	↑2	9	↑3	5	↑19	3
Citigroup	8	↑1	5		7	↑1	13	↑13	43
TD Bank	9	↑2	10	↑5	8	↓2	20	↑11	16
Morgan Stanley	10	↑7	18	↑2	4	↑3	17	↓3	42
Goldman Sachs	11	↓4	13	↓5	6	↓2	27	↓12	37
BNP Paribas	12		12	↑6	14	↓1	6	↓2	21
BBVA	13	↑13	8	↑22	20	↑1	7	↑6	26
BNY	14		20	↑3	13	↑2	10	↑2	11
Bank of America	15		11	↑3	11		18	-11	20
DBS	16	↓6	21	↑1	50	↓16	1		7
ING	17	↓9	9	↓3	19	↓9	9	↑10	24
NatWest	18	↑1	16	↓5	23	↑2	16	↑14	9
Standard Chartered	19		↑4 31	↓4	15	↑7	11	↑22	2
Scotiabank	20	↓2	22	↓1	36	↓8	22	↑1	4
Société Générale	21	↓1	17	↓8	30	↑2	23	↓7	29
CIBC	22	↑19	29	↑11	28	↑1	24	↑22	6
Intesa Sanpaolo	23	↑2	46	↑1	24	↓8	4	↑5	12
Bank of Montreal	24	↓8	27	↓2	10	↓1	39	↓17	17
Barclays	25	↓3	19	↓7	21	↓7	45	↓18	28
Deutsche Bank	26	↑3	25	↑7	16	↑24	31	↓21	33
Lloyds Banking Group	27	↑1	28	↓9	17	↑26	33	↑5	19
Santander	28	↓7	23	↑1	31	↓13	34	↓23	18
ABN AMRO	29	↓5	14	↓1	35	↓9	35	↑6	46
Rabobank	30	↑1	24	↓14	41	↑5	29	↑14	30
State Street	31	↑8	38	↓2	26	↑15	21	↑7	27
Crédit Agricole	32	↓5	34	↓3	25	↓8	28	↑6	40
ANZ	33	↑5	36	↑7	40	↑2	12	↑6	25
NAB	34	↑1	32	↑10	46	↑2	32	↓7	13
PNC Financial	35	↑1	26	↑8	38		44	↓7	35
US Bank	36	↓2	30	↓2	32	↑4	37	↑2	32
Westpac	37	↓4	37		39	↓15	19	↑17	22
KBC	38	↓1	45	↓7	22	↑11	26	↓6	39
Citizens Financial	39	↑10	33	↑8	47	↑3	30	↑14	47
Groupe BPCE	40	↑6	43	↑3	37	↓2	43	↑5	41
CaixaBank	41	↓1	40	↑8	45	↑2	38	↓32	38
Danske Bank	42	↑2	35	↑4	49	↓5	48	↓8	34
Truist Bank	43	↓13	44	↓15	34	↓7	46	↓17	45
Crédit Mutuel	44	↑1	50	↓1	27	↑3	25	↓4	31
Commerzbank	45	↓3	48	↓4	42	↓3	41	↑4	23
UniCredit	46	↓3	47	↓2	29	↓6	50	↓8	36
Raiffeisen Bank Intl	47	↓15	41	↓15	43	↓12	42	↓7	49
Nordea	48		39	↓4	48	↑1	47	↑2	44
Charles Schwab	49	↓2	42	↓9	44	↑1	40	↑7	48
First Citizens	50		49	↑1	33	↑4	49	↑1	50