accenture

In association with ELEVANDI

RESPONSIBLE AI IN FINANCE: NAVIGATING THE ETHICS OF GENERATIVE AI

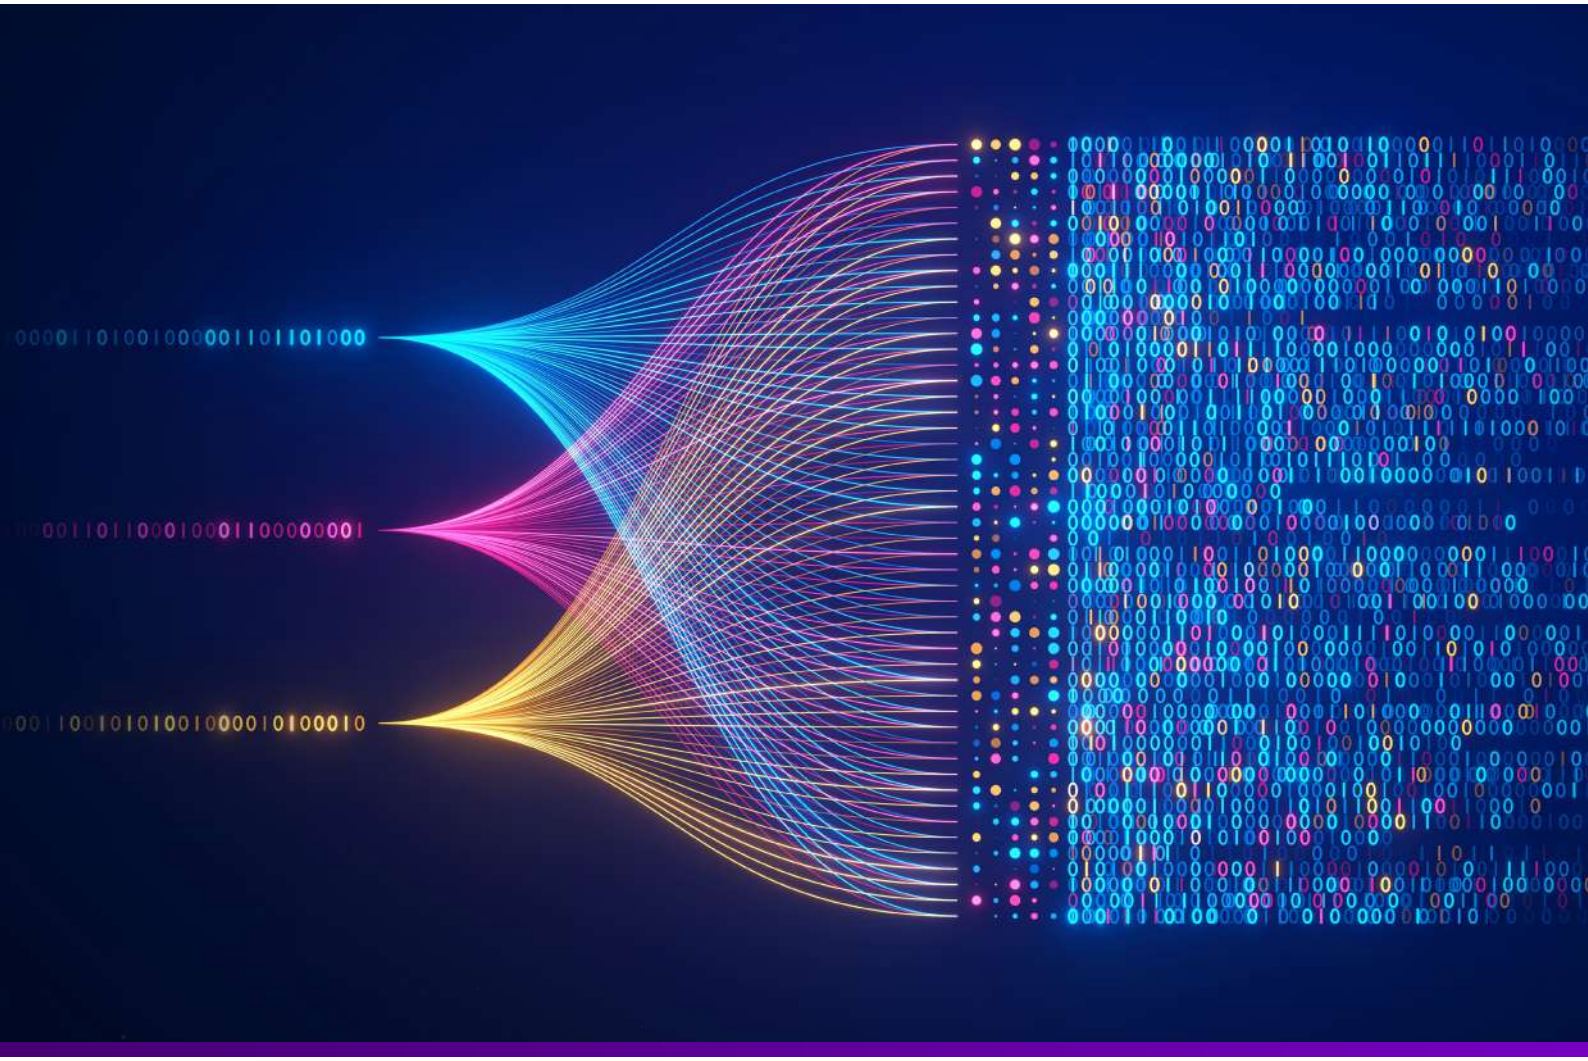

TABLE OF CONTENTS

INTRODUCTION

Generative artificial intelligence, or GenAI, has advanced rapidly, and its potential to disrupt industry has caused as much excitement as it has alarm.

Accenture noted in its Technology Vision 2023 report that the question businesses need to ask is not whether large language models (LLMs, which are the foundation of GenAI's output) "will impact their industry, but how", adding that such foundation models "have the potential to transform human-AI interaction".[1]

Banks in the US, Europe, and some growth markets have rolled out GenAI-powered solutions in areas like customer service, marketing, code creation, fraud detection, financial modelling, analysis of legal contracts, and risk management. There is much more to come (see graphic).[2]

At the same time, awareness of GenAI's risks has never been higher. Though talk of a robot revolution remains hyperbolic, the use of GenAI carries genuine risks—as even some of its inventors have warned. At the June 2023 Point Zero Forum, Singapore's Deputy Prime Minister Heng Swee Keat noted that: "The current debate on GenAI is instructive—the initial excitement has been followed by strong calls to put in place guardrails or even halt its development."[3]

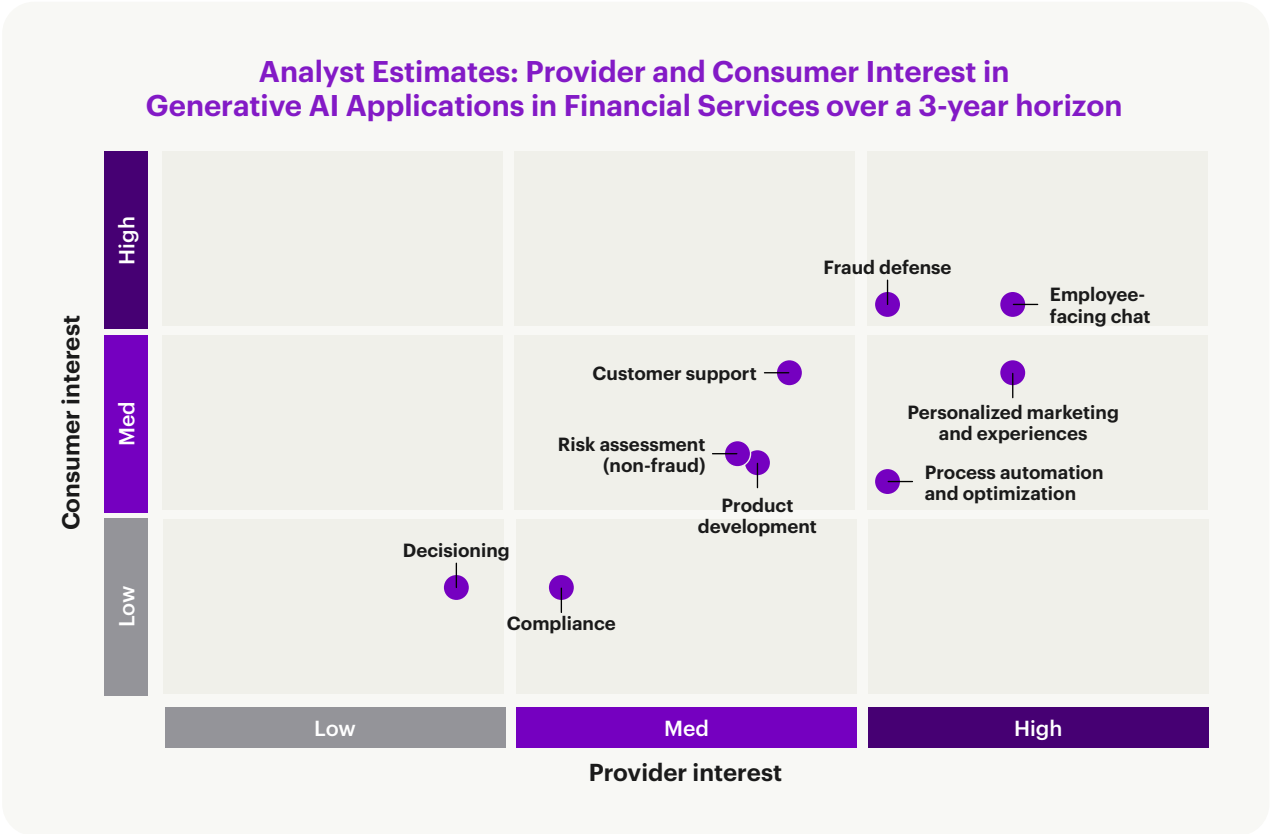

Source: Insider Intelligence, 2023

[1] *Technology Vision 2023: When Atoms meet Bits: The foundations of our new reality*
See: https://www.accenture.com/us-en/insights/technology/technology-trends-2023
[2] *ChatGPT and Generative AI in Financial Services: Reality, Hype, What's Next, and How to Prepare, Insider Intelligence (March 2023).*
See: https://www.businessinsider.com/chatgpt-and-generative-ai-in-financial-services-how-to-prepare-2023-4
[3] *Speech by Mr Heng Swee Keat, Deputy Prime Minister and Coordinating Minister for Economic Policies, at the Point Zero Forum on 26 June 2023, MAS (June 26, 2023). See: https://www.mas.gov.sg/news/speeches/2023/speech-by-dpm-heng-swee-keat-at-the-point-zero-forum*

Risks already exist in AI models, such as algorithmic bias, lack of transparency, and potential for misuse. GenAI is subject to those risks as well but brings with it additional risks like model confabulations (where erroneous outputs are generated due to harmful or erroneous training data) and opaqueness of data inputs. As GenAI use cases continue to expand rapidly, there is increasing attention around the responsible use of GenAI, whether in finance or elsewhere.

Against this backdrop, a roundtable of banks, regulators, and industry leaders convened at the Point Zero Forum to discuss *Responsible AI in Finance: Navigating the Ethics of Generative AI*.

This paper is based on the discussions at that roundtable, as well as extensive Accenture research. It touches on the trajectory of LLM deployment in financial institutions; how existing principles like fairness, ethics, accountability, and transparency can be applied and whether additional principles need to be introduced in the context of LLMs; and how financial institutions and regulators can prepare for the disruption to come. As the meeting did not include a formal mechanism to establish consensus on any specific issue, takeaways and findings presented in the paper are a consolidation of individual opinions and not representative of the group's collective stance.

In sum, all stakeholders will need to focus on areas including what data is ingested, how models are generated, and how outputs are used. The three key lessons that participants agreed upon were the need to: a) move cautiously and consider all aspects of benefits and risks; b) go beyond conceptual thoughts to start experimenting on pilots; and c) work collaboratively across organisations to share learnings. This will allow the financial services industry to position itself for the responsible and ethical use of GenAI.

GenAI: DEFINING PROGRESS & PROBLEMS

GenAI refers to a broad spectrum of learning algorithms that can make predictions, write text, or create visual media. These are powered by machine-learning models in a process known as large-language modelling (LLM), which is built upon huge amounts of data.[4] AI chipmaker Nvidia describes LLMs as "a deep-learning algorithm that can recognise, summarise, translate, predict, and generate text and other content based on knowledge gained from massive datasets".[5]

Unlike traditional AI, Gen AI can be used far more easily by non-technical specialists. Accenture's research shows financial services will be the most disrupted sector, with an estimated 54-66 percent of working hours affected. This makes defining the responsible use of Gen AI urgent.

Impactful: How GenAI will affect work-time distribution across different industries.

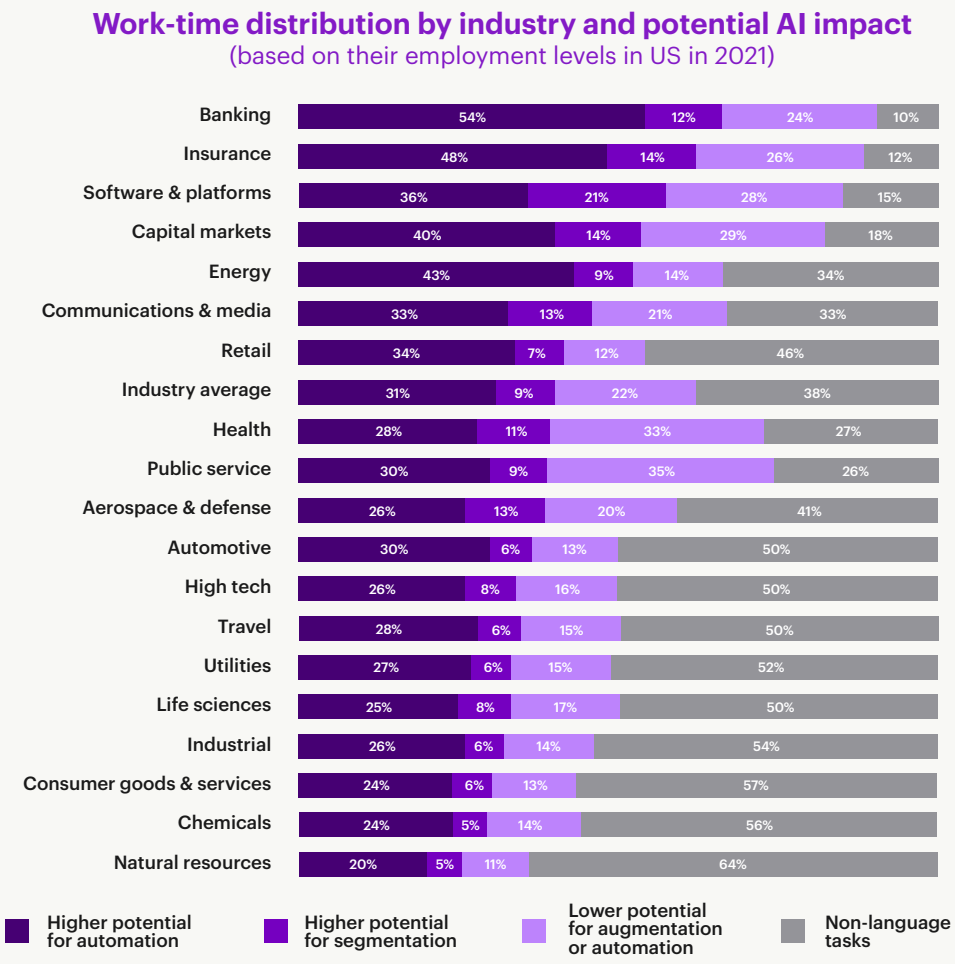

Source: Accenture Research based on analysis of Occupational Informational Network, US Dept. of Labor, US Bureau of Labor Statistics

---

[4]*A new era of generative AI for everyone, Accenture (2023). See: https://www.accenture.com/content/dam/accenture/final/accenture-com/document/Accenture-A-New-Era-of-Generative-AI-for-Everyone*

[5]*See: https://blogs.nvidia.com/blog/2023/01/26/what-are-large-language-models-used-for/*

[6]*A new era of generative AI for everyone, Accenture (2023). See: https://www.accenture.com/content/dam/accenture/final/accenture-com/document/Accenture-A-New-Era-of-Generative-AI-for-Everyone*

Exponential gains from GenAI are expected in the years to come, but that also brings with it exponential risks. In this context, the challenges most pertinent to the use of GenAI are:

First, **mistakes made by GenAI can still be deceptively convincing.** There are numerous examples of GenAI citing fictitious sources, for example, and generating errors—so-called confabulations. The detection of such confabulated outputs may be difficult for the untrained specialist.

Second, **GenAI can reinforce embedded biases** in data, and because GenAI is built on foundation models, any biases inherent in those datasets that they are trained on could be easily propagated to their derivatives as well.

Third, the larger volume and variety of data that GenAI is trained on means **issues around privacy, confidentiality, and copyright** are more pronounced. Adding data that is scraped from restricted sites and personal social media to GenAI's corpus may incidentally violate the contextual integrity for which such information was originally intended to be shared.

With the challenges posed and opportunities availed, roundtable participants represented by financial institutions and regulators discussed the crucial considerations and actions to take to embrace GenAI in a safe manner. There were five key takeaways.

TAKEAWAY 1: EMBRACING A WATERSHED MOMENT WITH GenAI

When asked about the potential use cases of GenAI, one participant from a global bank put it as, "an easier question would have been: What are not the use cases of GenAI? We could maybe answer that quicker."

Banks globally have deployed GenAI in a variety of areas, including customer service, marketing, and reviewing contracts. Industry leaders expect usage to expand rapidly to, as Nvidia puts it, "create hyper-personalised customer content, automate document summarisation to reduce manual work, and analyse terabytes of public and private data to generate investment insights".[7]

Where GenAI is having a profound impact is in helping banks to make use of their data. As much as 80 percent of a typical bank's data is unstructured,[8] and therefore difficult to analyse without extensive manual tagging. GenAI's power lies in helping banks to consume the data in a way that fully realises its potential.

Another area where banks are seeing significant efficiencies from deploying GenAI is in boosting developer productivity. In these early days, GenAI is best used to assist humans in developing code and with the testing lifecycle, and there is potential to oversee the end-to-end process in the future.

Among the many changes GenAI will bring are that repetitive tasks will disappear, along with certain jobs, while other roles will emerge. Planning for this shift starts by ensuring that more schools teach basic computer science. Other changes are needed further along the educational chain, including ensuring that more of those joining the workforce have the technical skills necessary to drive the next generation of data science tools.

Additionally, firms need to train staff in a range of specialist skills to transform their roles from creators of content to reviewers. Other key skills include data literacy, systems-level thinking, critical thinking, and data science ("to understand what's happening under the hood, everybody in the organisation needs some data science," said one participant from a financial institution). And firms need to equip staff with the right tools to ensure they can safely explore this technology.

Also crucial, said a quant specialist, is that "we need people with more mathematical literacy, because that is the underlying foundation for any of these skills". Those mathematical skills are vital for logical reasoning, as organisations seek to define criteria such as attribution, transparency, fairness, interpretability, privacy, and other terms in relation to GenAI in a measurable or quantifiable way.

It was noted that there are multiple ways that one could mathematically define fairness, and they are not compatible. This means that a key decision still needs to be made on which fairness criterion is most appropriate for a given use case. In addition, there are currently no mathematical criteria for accountability and interpretability.

---

[7] *AI-Fueled Productivity: Generative AI Opens New Era of Efficiency Across Industries, Nvidia (July 13, 2023).*
See: https://blogs.nvidia.com/blog/2023/07/13/generative-ai-for-industries/
[8] *Unlocking the benefits of unstructured data in banking, Fintech Futures (October 20, 2020).*
See: https://www.fintechfutures.com/2020/10/unlocking-the-benefits-of-unstructured-data-in-banking/

TAKEAWAY 2: GenAI IS NOT A PANACEA

Despite the clear potential of GenAI, the watchword among roundtable discussants was "caution". Many organisations have started deploying GenAI only in low-risk areas to better understand the potential issues. They are generally looking at this through three lenses:

- **Can they do it:** Do they have consent, and are the security and privacy conditions in place?
- **Should they do it:** Even if it is legal to do so, is it ethical to use AI in a particular context?
- **How do they do it?** Which modelling approach would work best? Do you need the best LLM available or would a more light-weight model suffice?

GenAI also has its limitations. Architectural limitations mean, for instance, that it is not suitable for processes requiring a real-time response with wall-clock latency below around 10 milliseconds.

Deploying a model is no substitute for getting the foundation database right in the first place. One participant reflected that: "It's not just about the modelling. It's also about data input, data sources, prompt engineering, knowledge bases. How do you keep them current? What are the metadata of the knowledge to make it useful?"

"We also have to think about the output. What's the traceability of the output? How do we make sure there's no toxicity? All of these are end-to-end considerations where the model is a small part," they said, advocating that guardrails be put up early on to avoid major unintended consequences.

There are also crucial considerations around managing third-party risks. These encompass aspects like data management, model risk management, governance, and outsourcing.

One regulatory participant highlighted two areas that require further thought. The first is that, given the amount of data, compute and expertise needed to build LLMs, at least some firms may buy them from a handful of dominant providers. These providers' desire to protect their intellectual property may hamper an individual firm's ability to adequately interrogate how the models work—thereby affecting their ability to comply with regulatory expectations. If those third-party models eventually came to support critical business operations at scale, those third-party providers could become systemically important.

Additionally, combining LLMs with other software systems could create new risks, moving LLMs away from being oracles to a place where their outputs could determine actions taken by other software. This may create cascading risks within and across firms with increasingly complex networks of interacting AI models.

Lastly, another participant pointed out that holding oneself accountable to customers will be increasingly important to reinforce trust in financial service provision, especially with the increase in associated risks.

TAKEAWAY 3: ADOPT A RISK-BASED APPROACH TO BALANCE BETWEEN INNOVATION AND POTENTIAL HARM

The use of GenAI brings new and complex challenges that banks must address to ensure they implement it in a responsible manner. Accenture grouped these risks broadly into two areas: inputs into GenAI and outputs from GenAI (see graphics).

Risks to consider for Inputs into GenAI

Inputs to GenAI systems

Emerging risks and challenges

Fairness | Ownership | Transparency | Inclusion

Emerging best practices for mitigation

Data

The models powering GenAI require huge amounts of data to function. Many of the leading systems today were trained on the public internet.

Computing power

Creating a GenAI model requires expensive supercomputing capabilities.

GenAI could learn from biased or harmful training data

Inputs could be biased, violate the privacy of individuals, or encourage the disrespect of fundamental rights.

As datasets grow and the degree of qualitative human input in model-building increases, it becomes challenging to supervise the training process.

GenAI could violate copyright laws and hurt creators

The use of public internet data to train GenAI has caused public backlash, especially among creatives, over claims of plagiarism and copyright infringement.

The legal status of GenAI models trained on unlicensed content remains ambiguous in many key jurisdictions.

GenAI's data inputs could be opaque to the public

The large data requirements for GenAI and the opacity of training processes, make public scrutiny of them challengitng.

Public scrutiny can help identify biases and limitations.

The public can struggle to identify the data inputs to GenAI models that they interact with and to impose scrutiny on how they operate.

GenAI could widen the barriers facing certain language and groups

The langugages with the largest repositories of digitally-readable training data - English, Chinese, Spanish, and French - have disproportionate access to quality GenAI models.

$4.6M
to train GPT-3

The high cost of the computing power behind GenAI training limits opportunities for SMEs or service-providers in minority languages.

- Robust, verified datasets, supported by human annotators
- Transparency around inputs into GenAI models and legal, consensual use of data
- Investment in technologies and capabilities for minority languages.

- Demonstrate suitability of data sources and governance
- Publish summary of copyrighted training data
- Measures for SMEs

- Copyrighted information must be licensed; private information must be consensual
- Use accurate and diverse data, annotated according to clear rules

Risks emerging from inputs to Generative AI

Input risks can be categorised into the following:

- Lack of fairness: Data could be biased or harmful.
- Unclear ownership: Humans are needed in or on the loop to ensure accountability.
- Lack of transparency: transparency and explainability could become increasingly challenging with more sophisticated technology.
- Lack of inclusivity: Much of the data being used to train models is in English and in the public domain, which could affect contextualisation around social norms and privacy.
- Need for finetuning: Finetuning models is hugely expensive, which is why some banks are training smaller models on very high-quality data for internal use.
- Model development risk: Given that model risk is crucial, model validation is key.

Risks to consider for Outputs from GenAI

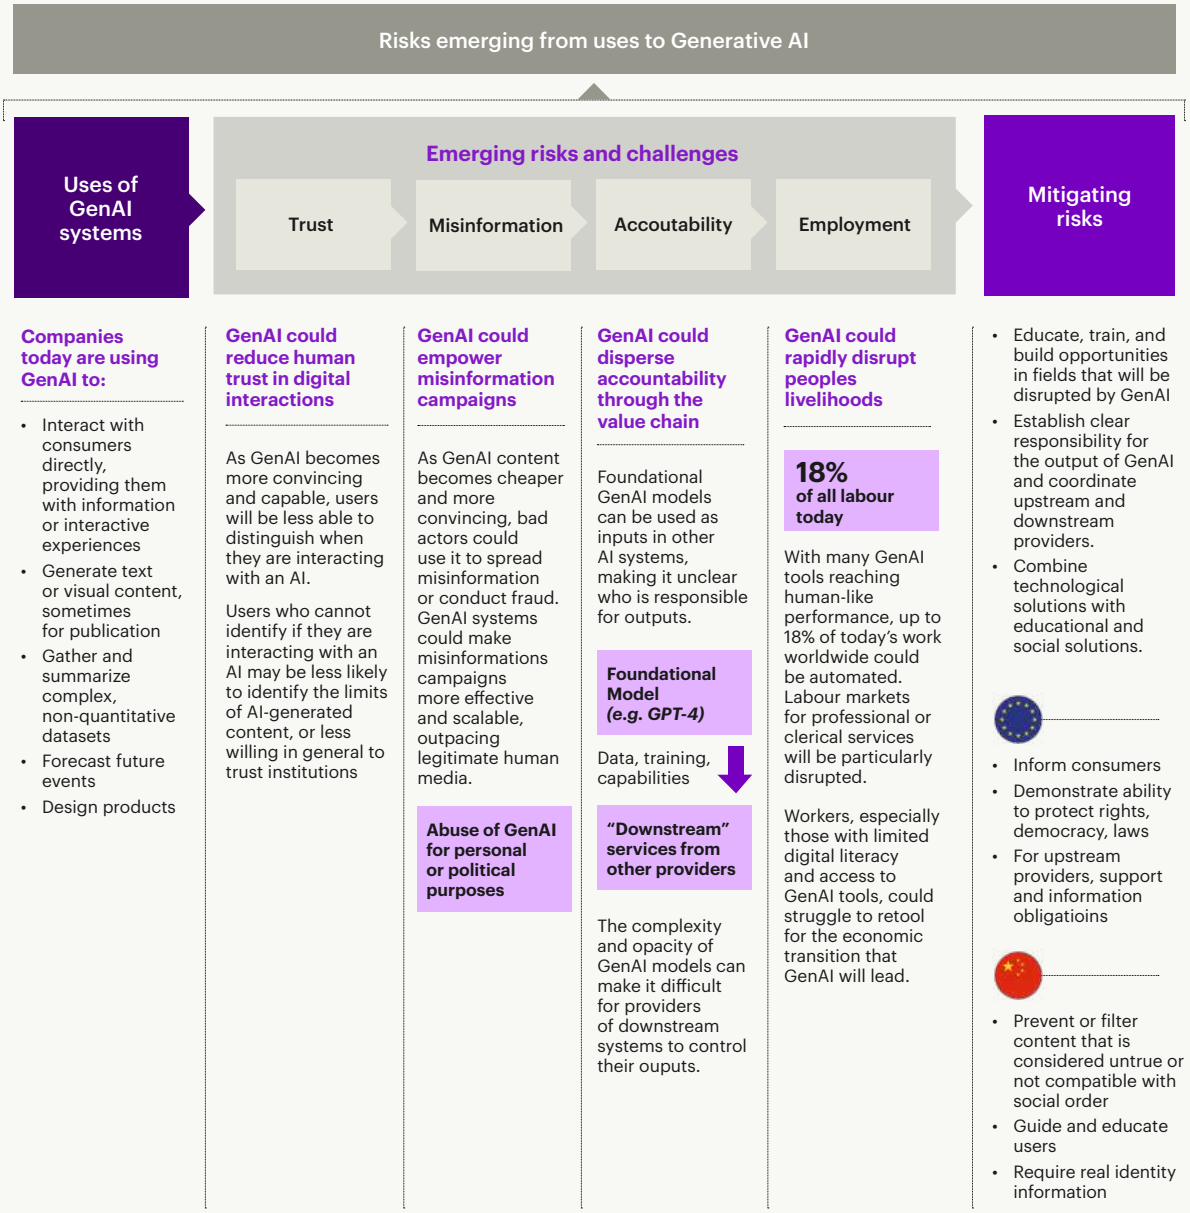

Output risks can be categorised into the following:

- Lack of trust: GenAI could lower trust in digital interactions—a particular concern for the financial sector given that its principal pillar is the trust that clients vest in it for their money and data. Linked to this, the centrality of robustness, stability, and trust to banks' operations reinforces the importance of dealing with confabulations.
- Misinformation: GenAI could empower misinformation campaigns as the generated content becomes more realistic or convincing.
- Unclear accountability: Current frameworks are largely designed for internal rather than external datasets. Dealing with the consumers of AI output is also crucial, especially in the trade-off between transparency, accuracy, and explainability, and in establishing data provenance. The concept of accountability becomes further challenged when outputs are also directly dependent on the inputs (questions) to the LLM models, and therefore users have an equal responsibility in the outputs by asking the right questions. Once again, having humans in or on the loop helps to ensure accountability.
- Employment disruption: GenAI could rapidly disrupt people's livelihoods as more roles get replaced or reduced. Participants agreed that upskilling existing staff as well as making changes to the educational curriculum in schools will be required for the next generation.
- Unfit-for purpose training data: Participants noted the importance, when generating training data for other models, of ensuring that knowledge distillation constitutes a perfect use case.

In the roundtable, it was emphasised that different use cases carry different risks.

Getting the balance right will not be straightforward. One approach could be to start with lower-risk use cases (e.g., those that drive internal efficiencies) before investing more heavily in customer-facing applications, which are typically higher risk. Deploying GenAI to do summarisation work, for example, and then using humans to make the final decision is far lower risk than deploying a chatbot that is directly exposed to consumers.

One participant also shared that their bank's code developers have numerous processes in place that are highly regulated and that work well. Using an LLM to write that code that is then subject to the same review process adds no incremental risk.

What is key from the above examples is to separate the space using a principles-based, risk-based approach and separate out low-risk models from others that have a greater risk. Doing so would accelerate the pace of responsible adoption of GenAI.

And, given the pace of change and disparity between the skills and resources available across the economy, monitoring and evaluating GenAI applications, along with a peer-review arrangement, are vital to ensure they operate safely and would help to address weakest-link issues across the lifecycle.

There are two types of technique available, each of which carries different risks. The Generative Model is one which creates new data points in an unsupervised machine-learning environment. This generative nature could create entirely new data from existing data. Its use raises questions around intellectual property, and because the mathematical model is intuitive, this could also give rise to transparency and explainability risks and undermine human accountability. On the other hand, the Discriminative Model is a conditional supervised machine-learning model and uses existing data sets to estimate outcomes with maximum likelihoods. The associated risk is that it could amplify existing biases.[9]

[9]*See Turing's Generative Models vs Discriminative Models: Generative models vs Discriminative models for Deep Learning. (turing.com)*

TAKEAWAY 4: HARMONISING AI GOVERNANCE FRAMEWORKS WILL BE CHALLENGING

Major jurisdictions have proposed or implemented principles that govern how firms deploy AI—with many of these overlapping—yet few have decided how to adapt those elements for GenAI or have fully considered whether what is currently in place for AI is sufficient for GenAI too.

When it comes to AI principles, the OECD's Principles on Trustworthy AI,[10] adopted in 2019, are the global standard and underpin AI governance models in key markets (see graphic). Here, Accenture shares a comparison of the select AI governance models across countries and their respective AI principles, and identifies key differences between them.

Comparison of select AI governance models

The basic principles underlying AI governance are strikingly similar, with differences emerging in regulation

	Voluntary best practices			Compulsory post-market	Compulsory hybrid control		Compulsory preemptive
	AUSTRALIA	SINGAPORE	UNITED STATES	CANADA	UNITED KINGDOM	EUROPEAN UNION	CHINA
OECD Principles on trustworthy AI	✓	✓	✓	✓	✓	✓	✓
Monitoring approach	Voluntary	Voluntary	Voluntary	Post-market monitoring	Preemptive and post-market	Preemptive and post-market	Preemptive monitoring
Scope	Applicable to all AI systems	Applicable to all AI systems	Applicable to all AI systems	Applicable to all AI systems	Sector-by-sector approach	Applicable to all AI systems	Sector-by-sector approach
Different governance for different risk levels				✓	✓	✓	
Morality rules							✓
Technology-agnostic	✓	✓	✓		✓		
Use case-agnostic	✓	✓	✓				
Scope of harm	Individuals, rights	Individuals, rights	Individuals, rights	Individuals	Individuals, rights, security	Individuals, rights, infrastructure	Individuals, rights, security, order, information
Prohibits some AI applications						✓	
Stage of completion	Promulgated	Promulgated	Promulgated	In discussion	Early guidance	In discussion	In force

[10] See: https://oecd.ai/en/ai-principles, and Humans keeping AI in check – emerging regulatory expectations in the financial sector (bis.org)

Comparison of select AI principles

The principles that countries have proposed for AI are broadly similar to one another and to the OECD standard

OECD	AUSTRALIA	SINGAPORE	UNITED STATES	CANADA	UNITED KINGDOM	EUROPEAN UNION	CHINA
1 Inclusive growth, sustainable development and well-being	1 Human, societal and environmental wellbeing	2 4 Fairness	4 Safe	1 4 Human oversight and monitoring	4 Safety, security and robustness	3 Human agency and oversight	4 Agile governance
2 Human-centred values and fairness	1 Human-centred values	1 2 Ethics	4 Secure & Resilient	3 Transparency	3 Appropriate transparency and explainability	4 Technical robustness and safety	* Open collaboration
3 Transparency and explainability	2 Fairness	5 Accountability	3 Explainable and interpretable	2 Fairness and equity	2 Fairness	4 Privacy and data governance	5 Shared responsibility
4 Robustness, security and safety	4 Privacy protection and security	3 Transparency	4 Privacy-enhanced	4 Safety	1 5 Accountability and governance	3 Transparency	3 4 Secure/safe and controllable
5 Accountability	4 Reliability and safety		2 Fair - with harmful bias managed	5 Accountability	3 5 Contestability and redress	2 Diversity, non-discrimination and fairness	4 Respect privacy
	3 Transparency and explainability		3 5 Accountable and transparent	4 Validity and robustness		1 Societal and environmental well-being	1 Inclusitivity and sharing
	3 Contestability		1 4 Valid and reliable			5 Accountability	2 Fairness and justice
	5 Accountability						1 Harmony and friendliness

Though the tenets of many of these initiatives are similar, a key question posed to the roundtable participants was whether a global set of harmonised standards could be feasible and practical, given regulators' varied risk appetites and socially accepted norms, and whether there are benefits to having a set of AI standards specific to the financial sector versus horizontal standards that apply across the economy.

The basis for a joint approach is clear: The technology is globally available and the issues regulators face are common across jurisdictions. As one regulatory discussant said, the benefits of GenAI can be realised more quickly "if firms can spend their time and money innovating and not navigating different jurisdictional approaches that needlessly conflict".

Nonetheless, a participant from a financial institution noted that the reality is that different jurisdictions may focus on different priorities, given varying levels of economic and technological maturities, and disparate social norms. Countries need to prioritise between the rights of individuals, the rights of corporations, and social harmony. The focus on each might vary depending upon the social and cultural norms of each society, and will have an impact on how regulations are formulated. In addition, the track record shows that regulators in the data space seldom come together with global standards.

Efforts are underway to tackle the issue on a multilateral basis. G7 leaders called for rules to govern AI, including GenAI, at their Japan meeting in May 2023. They established an intergovernmental forum, known as the Hiroshima AI Process,[11] to consider the development and adoption of technical standards to manage the problems created by GenAI tools and how the technology should be governed to be trustworthy. The forum will consider areas including disinformation, protection of IP, and AI governance, with a report scheduled for delivery to the G7 leaders by the end of 2023.[12]

Regulatory participants agreed there was a need to align at the principles level, even if not at a detailed technical standards level. There was also an opportunity for the regulators to agree on how to mitigate the risk of bad actors using this technology for their personal gain, and the collective safeguards needed to do so.

Finally, participants agreed that there is a pressing need to develop common definitions for the terminology relating to Responsible AI, such as "attribution", "transparency", "fairness", and "interpretability", so as to facilitate regulatory alignment.

Balancing regulation and innovation

It appears clear that the Responsible AI frameworks used by the financial services industry will need to be updated to keep up with the advancements brought about by GenAI. Applying the principle of proportionality to the uses of GenAI would go a long way to ensuring a good balance, said an advisor to a global financial institution. "The AI principles that we have are probably sufficient—we don't need to go overboard."

One challenge is that the EU's AI Act has already deemed that all LLMs are high risk regardless of the use case—which, as a first-mover in this regulatory space, will likely be used as a benchmark, shared another participant from a bank. "How do we work through that?" A regulatory participant added that detailed technical standards tend to take a long time to agree, which makes them less suitable to embed in regulation for this fast-developing space. Nonetheless, there are examples of a collaborative approach between regulators and innovators to establish and institutionalise high-level usage guidelines.

---

[11] *G7 officials to hold first meeting on AI regulation next week, Reuters (May 26, 2023).*
See: https://www.reuters.com/world/g7-officials-hold-first-meeting-ai-regulation-next-week-2023-05-26/
[12] *G7 officials to hold first meeting on AI regulation next week, Reuters (May 26, 2023).*
See: https://www.reuters.com/world/g7-officials-hold-first-meeting-ai-regulation-next-week-2023-05-26/

For example, the UK's Financial Conduct Authority (FCA) and the Bank of England (BoE) established an Artificial Intelligence Public-Private Forum (AIPFF) to explore the means to support a safe adoption of AI/ML technologies within financial services, and whether principles, guidance, regulation, and/or industry good practice could support this adoption.[13] The Monetary Authority of Singapore (MAS) has also worked with the financial services industry to develop its FEAT Principles, which promote fairness, ethics, accountability, and transparency in the industry's use of AI and data analytics.[14] Its Veritas Initiative—another collaboration with industry—provides a toolkit for financial institutions to evaluate their AI solutions against the FEAT Principles and integrate them into their internal risk governance.[15] Most recently, MAS has established Project MindForge, the purpose of which is "to examine the risks and opportunities of GenAI for the financial sector".[16]

Collaborative approaches like this ensure all participants are invested in coming up with high-level principles to provide a framework that allows the industry to innovate within guardrails.

If that could be replicated at a global scale, the results would be highly beneficial both defensively and offensively—particularly for global banks that must already deal with dozens of different regulations covering areas like data privacy, data-sharing, and records management, and not all of which are compatible. Navigating this array of complex regulations takes time and effort, which could be better spent solving bigger problems or innovating.

---

[13]The BOE and the FCA published the final report of the AIPPF in 2022: https://www.bankofengland.co.uk/-/media/boe/files/fintech/ai-public-private-forum-final-report.pdf?la=en&hash=F432B83794DDF3F580AC5A454F7DFF433D091AA5

[14]MAS introduces new FEAT Principles to promote responsible use of AI and data analytics, MAS (November 12, 2018). See: https://www.mas.gov.sg/news/media-releases/2018/mas-introduces-new-feat-principles-to-promote-responsible-use-of-ai-and-data-analytics

[14]Veritas Initiative, MAS (March 3, 2021). See: https://www.mas.gov.sg/schemes-and-initiatives/veritas

[15]Speech by Mr Heng Swee Keat, Deputy Prime Minister and Coordinating Minister for Economic Policies, at the Point Zero Forum on 26 June 2023, MAS (June 26, 2023). See: https://www.mas.gov.sg/news/speeches/2023/speech-by-dpm-heng-swee-keat-at-the-point-zero-forum

TAKEAWAY 5: CREATING AN ENVIRONMENT THAT FACILITATES THE RESPONSIBLE DEPLOYMENT OF GENAI INVOLVES MORE THAN JUST THE TECHNICAL EXPERTS

One participant from the financial sector likened the development of GenAI in a responsible manner to raising a child with the right values. That will prove challenging—particularly when situations are complex and trade-offs are needed. Additionally, the tools being developed now could be around far longer than a single human lifespan, which raises yet more considerations.

For firms in the financial services industry, founded as they are on trust, GenAI that is imbued with the wrong values could erode this trust. For an industry that is of global systemic importance, the stakes are even higher.

As the graphic makes clear, collaboration is crucial when countering the risks that GenAI can bring.

Responding to the risks of GenAI is a collaborative effort

Public Institutions

- **Collaborating** with companies, innovators, and citizens to chart a path forward for GenAI
- Building a fair, effective **regulatory environment**
- Preparing workers and citizens with the **tools to thrive** in a GenAI-powered world
- Supporting minority languages, SMEs, and other potentially **excluded groups**
- Supporting continued **interdisciplinary research and innovation** into GenAI and its risks

Inputs that are unbiased, legally obtained, transparent, and inclusive

Outputs that are identifiably AI, resistant to bad actors, clearly accountable, and socially responsible

Companies using GenAI

- Using GenAI in a **transparent** way, with clear accountability for outputs
- Using GenAI for its intended purpose
- **Educating** customers and employees on how to engage responsibly with GenAI
- Proactively identifying and **preventing abuse** of their GenAI products and services
- **Engaging** with innovators, regulators, content creators, and civil society

Innovators building GenAI

- Designing GenAI systems that are **fundamentally responsible,** unbiased, and respectful or rights
- Establishing **clear responsibility** for the outputs of GenAI in their value chain
- **Supporting downstream implementers** of their GenAI products and services
- Rigorously screening and maintaining **data sources** for fairness and copyright
- Designing GenAI to be **broadly inclusive**

Ultimately, ensuring an ethical approach to the use of GenAI in the financial services industry will depend upon fostering a culture that attracts responsible and accountable usage of the technology. This requires a shared responsibility among everyone in the organisation, so that all individuals are properly equipped to contemplate the appropriate applications for GenAI.

Concluding remarks

Responsible AI is not an isolated goal but an ongoing journey. By acting collectively, all parties can create solutions that serve their organisations and society in an ethical, unbiased, and beneficial manner.

The key factors to ensuring a balance between innovation and managing the risks involved are:

- Clarifying the application of existing governance models and principles for AI on GenAI use cases.
- Being intentional about building up the skills needed to understand and deploy AI in a responsible manner within a firm
- Collaborating between regulators, industry, and other interested parties to determine the appropriate guardrails in a fast-evolving space.
- Exploring the possibility of harmonising global principles that would ensure a level playing field for the sector.

ACKNOWLEDGEMENT

Discussants

Andrew Walters
Senior Fintech Analyst,
Bank of England

Calissa Aldridge
Senior Executive Leader – Market Supervision
Australian Securities & Investments Commission

Gary Kazantsev
Head of Quant Technology
Bloomberg

Jeffery Yong
Principal Advisor
Bank for International Settlements

Dr Li Xuchun
Head of AI Development Office
Monetary Authority of Singapore

Musa Parmaksiz
Group Data Science and Analytics Lead
UBS

Pier Luigi Culazzo
Chief Data Officer
Standard Chartered Bank

Dr Prag Sharma
Global Head of Artificial Intelligence
Citigroup

Sameer Gupta
Managing Director, Chief Analytics Officer
DBS Bank

Chair

Marcus Lim
Assistant Managing Director (Banking & Insurance)
Monetary Authority of Singapore

Moderator

Irene Liu
Managing Director, Digital Risk and Compliance
Accenture

About Accenture

Accenture is a global professional services company with leading capabilities in digital, cloud and security. Combining unmatched experience and specialized skills across more than

40 industries, we offer Strategy and Consulting, Technology and Operations services and Accenture Song — all powered by the world's largest network of Advanced Technology and Intelligent Operations centers. Our 721,000 people deliver on the promise of technology and human ingenuity every day, serving clients in more than 120 countries. We embrace the power of change to create value and shared success for our clients, people, shareholders, partners and communities. Visit us at www.accenture.com.

About Elevandi

Elevandi is set up by the Monetary Authority of Singapore (MAS) to foster an open dialogue between the public and private sectors to advance FinTech in the digital economy.

We work closely with governments, founders, investors, and corporate leaders to drive collaboration, education, and new sources of value at the industry and national levels.

Our initiatives have convened over 350,000 people since 2016 to drive the growth of FinTech through events, closed-door roundtables, investor programmes, educational initiatives, and research. Our flagship product is the Singapore FinTech Festival alongside fast-rising platforms, including the World FinTech Festival and Point Zero Forum.