



# Financial Services in the Era of Generative AI

## Facilitating Responsible Adoption

April 2025

# Contents

	Pages
<b>Foreword</b>	<b>2</b>
<b>Acknowledgements</b>	<b>3</b>
<b>Executive Summary</b>	<b>4</b>
<b>Chapter 1      Generative Artificial Intelligence: An Overview</b>	<b>6</b>
Evolution, key technical breakthroughs, and implications for the financial services industry	
<b>Chapter 2      Implications for Financial Services Industry Risk Management and Talent Strategies</b>	<b>18</b>
The importance of risk management and talent	
<b>Chapter 3      Insights from Market Participants in Hong Kong</b>	<b>31</b>
GenA.I. adoption in the financial services industry	
<b>Chapter 4      The Role of Regulation</b>	<b>45</b>
The international and local experience	
<b>Chapter 5      Considerations For Fostering Responsible GenA.I. adoption in Hong Kong</b>	<b>55</b>
The scope for further facilitation and cooperation	
<b>Conclusions</b>	<b>63</b>
<b>Appendix A:    Background of the GenA.I. Survey</b>	<b>64</b>
<b>Appendix B:    Glossary of Technical Terms</b>	<b>66</b>
<b>Appendix C:    References</b>	<b>70</b>

# Foreword

Generative Artificial Intelligence (GenA.I.), a rapidly evolving technology in its early stages of development and adoption, represents the latest evolution in Artificial Intelligence (AI). Globally and in Hong Kong, financial institutions are swiftly exploring GenA.I.'s capabilities, with many promising use cases emerging.

In the financial services industry, GenA.I. applications have the potential to lead to more curated customer experience, and more efficient ways of processing and utilising digital information. However, the adoption of GenA.I. could also give rise to new risks and challenges. This suggests that amidst accelerated GenA.I. innovation, its adoption requires a critical focus on safety, trust and integrity. In Hong Kong, authorities have begun to update regulatory guidelines and to launch initiatives in support of responsible GenA.I. adoption and innovation.

Motivated by recent developments, this report provides a comprehensive overview of the evolution of GenA.I. and its broader implications for both the financial services industry and financial regulators. The report draws on the findings from a survey and interviews commissioned by the Hong Kong Institute for Monetary and Financial Research, which gathered the views of market participants on the current state of GenA.I. adoption among local financial institutions, the expected trajectory of GenA.I. development in Hong Kong, and the strategies employed for risk management and talent development.

The report finds that the adoption of GenA.I. is progressing steadily across the financial services industry in Hong Kong. However, there are challenges hindering adoption, including concerns regarding model accuracy, data privacy and security, as well as constraints related to resources and talent. The emergence of less resource-intensive models and maturing technology, coupled with regulatory engagement, is likely to contribute to the broadening of GenA.I. adoption over time. Based on these findings, the report outlines some considerations aimed at facilitating responsible GenA.I. adoption by the financial services industry in Hong Kong.

We hope the findings of this report can help inform best practices for addressing GenA.I. adoption challenges in the financial services industry, and contribute to discussions on responsible innovation and adoption, as well as industry-wide capacity building. Looking ahead, with maturing GenA.I. technology and evolving regulation anticipated to support a broadening of GenA.I.'s application to a wider spectrum of activities, further research may also be warranted to understand the potential for GenA.I. to support RegTech, SupTech, and the broader policy setting.

## **Mr Enoch Fung**

Chief Executive Officer  
Hong Kong Academy of Finance

Executive Director  
Hong Kong Institute for Monetary and Financial Research

# Acknowledgements



We thank the Insurance Authority, Mandatory Provident Fund Schemes Authority, Securities and Futures Commission, as well as the Banking Conduct Department, the Banking Policy Department, the Banking Supervision Department and the Digital Finance Division of the Hong Kong Monetary Authority for their valuable comments and suggestions. We are also grateful for our collaboration with Ernst & Young Advisory Services Limited in designing and administering a survey, titled *Financial Services in the Era of Generative Artificial Intelligence: Opportunities and Risk Management*, from October 2024 to January 2025, and in conducting interviews with various market participants, including banks, insurers, and asset and wealth managers, as well as GenA.I. service providers. Finally, we thank the Hong Kong Institute for Monetary and Financial Research Council of Advisers for Applied Research for their continuing support and guidance related to the Institute's research activities.



# Executive Summary

**Generative Artificial Intelligence (GenA.I.), a rapidly evolving technology in its early stages of development and adoption, represents the latest evolution in Artificial Intelligence (AI).** GenA.I. models can learn from the patterns and structures of their training data to generate outputs with similar characteristics, be it text, images, audio, or video. GenA.I.'s key attributes of accessibility, versatility, and adaptability thus broadens the potential to automate, innovate, and enhance productivity in the financial services industry.

**The adoption of GenA.I. is progressing steadily across the financial services industry in Hong Kong.** 75% of the surveyed financial institutions have already implemented at least one GenA.I. use case, or are currently piloting and designing use cases, and exploring potential investment areas. This ratio is expected to increase to 87% within the next three to five years.

**GenA.I. adoption has been somewhat higher among the larger surveyed financial institutions.** Among surveyed firms, 83% of large firms have rolled out at least one GenA.I. use case or are taking steps towards adoption, compared to 63% of small firms. Larger firms were typically more advanced in their GenA.I. adoption and general preparedness, while smaller firms with comparatively less resources faced greater adoption hurdles.

**The primary implementations of GenA.I. in financial services remain largely internal and non-customer facing.** 75% of the surveyed financial institutions viewed GenA.I. as a tool to enhance productivity and operational efficiency, followed by 53% who viewed GenA.I. as empowerment for innovation and decision-making. The most common GenA.I. use cases are virtual assistants for employees, with GenA.I. use in complex, higher-risk, and external customer-facing applications dependent on further improvements in the accuracy of the technology.

**There are a number of risk management challenges hindering adoption, including concerns regarding model accuracy, data privacy and security, as well as constraints related to resources and talent.** When adopting GenA.I., financial institutions considered model performance and accuracy (highlighted by 95% of surveyed firms), model transparency and explainability (65%), and data privacy and security (64%) as the top three risk-management considerations.

**To strengthen risk management, financial institutions in Hong Kong have made solid first steps towards responsible GenA.I. adoption and development, supported by updated regulatory guidelines.** There is a clear prioritisation of transparency and accountability in GenA.I. tools, alongside a strong emphasis on data protection and safeguarding customer information. Ongoing data monitoring and governance of model outputs, and enhancing data quality controls and checks, were common priority areas identified for improvement. GenA.I.-related cybersecurity awareness training and regular security assessments to identify vulnerabilities are also being introduced.

A 'human-in-the-loop' approach is also considered necessary to ensure proper risk management and controls, especially at this stage of GenA.I. technology maturity. Consistent with this approach, about 80% of the surveyed financial institutions identified technical skills required in the use and development of GenA.I. as some of the top skills gaps faced by the industry. 60% of survey respondents also highlighted compliance skills as a key skills gap in supporting GenA.I. initiatives.

**Reskilling and upskilling appears to be a core part of financial institutions' talent strategy.** To help bridge skills gaps, the surveyed financial institutions are opting for a combination of

upskilling existing employees, hiring new talent, and establishing external partnerships, especially with GenA.I. service providers. There is also an emphasis on continuous learning and innovation and on upskilling existing data and technology teams in areas of data science, cloud computing and business analytics.

**The emergence of less resource-intensive models and maturing technology, coupled with regulatory engagement, are likely to contribute to the broadening of GenA.I. adoption over time.** The recent emergence of less resource-intensive models, such as DeepSeek-R1, is challenging the prevalent view that scaling GenA.I. requires vast computing power and investment. New approaches to language modelling are also improving model accuracy and general performance. The trajectory of these developments should support a broadening of GenA.I. use over time.

**In recognition of the transformative potential of GenA.I., and to facilitate responsible adoption, the regulatory and policy landscape of AI regulation has been an evolving process across jurisdictions worldwide.** Regulators globally have begun re-examining their regulatory frameworks governing AI development and adoption. Multilateral discussions among global policymakers are also underway to understand the potential implications of GenA.I. adoption for financial stability and market integrity.

Although jurisdiction-specific regulatory frameworks are currently diverse in terms of the degree of codification, industry coverage, and sandboxing and facilitation, these continue to be guided by an end goal of facilitating responsible adoption, and balancing between innovation and safety. Achieving greater cross-jurisdiction harmonisation over time can help reduce the costs of compliance

for financial institutions, especially those with a substantial global footprint, as well as help prevent regulatory arbitrage.

**Based on these findings, the report outlines some considerations for facilitating responsible GenA.I. adoption by the financial services industry in Hong Kong.** The Hong Kong government and financial authorities have been active in undertaking a multi-pronged approach in supporting responsible GenA.I. adoption and development, with the release of policy statements, regulatory circulars and guidelines surrounding the use of GenA.I., and the launch of facilitation measures such as the Hong Kong Monetary Authority (HKMA) and Cyberport GenA.I. Sandbox for banks and the Cyberport AI Supercomputing Centre.

In the near-term, continued adjustments of financial institutions' risk management frameworks to align with best practices can foster further adoption. Expanding the scope of facilitation measures, as well as supporting technology and implementation infrastructure, can also facilitate responsible GenA.I. adoption and spur innovation and development.

In the medium-term, survey and interviews suggest greater collaboration among regulators, industry players and developers, cross-jurisdiction regulatory engagement, and sustained long-term investment in digital infrastructure are crucial as GenA.I. adoption broadens.

The surveyed financial institutions viewed the development of more advanced use cases and applications as the top area where industry-regulator-developer cooperative gains are most likely to occur, followed by talent development, enhanced public understanding, and infrastructure.

# Chapter 1

## Generative Artificial Intelligence: An Overview

### HIGHLIGHTS:

- GenA.I.'s key attributes of accessibility, versatility, and adaptability broadens the potential to automate, innovate, and enhance productivity in the financial services industry. By automating labour-intensive tasks and augmenting workers in cognitive tasks, GenA.I. can also free up human capital for strategic planning.
- However, similar to other technological advancements in the early stages of development and adoption, GenA.I. also introduces potential risks that pose challenges for financial institutions as prospective end-users, and for financial authorities through the implications for financial stability, and consumer and investor protection.
- Responsible GenA.I. adoption thus requires financial institutions to consider the robustness of their risk management and talent strategies. Financial regulation is also anticipated to be a dynamic process, owing to the nascent nature of the technology, and the speed and breadth of the integration of GenA.I. into the financial services industry.


Generative Artificial Intelligence (GenA.I.), a rapidly evolving technology in its early stages of development and adoption, represents the latest evolution in Artificial Intelligence (AI). Its development has been underpinned by advancements in computing power and efficiency, breakthroughs in deep machine learning architecture, and the availability of large datasets that have facilitated the training of complex foundation models. GenA.I. models can learn from the patterns and structures of their training data to 'generate' outputs with similar characteristics, be it text, images, audio, or video. Globally and in Hong Kong, financial institutions are swiftly exploring GenA.I.'s capabilities, with many promising use cases emerging<sup>1</sup>. This chapter provides an overview of GenA.I., including the key milestones and technical breakthroughs, its key attributes and risks features, and its broader implications for both the financial services industry and financial authorities.

## 1.1 A BRIEF HISTORY OF AI

### The birth of AI

AI is intelligence exhibited by machines and the ability of machines to perform tasks commonly associated with human intelligence. The invention of the programmable electronic computer at the beginning of the 20<sup>th</sup> century first prompted serious consideration of AI and its possibilities. Since then, AI research has progressed through a number of distinct developmental paradigms<sup>2</sup>, in a reflection of both shifts in AI ideology and the gradual advancement of the technology (Figure 1.1).

**Figure 1.1: Major stages in AI development**



Source: HKIMR staff compilation.

Initial AI research was focused on how to encode logic reasoning into machines, and the design of AI systems that follow specific rules and perform well-defined tasks<sup>3</sup>. One of the first natural language processing (NLP) chatbots, ELIZA, was developed during this period. Although this period

culminated in the arrival of expert systems that emulate the knowledge and reasoning abilities of human experts in specific domains requiring narrow but deep knowledge, the difficulty of such rule-based AI systems in handling ambiguous and noisy data<sup>4</sup> ultimately limited any real-world applications and adoption.

<sup>1</sup> McKinsey & Company (2023).

<sup>2</sup> European Commission (2020).

<sup>3</sup> Two attendees of the Dartmouth Summer Research Project on Artificial Intelligence, Herbert Simon and Allen Newell, proposed more specifically that human minds and modern digital computers were 'species of the same genus,' namely symbolic information processing systems – both take symbolic information as input, manipulate it according to a set of formal rules, and in so doing can solve problems, formulate judgements, and make decisions.

<sup>4</sup> The translation of human knowledge into logical rules to solve ever more complex real-world problems proved computationally expensive and impractical, contributing to 'knowledge bottlenecks'.



## Machine learning

From the 1990s onwards, steady advances in computational power to train and run AI models and progress in machine learning (ML) algorithms that leverage statistical techniques to process information began to support better predictions and decision-making based on historical data. The use of backpropagation, an iterative algorithm that helps to minimise the cost function by determining which weights and biases should be adjusted by moving down the gradient of the error, also experienced a resurgence over this period. By enabling more efficient training, backpropagation helped to inspire the renaissance and interest in artificial neural networks in AI research. As a result, ML definitively shifted from a rule-based knowledge-driven approach to a data-driven approach.


### Artificial neural networks (ANNs)

Artificial neural networks<sup>5</sup>, with their topology of interconnected functions ('neural nodes'), aim to simulate the way neurons in the human brain decipher information and signal to one another, and are trained to recognise patterns and

infer rules through processing large amounts of input data across successive and 'hidden' neural network layers (Figure 1.2). The integration of ML algorithms into this type of AI model architecture enabled the model to learn mathematically complex relationships between data points through numerous iterations, and has greatly expanded the potential for AI in real-world applications and adoption.

In each layer, the weights of the neural nodes are adjusted through the application of statistical methods to prioritise those input data from preceding nodes that contribute to improving the desired output, until the best fitting model that minimises the cumulative error from all the training data points is found. Once trained, neural networks are further calibrated and validated to minimise error on a previously 'unseen' segment of the input data set (the 'test data'), to improve their robustness to noise. The trained and validated neural network model can then be used to interpret new inputs and make decisions or predictions. The more layers and interconnected neurons per layer a neural network has, the larger and more diverse the training data and computing power it requires.

**Figure 1.2: An example schematic of the working of artificial neural networks**



Source: HKIMR staff compilation.

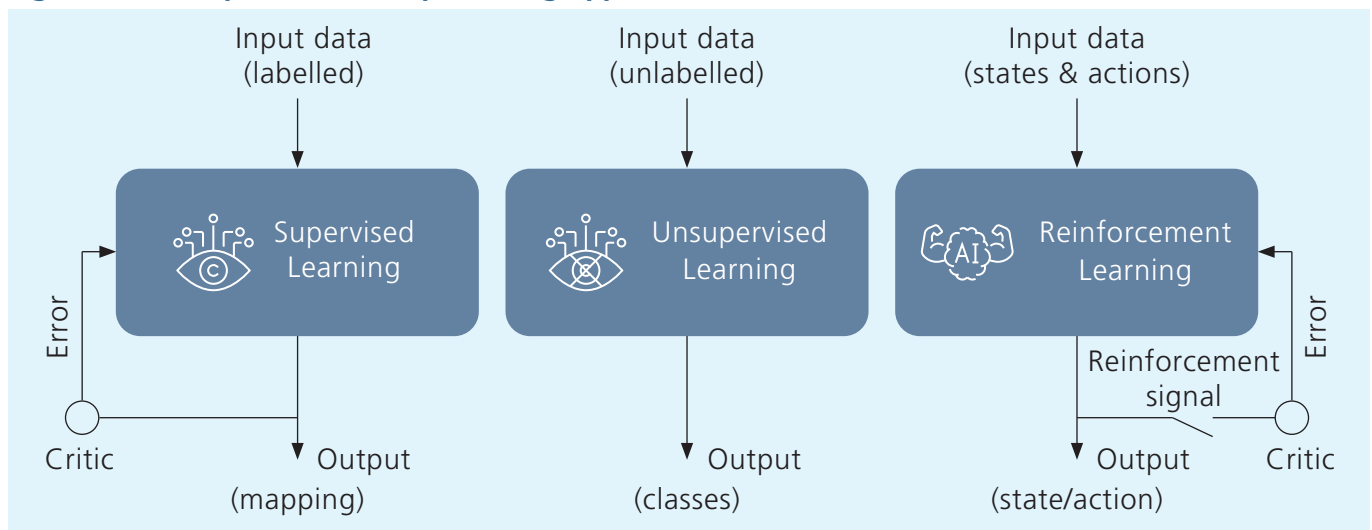
<sup>5</sup> Artificial neural networks were first proposed in 1943 by Warren McCulloch and Walter Pitts from the University of Chicago, with the publication of the first mathematical modelling of a neural network and how it might perform simple logical functions. The first trainable neural network, the Perceptron, was demonstrated by Cornell University psychologist Frank Rosenblatt in 1957.

## Deep learning

The ability of neural networks to take on additional layers has improved over time through new training techniques<sup>6</sup>. However, it was the integration of high-performance graphics processing units (GPUs) and thus more efficient parallel processing<sup>7</sup>, as well as additional access to 'big data' sets through cloud services, that really facilitated the training and running of multi-layered ('deep') neural networks. Greater computational power also supported more sophisticated learning approaches, expanding the type of data used to train the algorithms, while reducing the degree

of human intervention and increasing in-model feedback (Figure 1.3). In recent years, deep learning has come to underpin the AI systems behind nearly all high-performing predictive and prescriptive AI applications, from forecasting and fraud detection to optimisation and logistics. Today, large-scale AI systems can have hundreds or even thousands of hidden neural network layers that mimic the complex decision-making powers of the human brain, and excel at the type of real-world tasks that earlier rule-based AI systems struggled with.

**Figure 1.3: Comparison of deep learning approaches**



Sources: IBM and HKIMR staff compilation.

## The era of GenA.I.

Evolution in GenA.I. has mirrored the evolution in broader AI research. There was early GenA.I. experimentation with statistical models for speech recognition in the 1950s and some success in language modelling tasks by the late 1980s<sup>8</sup>. However, GenA.I. applications for a long time struggled to move beyond simple detection to the

efficient generation of high-quality and diverse new content. This reflects the even higher demand of GenA.I. than traditional AI for large and diverse training datasets that can adequately capture the intricacies and variations present in the real world<sup>9</sup>, and for computational power that can train and operate the complex integrated algorithms of large-scale foundation models<sup>10</sup>.

<sup>6</sup> The invention of the 'greedy layer-wise pre-training' technique in 2006 allowed each layer of a neural network to be trained individually, thereby reducing the aggregate amount of training time.

<sup>7</sup> GPUs, originally designed for use in computer games, pack thousands of processing cores onto a single chip and their architecture is similar to that of neural networks. High-performance GPUs have a parallel architecture that allow parallel processing, that is, running two or more central processing units (CPUs) to handle separate parts of an overall task, reducing the amount of time to run a programme.

<sup>8</sup> The introduction of recurrent neural networks in the late 1980s and long short-term memory networks in 1997 enhanced the ability of AI systems to process sequential data. Between the 1980s and 2000s, the shift towards statistical models and ML begin to lead to more practical NLP applications like translation services, search engines, and voice-activated assistants.

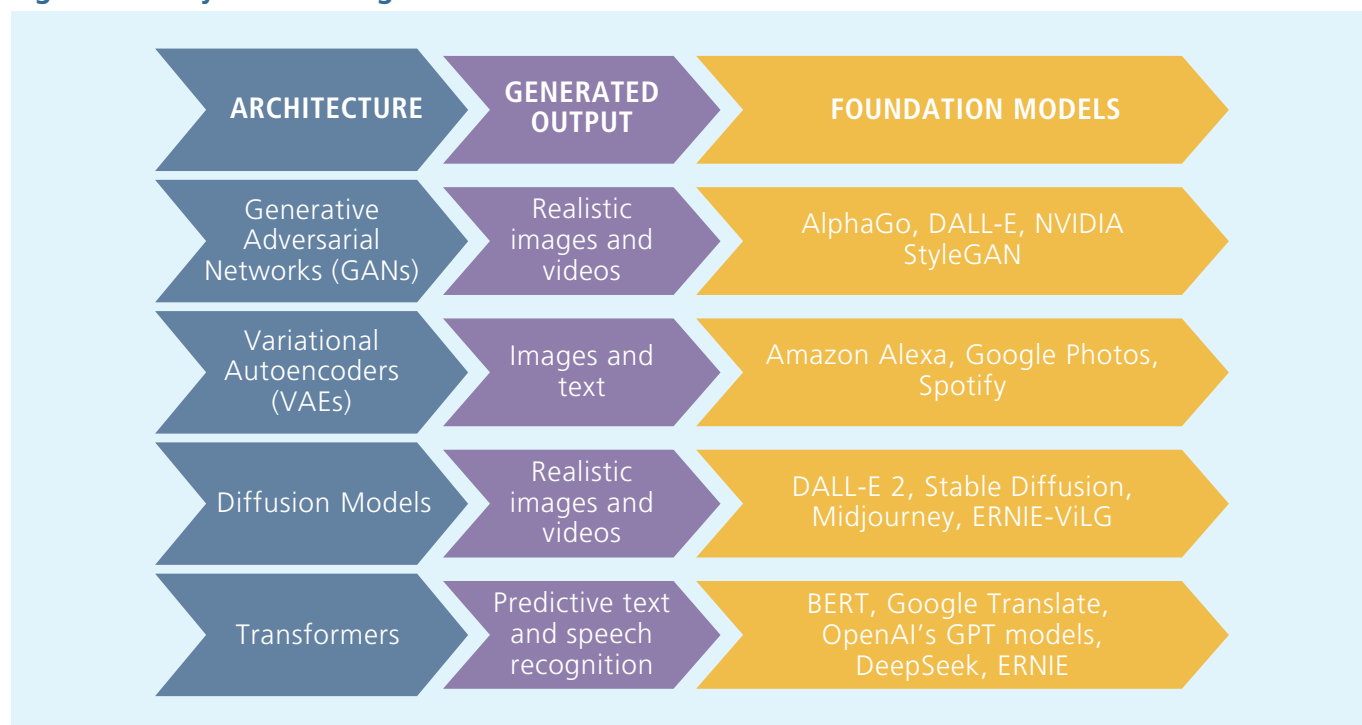
<sup>9</sup> Alzubaidi, L. et al. (2023).

<sup>10</sup> Hu, Q. et al. (2024).

The accelerated progression of deep learning and probabilistic modelling techniques in the early 2010s thus also led to a jump in the productivity of GenA.I. models, through facilitating a series of breakthroughs in GenA.I. model architecture

(Figure 1.4). New transfer learning techniques<sup>11</sup> and the arrival of multimodal models (e.g. text-to-image, text-to-music) further helped to drive efficiency and diversity in new content generation.

**Figure 1.4: Key breakthroughs in GenA.I. architecture: use cases and models**



Source: HKIMR staff compilation.

In the visual domain, generative adversarial networks, variational autoencoders, and diffusion-based models have made it possible to create and generate novel images, audio, and video content. However, it was the introduction of the transformer architecture that has been viewed by some AI researchers as paradigm-shifting<sup>12</sup> for NLP, an area of long-standing AI research interest that aims to automate the analysis and synthesis of natural languages and speech. Transformer-based large language models (LLMs) can process and track relationships in sequential data and undertake parallel processing, thus overcoming

limitations of earlier models in handling long text sequences while reducing computational demands<sup>13</sup>. Today, prominent transformer-based LLMs – such as Google's BERT, OpenAI's GPT models, and the more recent DeepSeek models – can produce coherent and contextually relevant text as well as generate conversation on the basis of ordinary language prompts.

The success of early GenA.I. breakthroughs has given impetus to greater competition among both established technology firms and new entrants globally. For instance, in addition to a focus on

<sup>11</sup> Under transfer learning, neural network models pre-trained for one task can be fine-tuned for a new, similar task by re-training just a subset of their layers with a smaller number of data samples, reducing the demand for data and computational resources.

<sup>12</sup> Bommasani, R. et al. (2022).

<sup>13</sup> Vaswani, A. et al. (2023).

AI developments among US technology firms, Chinese technology firms and academia have also been active on the AI front. In 2017, Baidu, China's leading online search engine provider, in partnership with leading Chinese universities and research institutions, established Baidu Research as the country's first national laboratory aimed at advancing deep learning technologies<sup>14</sup>. Other major technology players, such as Alibaba and Tencent, have also been undertaking AI research over the past decade with the establishment of the DAMO Academy and the Tencent AI Lab respectively<sup>15</sup>. As a result of broad-based efforts to strengthen AI development, China accounted for just over 60% of global AI patent origins in 2022<sup>16</sup>. More recently, the emergence of DeepSeek-R1 has challenged the prevalent view that scaling GenA.I. requires vast computing power and investment. New approaches to language modelling are also improving model accuracy and general performance. These developments are further supported by broad-based and long-term facilitation measures, such as China's AI Plus Initiative that was first announced in 2024, which aims to advance research and application of AI, as well as to combine digital technologies with China's market strengths.

In Hong Kong, locally developed GenA.I. tools are also spurring innovation in the local AI ecosystem. In early 2025, the Hong Kong Generative AI Research and Development Center, a government-backed joint-university collaborative venture led by the Hong Kong University of Science and Technology, launched the HKGAI V1<sup>17</sup>. The LLM, trained with the assistance of Chinese AI firm DeepSeek, can support Cantonese, Putonghua, and English, thus offering substantial applicability to local linguistic and cultural contexts as well as

support to local start-ups and enterprises as an open-source AI tool for business applications.

## 1.2 GenA.I. OPPORTUNITIES FOR THE FINANCIAL SERVICES INDUSTRY

GenA.I.'s key attributes of accessibility, versatility, and adaptability are raising prospects of more intuitive human-computer interactions, as well as reshaping how digital information can be processed, utilised, and presented with greater efficiency (Figure 1.5). This broadens the potential to automate, innovate, and enhance productivity in the financial services industry. The appropriate use of GenA.I. is also anticipated to boost worker productivity, by further automating labour-intensive tasks and augmenting workers in cognitive tasks<sup>18</sup>, and by freeing up human capital for strategic planning and informed decision-making. In some business operations with more complex and flexible cognitive requirements, hybrid human-AI teams may also be deployed to further enhance productivity<sup>19</sup>.

In the front-office, GenA.I.'s accessibility offers the possibility of enhanced and personalised customer experience. GenA.I.-enabled chatbots skilled at contextual issues can better assist customers through complex business activation and troubleshooting processes. Staff can also leverage GenA.I.-enabled virtual assistants to access institutional knowledge through conversation-like prompts that can contribute to more effective interactions with customers. Meanwhile, GenA.I.'s added versatility and adaptability can facilitate a broad range of highly technical tasks with the potential for improved outcomes. For example,

<sup>14</sup> *South China Morning Post* (2017).

<sup>15</sup> *MIT Technology Review* (2017).

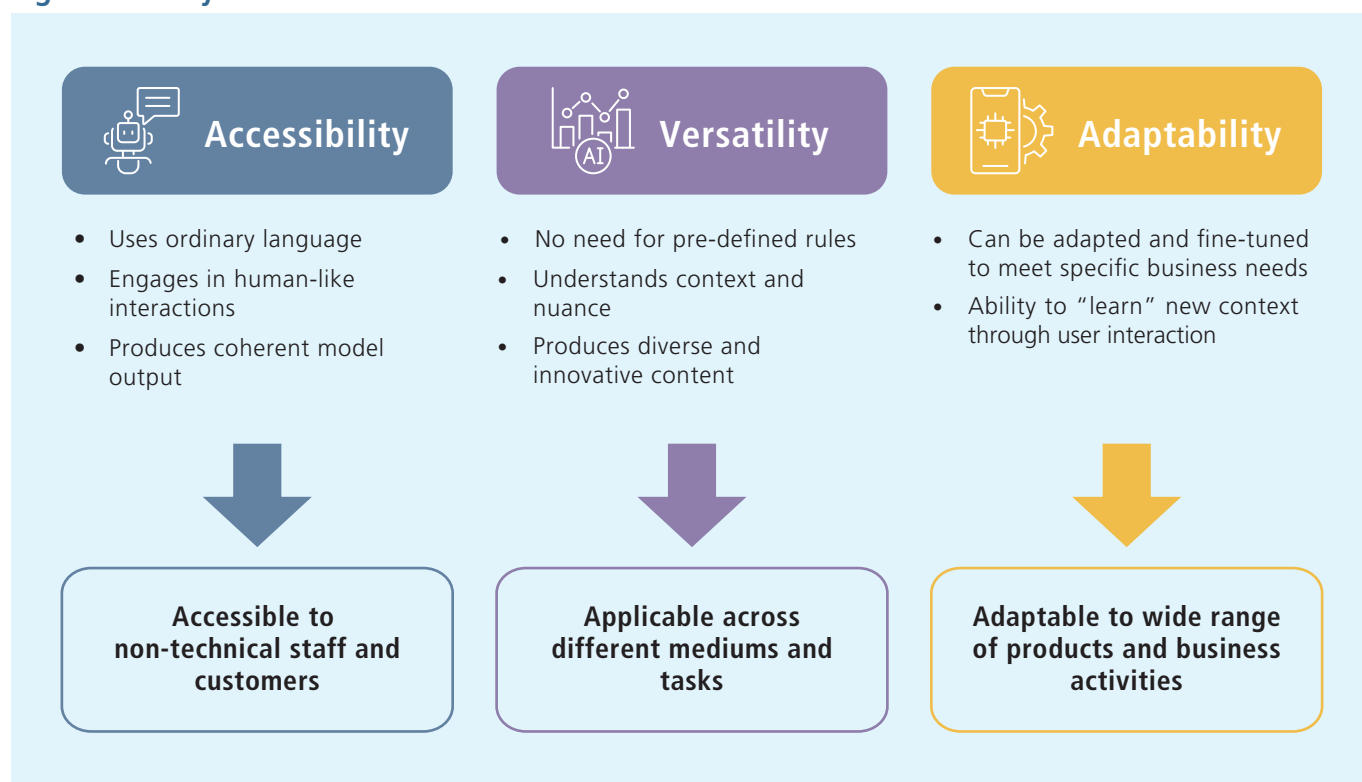
<sup>16</sup> *Stanford University* (2024).

<sup>17</sup> *The Standard* (2025).

<sup>18</sup> Yang, C. H. (2022).

<sup>19</sup> Fuchs, A., Passarella, A., & Conti, M. (2024).

Figure 1.5: Key attributes of GenA.I.



Source: HKIMR staff compilation.

human-AI hybrid trading teams leveraging GenA.I. can uncover strategies that were previously computationally difficult to process. GenA.I. also raises the possibility of customised credit and insurance products based on individual customer data on a scale that human-only product teams would find difficult to achieve.

In the middle and back office, GenA.I.’s versatility and adaptability can improve the efficiency of compliance processes and financial reporting, as well as support more rapid real-time fraud detection and prevention. The automation of some claims processing and other routine and repetitive tasks can also free staff for more complex risk management and compliance issues. Moreover, GenA.I., in supporting informed decision-making as an advanced knowledge repository and search engine<sup>20</sup>, allows for more accessible flow of information within an organisation.

### 1.3 GenA.I. RISKS AND CHALLENGES

However, similar to other advancements in technology in the early stages of development and adoption, GenA.I. also comes with a number of risks and challenges<sup>21</sup> (Figure 1.6). First, GenA.I. solutions have been predominantly known to be **computationally expensive and resource intensive** to develop and to adopt, notwithstanding recent competition and progress towards more efficient models. In the last few years, only a limited set of leading AI developers have had the requisite resources and talent to build cutting-edge GenA.I. models. Likewise, high upfront resource commitments can present significant barriers to entry for some financial institutions, and may limit GenA.I. adoption to larger and better resourced firms. This can result in a concentration of GenA.I. developers and service providers and a market

<sup>20</sup> Morgan Stanley (2023).

<sup>21</sup> International Monetary Fund (2023).



dominance of top foundation models, as well as an unequitable adoption of GenA.I. across the financial services industry.

Second, the **dependency on large training datasets** often necessitates the scraping of unlabelled public sources of data, such as unstructured text (e.g. Wikipedia), large image repositories, or audio recordings. The growing volume of potentially personal and proprietary data involved in training can also increase the risks around data privacy, misuse and leaks, as well as give rise to intellectual property challenges<sup>22</sup>. By comparison, the training data sets used in traditional AI models tend to be significantly smaller, labelled for supervised learning, and curated and pre-processed to fit a specific task.

Third, the use of statistical methods to generate the most plausible output through the optimisation of a large number of parameters can result in **hallucination and bias**<sup>23</sup>, especially if the training data are incomplete and have in-built bias to begin with. Indeed, there is already ample evidence of LLMs generating a response that is either factually

incorrect, nonsensical, or disconnected from the input prompt, or in extreme cases even malign<sup>24</sup>. This is because LLMs are unable to distinguish between what is linguistically probable and what is factually correct, with the implication that hallucination and bias may be architectural features rather than anomalies of some GenA.I. models<sup>25</sup>. It also leaves open the question of whether these problems merely reflect the limits posed by the size of the training data and the number of model parameters, or if they indicate more fundamental limits to knowledge that is acquired through language alone<sup>26</sup>.

Fourth, **complex and opaque model architectures** can make the results of GenA.I. models difficult to explain, as the multiple neural networks and parameters used for generating output complicate their interpretability relative to linear models (the 'black box effect'). The use of large and non-traditional and unstructured data sets can exacerbate these issues, and – when the input signals are incorrect (or 'poisoned') – can make detection of the appropriateness of the results difficult.

<sup>22</sup> Harvard Business Review (2023).

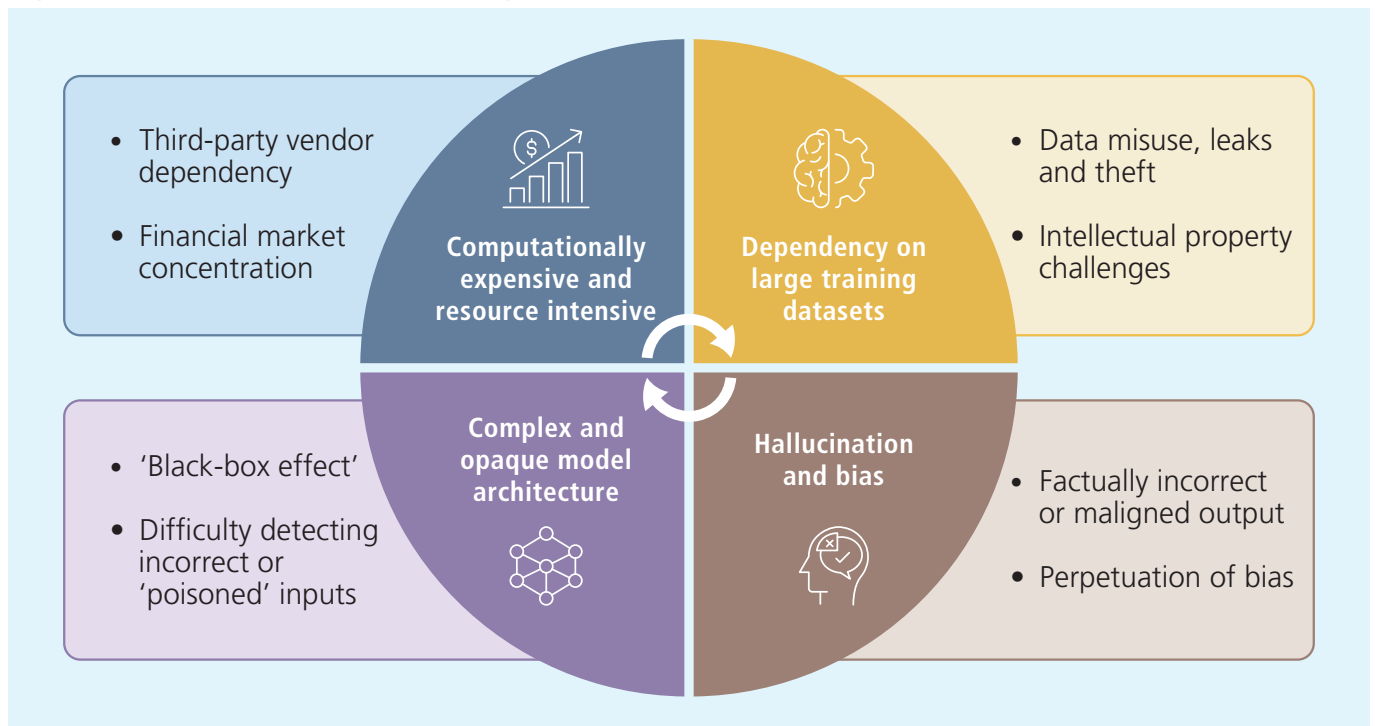
<sup>23</sup> Heikkilä, M. (2023).

<sup>24</sup> In 2023, with some clever prompt engineering, users convinced Microsoft's AI-powered chatbot to share information meant to be kept secret. More recently a new technical report from Apollo Research demonstrates a situation in which GPT-4 deployed as an autonomous stock trading agent can strategically display malign behaviour. Thus, despite the current safeguards in place, it is clear that LLMs have the potential to pose substantial risk with respect to confidential or sensitive information that passes through these systems.

<sup>25</sup> Kaddour, J. et al. (2023).

<sup>26</sup> Bank of International Settlements (2024a).

Figure 1.6: GenA.I. risks and challenges



Source: HKIMR staff compilation.

### Challenges for financial institutions

From the perspective of financial institutions exploring GenA.I. adoption, the use and storage of large data sets associated with the training and processing of GenA.I. models, and the need to safeguard intellectual property and privacy, pose immediate challenges for firms' internal compliance and data management and governance processes (Figure 1.6). The risk of model output errors and biases that are difficult to detect and explain poses challenges for firms' model-risk management, including around model validation, verification and interpretation, and the establishment of clear reporting and escalation procedures. While cybersecurity risks have generally risen as a result of the broader digitalisation of the financial services industry, GenA.I. can exacerbate such risks through the use and storage of large data sets and a reliance on

a concentrated and unregulated set of third-party technology service providers<sup>27</sup>. A reliance on a small handful of proprietary models can also increase the risk of vendor lock-in.

More generally, GenA.I. adoption shares a number of similarities with broader fintech adoption in requiring upfront senior leadership alignment and sponsorship, the setting of clear goals and value-centred use cases, a suitable integration strategy, and clear accountability for results. Novel risks, and the set of legal, ethical, and regulatory grey areas in GenA.I.'s early-stage developments, can also increase the need for specialised technology talent and talent structures that can ensure sufficient human oversight. For financial institutions planning to implement GenA.I. solutions, the above challenges and considerations will likely feature in their decision-making on whether to do so via open-source or proprietary models (Box 1).

<sup>27</sup> Banque de France (2024).

**Box 1: GenA.I. model adoption dilemma: open-source versus proprietary**

The choice of GenA.I. model can be a strategic decision for financial institutions. Open-source (e.g. DeepSeek's R1 and Meta's Llama) versus proprietary (e.g. OpenAI's ChatGPT and Google's Gemini) GenA.I. models can be broadly distinguished across six key characteristics – accessibility, transparency and explainability, adaptability, assumption of liabilities, aftermarket support, and security.

- **Accessibility:** Open-source is arguably more accessible, as the code (building blocks) of open-source models is freely provided and no direct cost is incurred. However, users may still incur infrastructure (e.g. computing servers, data centres) and maintenance costs, especially for complex use cases. Talent needs may also be correspondingly higher. In contrast, institutional users of proprietary GenA.I. models incur the direct cost of using the model in the form of a licence fee. However, ancillary costs (e.g. infrastructure, maintenance, and talent) tend to be lower as these are typically met by the service provider.
- **Transparency and explainability:** Open-source models tend to provide better transparency and explainability, given free access to code, model architecture, and sometimes the parameters around the training data<sup>28</sup>. Hence, the mechanisms by which a model output is generated are better understood. Conversely, proprietary models have strict copyrights in place. Users must pay for often restricted use. The technological infrastructures of proprietary models are also typically trade secrets. Service providers charge users a fee and grant them access only to an interface and the model's output, with no visibility of any 'nuts and bolts'.
- **Adaptability:** Both types have comparable adaptability, albeit in different ways. Users can freely customise open-source models to suit tailored use cases, as long as the right talent and infrastructure are available. Users of proprietary models can fully outsource these demands to the service provider, who in turn may charge additional fees for the corresponding development required.
- **Assumption of liabilities:** On this front, proprietary models provide an added layer of reassurance for financial institutions. While open-source models are copyrighted in a way that allows free use, such as with the MIT licence<sup>29</sup>, the liabilities arising from model failures are usually fully assumed by the user. This contrasts with proprietary licences, where, subject to jurisdiction-specific rules, the provider may either fully or partially assume the liabilities arising from model failures.
- **Aftermarket support:** A key appeal of proprietary models to users is the full provision of development, support, and customisation services. Most open-source implementations lack this feature, although some providers may also provide dedicated aftermarket services, such as hosting, customisation and troubleshooting, at a fee<sup>30</sup>. How critical this consideration is depends on the financial institution's pre-existing GenA.I. expertise.

<sup>28</sup> Often, only the trained model is shared openly, with the training parameters and input data summarised to preserve data privacy.

<sup>29</sup> A short and permissive software licence that permits the user unrestricted and repeated use, including for commercial purposes, of the copyrighted material, as long as the original copyright (hence, author attribution) is included. Many open-source materials, such as those hosted on GitHub, are copyrighted with the MIT licence, including DeepSeek R1.

<sup>30</sup> This business model is akin to Red Hat, which provides technical customisation and support services to corporate clients that host servers based on its own customisation of Linux operating systems, which itself is an open-source software.

- **Security:** Open-source models derive security from a large user base, a non-trivial share of whom may have a vested interest in reviewing and ensuring that their own implementations, hence source codes, are secure. A litmus test is a large and active expert user base capable of stress testing, detecting threats, and patching deep technical issues. In contrast, owing to an active desire to attract and retain customers, as well as the providers' own governance needs, proprietary models may be subject to systematic audit, testing, and maintenance. The specific mechanisms may vary, which demands further due diligence from users.

In the near-term, the choice of open-source versus proprietary, is likely to hinge on compatibility with existing data analytics and AI environments, which is highly specific to each user. Financial institutions with little pre-existing deep GenA.I. expertise may prefer proprietary models, given the lower requirements around talent and infrastructure. Smaller firms wishing to avoid high licence fees may opt for ready-to-use open-source models, though there are still other upfront infrastructure, maintenance, and talent costs to consider.

Over the medium-term, financial institutions may need to consider the risk of vendor lock-in that affect primarily proprietary models, and the risk of infrastructure lock-in<sup>31</sup> common to both types of models. Once the right talent pool, especially trainer-type experts<sup>32</sup>, and suitable infrastructure are developed, and tenable use cases are scoped, financial institutions may also increasingly look to develop their own proprietary GenA.I. models or hybrid deployments for complex and multi-layered use cases.

### Challenges for financial regulators

GenA.I.-related risks can also pose challenges for financial authorities (Figure 1.6). Notably, the current market concentration of GenA.I. developers and service providers may result in the emergence of new systemically important players that fall outside the regulatory perimeter<sup>33</sup>, as well as increase the risk of a failure or cybersecurity breaches of critical AI systems and solutions transmitting into wider operational disruptions. The market dominance of top foundation models trained on common datasets may increase the likelihood of correlated predictions and outcomes<sup>34</sup>, with pro-cyclical market behaviour in turn increasing the risk of systemic events<sup>35</sup>. Autonomous trading powered by GenA.I. models may also use private information to 'inadvertently or intentionally engage in market manipulation and tacit collusion' to profit maximise<sup>36</sup>, thereby undermining market integrity, while the 'black

box' nature of GenA.I. algorithms can complicate the detection and regulation of such behaviours<sup>37</sup>.

The prospective rise in GenA.I. adoption in financial service is also raising broader concerns around how to ensure responsible behaviour in the interests of consumers. GenA.I. models can perpetuate bias and inaccuracies in the data they are trained on, leading to discrimination against some social groups and their access to credit and financial services. Although these issues also exist in traditional AI models, the opacity and reduced explainability of GenA.I. models, and their potential to hallucinate, can make the identification and correction of bias and inaccuracies especially difficult. The challenge of ensuring data privacy and confidentiality when dealing with growing volumes of data, including proprietary data, is another concern. There is a question of how to address these issues in the context of regulatory guidelines and principles, including in Hong Kong<sup>38</sup>.

<sup>31</sup> Costs associated with moving over to new digital and physical infrastructure to facilitate a change in the choice of GenA.I. model.

<sup>32</sup> GenA.I. talent responsible for core development.

<sup>33</sup> Financial Stability Board (2017).

<sup>34</sup> Financial Stability Board (2024).

<sup>35</sup> Georges, C. & Pereira, J. (2021).

<sup>36</sup> Calvano, E. et al. (2020).

<sup>37</sup> OECD-FSB Roundtable on Artificial Intelligence (AI) in France (2024).

<sup>38</sup> Hong Kong Monetary Authority (2024a).

Financial regulators globally have begun revisiting their domestic regulatory and supervisory frameworks in support of responsible AI and GenA.I. implementation. Multilateral regulatory and standard-setting bodies such as the Bank for International Settlements (BIS), the Financial Stability Board (FSB), and the International Organization of Securities Commissions (IOSCO) have also begun assessing the implications of GenA.I. adoption in financial services for financial stability and market integrity. For now, however, the regulatory approach to GenA.I. taken by major AI jurisdictions remains independent. A fragmented global regulatory landscape in relation to GenA.I. may pose challenges<sup>39</sup> not only in leaving open the potential for regulatory arbitrage and a ‘race to the bottom’, but also when regulating financial institutions that are adopting GenA.I. solutions across jurisdictions.

### Preview of report coverage

GenA.I. technology is still at a nascent stage of development, and our understanding of its opportunities and risks is still evolving. Over time, rising competition in model development and lower costs of access for financial institutions wishing to adopt GenA.I. should contribute to a more level-playing field, while progress in the technology should also reduce concerns around some of its current risks.

Against this dynamic backdrop, Chapter 2 discusses the importance of a robust risk management and talent strategy for prospective GenA.I. adopters. Chapter 3 examines current GenA.I. adoption trends in the Hong Kong financial services industry, based on findings from a survey and interviews commissioned by the HKIMR, covering a diverse group of respondents from the banking, insurance, and wealth and asset management (WAM) sectors in Hong Kong. Chapter 4 discusses the current regulatory landscape and its anticipated trajectory. Chapter 5 concludes with a set of considerations for responsible GenA.I. adoption in the financial services industry in Hong Kong.

<sup>39</sup> Asia Securities Industry and Financial Markets Association (2024).



# Chapter 2

## Implications for Financial Services Industry

### Risk Management and Talent Strategies

#### HIGHLIGHTS:

- Financial institutions adopting GenA.I. should develop a robust data governance strategy to identify and mitigate risks, align with business objectives, and refine data processes and standards.
- Effective governance of GenA.I. models requires greater coordination between model developers, validators and users to address human bias, technical defaults and security risks, as well as a strong relationship between data governance and model-risk management.
- To securely integrate GenA.I., firms should adopt a multi-faceted approach, including 'hard defence' measures to strengthen digital infrastructure and 'soft defence' measures to build GenA.I. talent and literacy, while also reskilling and upskilling staff to leverage the benefits of GenA.I..

The speed and breadth of GenA.I. adoption, for financial institutions wishing to integrate the technology into their operations, are likely to depend on a balance between harnessing the benefits of the technology and minimising the risks and costs of implementation. Against this backdrop, a robust risk management and governance framework with clear oversight, accountability, and reporting procedures is key to ensuring safe and compliant GenA.I. adoption. The adoption of GenA.I. by financial institutions can also create new talent demands and challenges that require an effective talent strategy.

## 2.1 STRENGTHENING DATA GOVERNANCE

Data risk management and security have become an integral part of corporate governance amidst greater use of data in business analytics and growing public awareness of data privacy issues. Robust data governance for financial institutions thus serves three key purposes. First, it protects **data quality** and ensures that the data are suitable for their intended use. Second, it safeguards **data privacy** and personal information in the digital age. Third, it supports **data security** and the prevention of data breaches or leaks, which can lead to unauthorised data use and expose customers to identify theft and fraud. Risks to data privacy and security can be especially prohibitive from a compliance perspective and have serious impacts on firms' reputation.

### GenA.I.-related data governance challenges

GenA.I.'s ability to work with unstructured data such as text, images and audio – estimated to represent around 80% to 90% of enterprise data created these days<sup>40</sup> – represents an opportunity for financial institutions to further improve access and enable more financial services tailored to customers. However, the use of unstructured data for decision-making also creates a number of data-related challenges<sup>41</sup>. Notably, the various formats of unstructured data makes it difficult

to apply a uniform definition of data quality through traditional metrics such as completeness, consistency, or validity<sup>42</sup>. Data quality issues can also contribute to GenA.I. hallucination and bias.

Meanwhile, the democratisation of GenA.I. tools can support increased staff efficiency, but it can also increase the risks of data mishandling by staff. The use of previous customer conversations and transactions as input into GenA.I. applications can also challenge perceived notions of customer privacy, compelling greater clarity on data ownership, accessibility, storage, and security. Data security risks may also increase as the transition of data stored by or transmitted to third parties, such as providers of software-as-a-service models and cloud services, further broadens the channels for data theft.

According to the latest statistics from the EU General Data Protection Regulation (GDPR), the cumulative fines for data breaches had risen to EUR4.48 billion as of March 2024, with a total of 2,086 cases<sup>43</sup>. Data processing with insufficient legal basis or non-compliance with general data processing principles continue to be the key drivers of data breaches and fines. Although the highest average fines remain concentrated in 'Media, Telecoms and Broadcasting', large-scale data breaches and increased public awareness of data privacy issues have put the use and protection of personal data higher up the agenda.

### Strengthening data governance

An effective data governance strategy is thus even more critical in the era of GenA.I. as an anchor for responsible innovation and adoption. In addition to the inclusion of key functions such as data risk management, data access and security, and data warehousing as part of a robust data governance framework, the speed of GenA.I. innovations indicate the need for an iterative and proactive approach that can better respond to the emergence of new risks and challenges, shifts in customer trust and ethical expectations, and new compliance requirements.

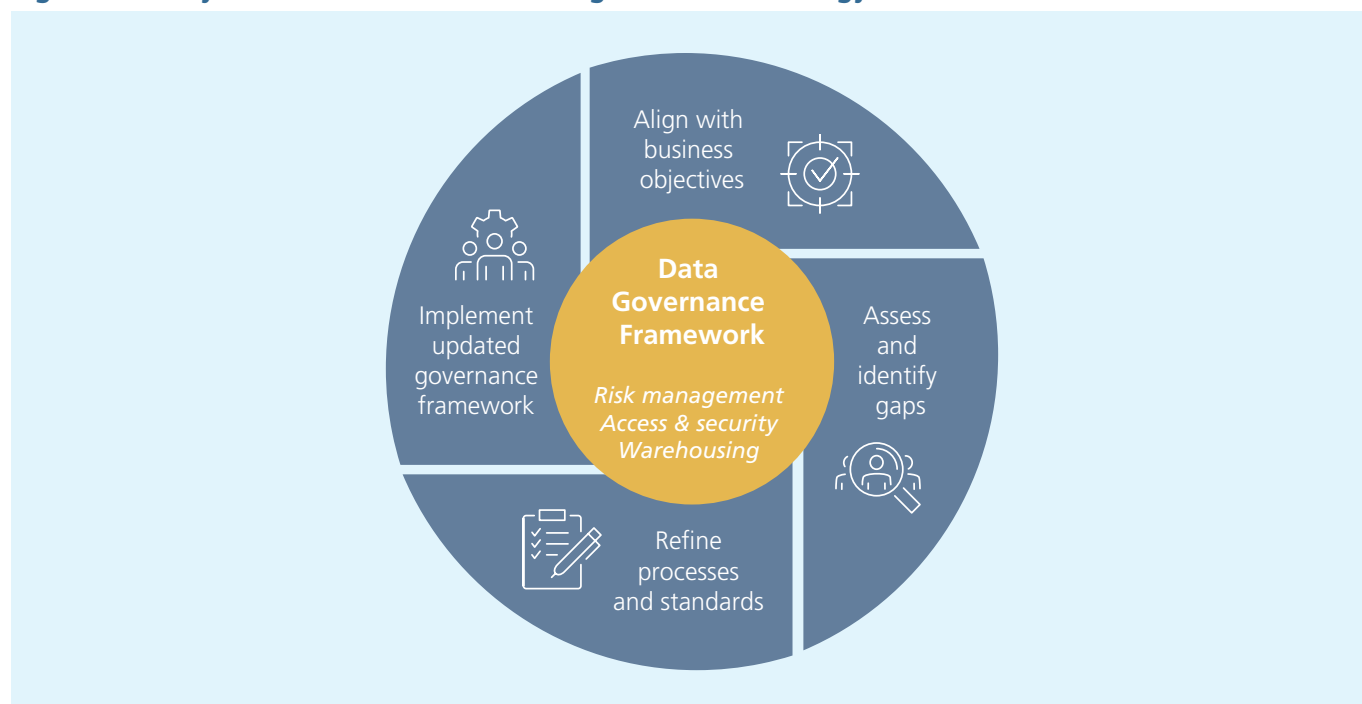
<sup>40</sup> MIT Management Sloan School (2021).

<sup>41</sup> Harvard Business Review (2024).

<sup>42</sup> Securiti (2024).

<sup>43</sup> CMS GDPR Enforcement Tracker Report (2024).

Figure 2.1: Key elements of a robust data governance strategy



Source: HKIMR staff compilation.

Notably, there should be an **alignment of the data strategy with business objectives**. Firms should prioritise particular business areas and data domains where they expect to implement GenA.I. in the near future, alongside a data strategy roadmap for longer-term GenA.I. adoption (e.g. start small and scale up over time). The objectives of data governance should be aligned with specific business aims, such as improving the efficiency of data management or meeting compliance requirements. Management also needs to develop a suitable communication strategy with major stakeholders, and find the right methods, technologies, and people to implement the data strategy throughout the organisation.

This should be followed by the **assessment and identification of risk management gaps**, through an evaluation of the risks to data accuracy, privacy, and security that may arise through GenA.I. adoption. This includes reviewing the adequacy of existing data-related policy and their preparedness for broader GenA.I. adoption and reaching consensus with major stakeholders such as IT, the

risk control office, and business units on how to address gaps in the risk management framework. Additional investment in data infrastructure and in human oversight (e.g. centralised data governance committee) may be needed.

There should also be a **refinement of data processes and standards** – from data collection and transformation to classification and consumption – with clear delineation of accountability and reporting procedures. Data integration and sanitisation of datasets that may support the use of GenA.I. should be done pre-emptively, such as by curating documents or text to prepare them for domain-specific GenA.I. models and applications. Clear standards and processes should be consistent across the organisation. To ensure data privacy and security, it may be necessary to set separate layers of access control based on confidentiality and materiality across different data categories such as customer, corporate, and analytical data to enhance data security. Having in place the procedures to enable incident reporting of data privacy violations and remediation measures is also key.

Finally, **implementation** should take place through an improvement of data architecture, data access and security controls, and expertise. There needs to be adequate oversight in the form of data risk officers to ensure that the collection and use of customer data comply with relevant data privacy and regulatory requirements, and that data sharing with third parties, such as a credit reference agency, is in line with company policy. Data risk officers should be equipped to detect irregularities or defects in data series, and liaise with data vendors to assure and improve data quality. With the increasing use of big data in AI models, data risk officers also need to monitor the proper use of third-party information from the media and the internet, and take remedial actions if a data breach or misuse is identified.

## 2.2 EFFECTIVE RISK MANAGEMENT OF MODEL-RISK

Available evidence suggests there has been an acceleration in the adoption of AI in financial services in recent years<sup>44</sup>. By some estimates, global spending on AI could exceed USD600 billion by 2028<sup>45</sup>, with the financial services industry accounting for one-fifth of total spending. However, an increase in model-driven activities and over-reliance on quantitative models can also lead to significant challenges. AI models can support financial decision-making, but only if the algorithms and statistical techniques used and their implementation are clearly aligned with real-world business problems and if the model results are relevant, accurate, and informative in supporting the desired business outcome.

As a subset of operational risk, AI model-risk – together with the consequences of poor design, calibration, or implementation – can lead to inadequate or erroneous decision-making. AI models can also be used inappropriately if there

are issues related to data availability, data quality, or representativeness that lead to sampling bias and a lack of fairness. In this respect, AI models are no different than traditional financial models. However, unlike traditional models, AI models' reliance on high-dimensional data, and greater model complexity and opacity, can make model-risks much more difficult to identify and assess. AI models also tend to undergo recalibrations more often than traditional financial models.

Mistakes in model-based decision-making can lead to financial losses and regulatory fines, and can impact negatively on firm reputation. Effective management of AI model-risk is thus critical to creating stakeholder trust and accountability, and has become a key area of focus for financial institutions looking to accelerate their use of AI in support of greater efficiency<sup>46</sup>. This typically requires a robust framework of standards, testing, and controls both at inception and throughout the AI model life-cycle to achieve proper coverage of risks.

### GenA.I. implications for model-risk

Previous research<sup>47</sup> has identified four risk factors that may lead to the improper use of AI models: human bias, technical defaults, usage defaults, and security defaults. The complexity of GenA.I. models, and their more widespread use of massive, unstructured data sets and the lack of transparency in training data sources, can further increase the likelihood and impact of these risk factors crystallising, as outlined in Figure 2.2. As the challenges specific to AI systems may be further amplified under GenA.I. approaches, the need for firms to have a robust and strategic approach to the monitoring and management of model-risk should be a prioritised aspect of GenA.I. adoption by financial institutions.





<sup>44</sup> Financial Stability Board (2024a).

<sup>45</sup> International Data Corporation (2024).

<sup>46</sup> Ernst & Young/EY (2020).

<sup>47</sup> Hong Kong Institute for Monetary and Financial Research (2020).

Figure 2.2: Re-evaluation of risk factors that may lead to problematic AI models in the era of GenA.I.

<div>Human bias</div> <div></div>	Biases of model developers and users can affect the outputs, and the black box nature of GenA.I. models can make it even more difficult to ascertain whether models are performing without such bias.
<div>Technical defaults</div> <div></div>	GenA.I. models' greater complexity and lack of transparency, and the use of a wider variety of data types and sources, can further increase the difficulty of finding independent and knowledgeable model validators who can effectively challenge model development approaches <sup>48</sup> .
<div>Usage defaults</div> <div></div>	GenA.I. models trained on proprietary or confidential data can inadvertently leak sensitive information through their outputs <sup>49</sup> . Hallucination can also be difficult to detect and evaluate, with limited explainability of some GenA.I. approaches further impeding the evaluation of output suitability and soundness.
<div>Security faults</div> <div></div>	GenA.I. models may be particularly susceptible to adversarial attacks and manipulation due to the expansion of the attack surfaces <sup>50</sup> . Maintaining up-to-date world knowledge also requires continual training on new data, which can turn these risks into a persistent threat.

Source: HKIMR staff compilation.

Strengthening GenA.I. model-risk management

Some jurisdictions, including Hong Kong, have already introduced specific standards and high-level principles for regulated financial institutions aimed at the management of AI model-risk, and these typically articulate specific expectations around (i) governance and accountability; (ii) model validation; (iii) model-risk assessment; (iv) model life-cycle monitoring and performance tracking; and (v) risk mitigation and contingency planning (Figure 2.3).

Model introductory stage

**Governance and accountability:** Senior management should have good oversight and understanding of use cases being explored, and consider whether these require enhancements to existing model-risk and related (e.g. data and information security) control processes. In-house skills training or the hiring of external experts familiar with GenA.I.-specific risk management issues and model development and validation techniques should also be considered. In addition, there needs to be clear articulation and delineation of responsibilities across the model life-cycle and across the various risk management functions (e.g. model-risk, model validation, and compliance).

<sup>48</sup> Financial Stability Board (2024a).  
<sup>49</sup> Wu, X., Duan, R., & Ni, J. (2024).  
<sup>50</sup> Zhu, B., Mu, N., Jiao, J., & Wagner, D. (2024).



*Model implementation stage*

**Model validation:** This refers to a set of processes designed to verify the robustness of models and ensure that these perform as intended, which is an essential step in identifying irregularities in model outcomes and towards initiating early remedial actions. This includes examining the model construction and data used, back-testing, and ensuring that the model meets regulatory compliance and internal governance requirements<sup>51</sup>. As the use of GenA.I. models can create 'black boxes' in decision-making, there is an added need to consider where data limitations can skew a model's outputs. It may be prudent to involve an independent party (e.g. the second or third line of defence or an external consultant) in the model validation process.

**Model-risk assessment:** This involves an assessment of the overall risk profile of models across a basket of factors such as materiality, complexity of methodology, financial impact, and performance soundness. In the case of GenA.I., the aggregate risk profile of a particular model and its intended application should be aligned with financial institutions' internal risk appetite thresholds, and this risk profile can provide a useful benchmark to rationalise the amount of resources allocated to tracking the performance of applications that pose high risk.

*Ongoing monitoring and maintenance*

**Model life-cycle monitoring and performance tracking**<sup>52</sup>: Due to the complexity of GenA.I. models and their capacity for frequent recalibration in response to new data, it is important that risk-mitigating controls are embedded into all stages of a GenA.I. model's life-cycle, from development to deployment and use. Periodic reviews that include a re-validation of the model may also be appropriate alongside ongoing monitoring and performance tracking to ensure that the applications continue to perform as intended and to align with compliance and governance requirements.

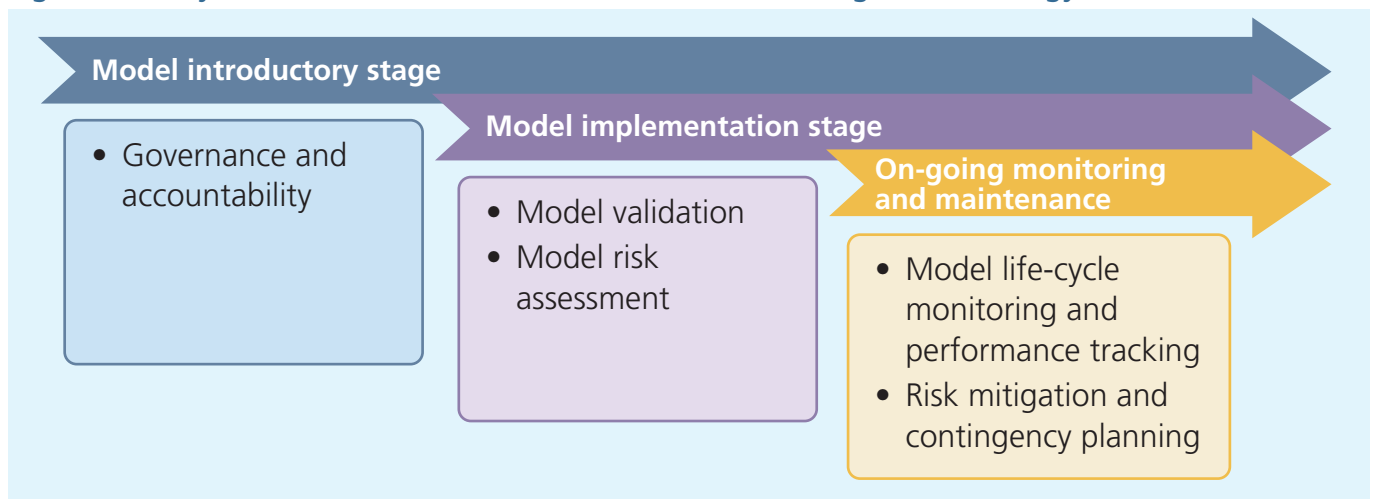
**Risk mitigation and contingency planning:** Even the most robust GenA.I. applications may deliver unintended outcomes. In addition to appropriate risk-mitigating controls (e.g. 'human-in-the-loop' mechanism, prudent risk limits, and sample quality assurance checks), there should be contingency processes in place that can promptly suspend GenA.I. applications and trigger fall-back procedures (e.g. human intervention or conventional processes).

While the design and governance of AI models throughout the model life-cycle are key to the effective and safe development of AI and its broader use in financial services, the added opacity of GenA.I. models may require greater coordination between model developers, validators, and users than is currently in place for traditional AI models. The higher frequency with which GenA.I. models continuously learn from data to refine their decision-making process also indicates a more intricate relationship between data governance and model-risk management in the era of GenA.I..

<sup>51</sup> Weedmark, D. (2022).

<sup>52</sup> KPMG (2022).

Figure 2.3: Key elements of a robust GenA.I. model-risk management strategy



Source: HKIMR staff compilation.

## 2.3 MITIGATING NEW SOURCES OF CYBERSECURITY RISKS

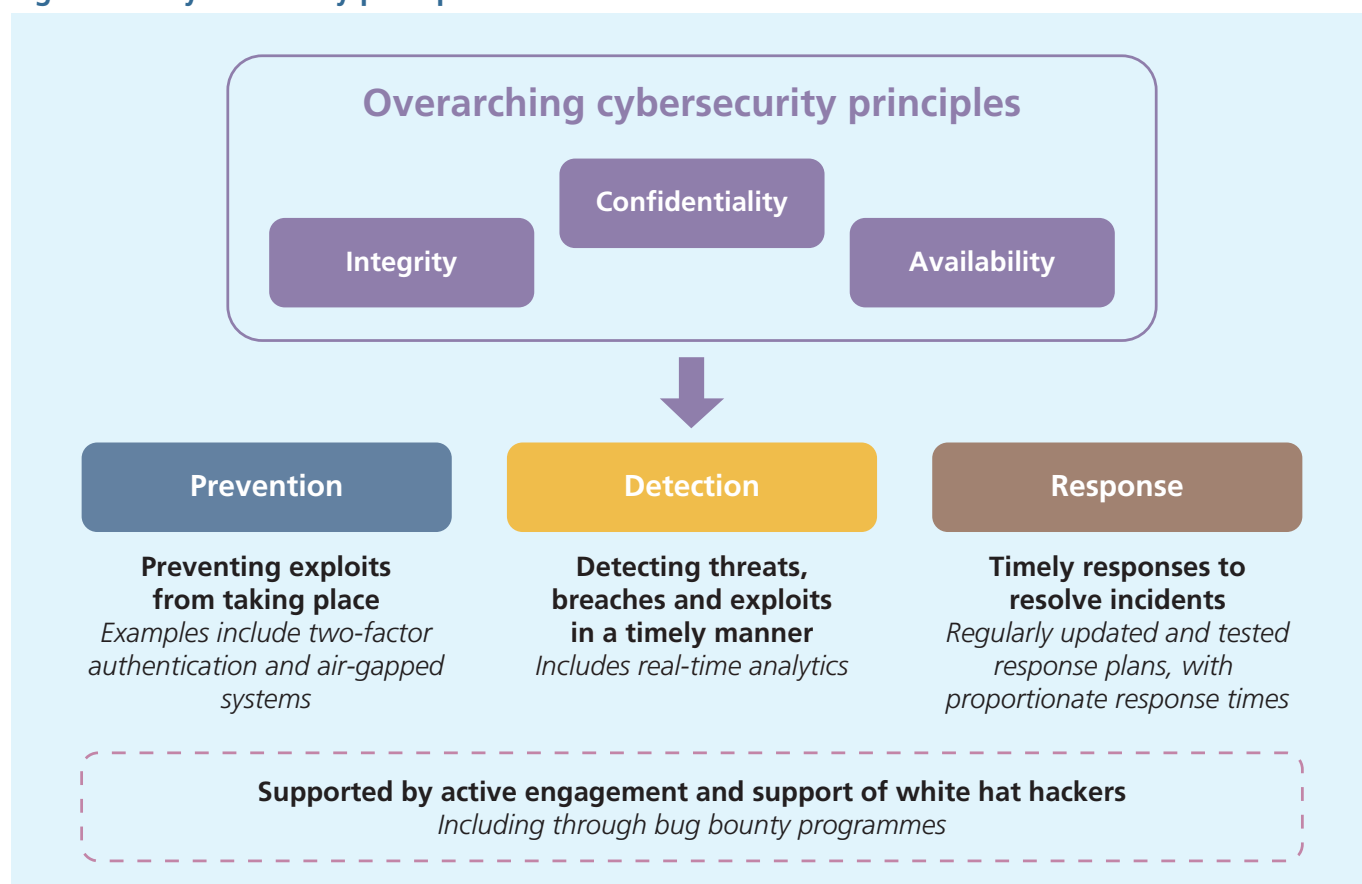
Modern financial institutions operate on a deep network of digital infrastructure to support their daily business activities and tend to be at the frontier of technology. As such, financial institutions are inevitably confronted with multi-faceted cybersecurity risks in their day-to-day operation. And while simple data breaches and system outages may only pose inconveniences in other industries, these can result in severe reputational damage and financial losses for affected financial institutions, which may take years to recover from. This elevated risk profile is reflected in the generally extensive investment made by the financial services industry in state-of-the-art cybersecurity systems<sup>53</sup>. Cybersecurity breaches can also pose risks to broader financial stability.

Cybersecurity systems are guided by three high-level principles: confidentiality, integrity, and availability. First, cybersecurity systems should keep secure sensitive information and trade secrets ('confidentiality'). Second, the data, output, and underlying processes should be consistent and trustworthy throughout the system's life-cycle until retirement ('integrity'). Third, technology tools and data should be made as widely available as possible without compromising confidentiality and integrity ('availability').

In the financial services industry, these three high-level principles in turn guide the implementation of tools geared at **prevention** against exploitation, **detection** against threats and breaches, and timely and robust **response** (Figure 2.4). The use of ethical hackers ('white hat hackers') to test the robustness of cybersecurity systems (e.g. the Meta Bug Bounty programme) can further ensure vulnerabilities can be revealed and resolved appropriately before malicious attacks can take place.

<sup>53</sup> McKinsey (2024) estimated that financial institutions allocated around 13% of their IT budget to cybersecurity.

Figure 2.4: Cybersecurity principles and measures



Source: HKIMR staff compilation.

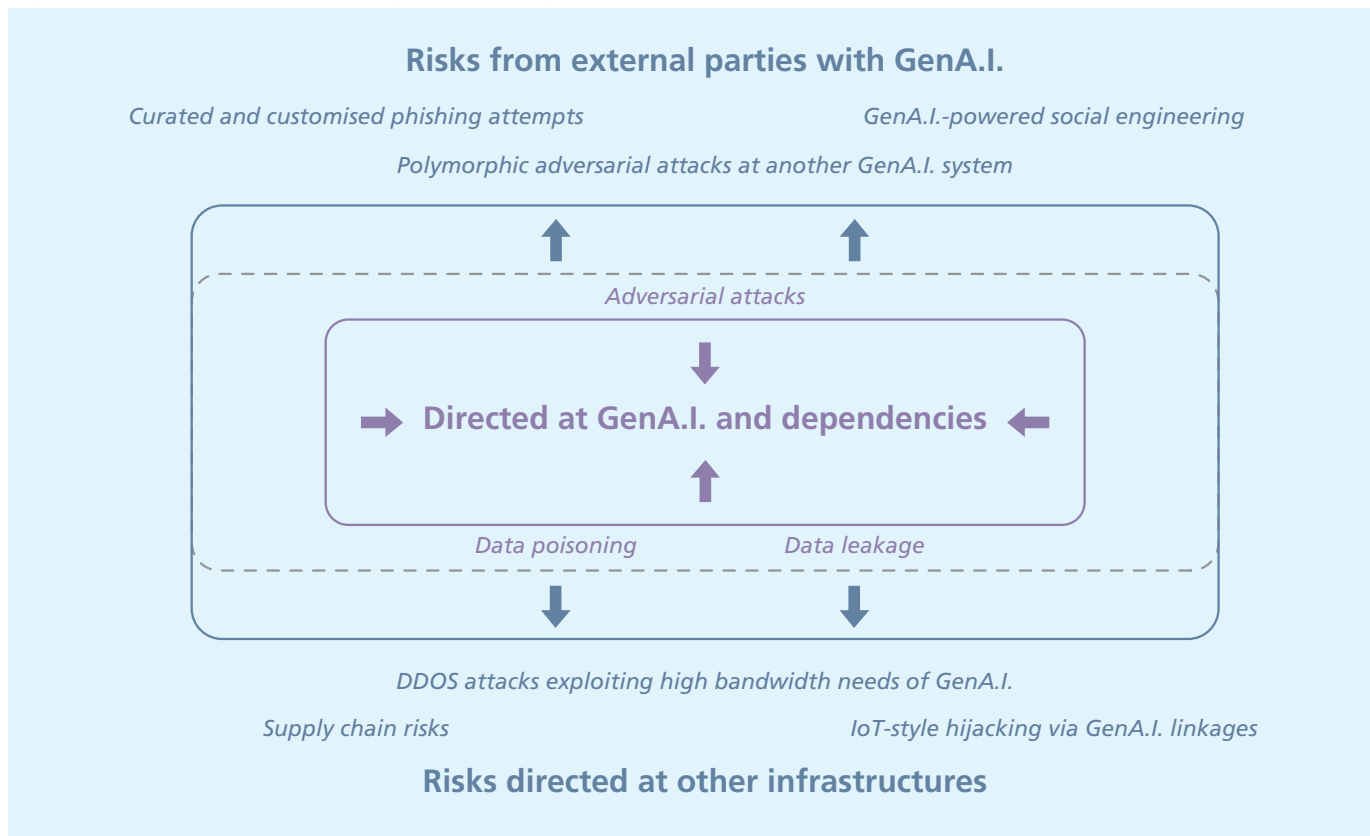
### GenA.I.-related cybersecurity challenges

The adoption and integration of GenA.I. within the financial services industry's day-to-day technology infrastructure introduce new cybersecurity threats as the technology evolves. Such threats can be broadly separated into two groups (Figure 2.5).

First, those targeting GenA.I. systems or infrastructures integral to GenA.I. systems. The most common are adversarial attacks aiming to produce incorrect or harmful output or to trigger data leaks via careful prompt engineering. In the absence of appropriate input filters, GenA.I. tools may also be prone to data poisoning where malicious input is injected or simply fed as a carefully crafted prompt to trigger a retraining of the GenA.I. model to produce malicious or incorrect output. This may be a particular concern for GenA.I. models that are continuously retrained with live updates to their training data, and can lead to serious financial losses if placed in business-critical systems, such as GenA.I.-enabled portfolio management algorithms.

Second, those arising from malicious GenA.I. users or that exploit the linkages between GenA.I. applications and the rest of the targeted institution's infrastructure. For instance, GenA.I. can be leveraged to produce malwares and social engineering and phishing strategies with greater scale, efficiency, and precision, such as the design of curated phishing hooks across a wide range of victim profiles. Distributed denial-of-service (DDOS) attacks can also exploit the high processing bandwidth needs of GenA.I. to stall other essential parts of the business operation, such as online banking services, that rely on the same or related computing infrastructure. Toolkits that connect and exchange data with GenA.I. applications through 'Internet-of-Things'-style integration, such as investment analytics tools with optional GenA.I. modules, may also be hijacked. Such threats could further give way to concentrated supply chain risks, especially if critical workflows exposed to GenA.I. do not have immediate workarounds.

Figure 2.5: GenA.I.-related cybersecurity threats



Source: HKIMR staff compilation.

### Key elements of an effective cyber defence strategy

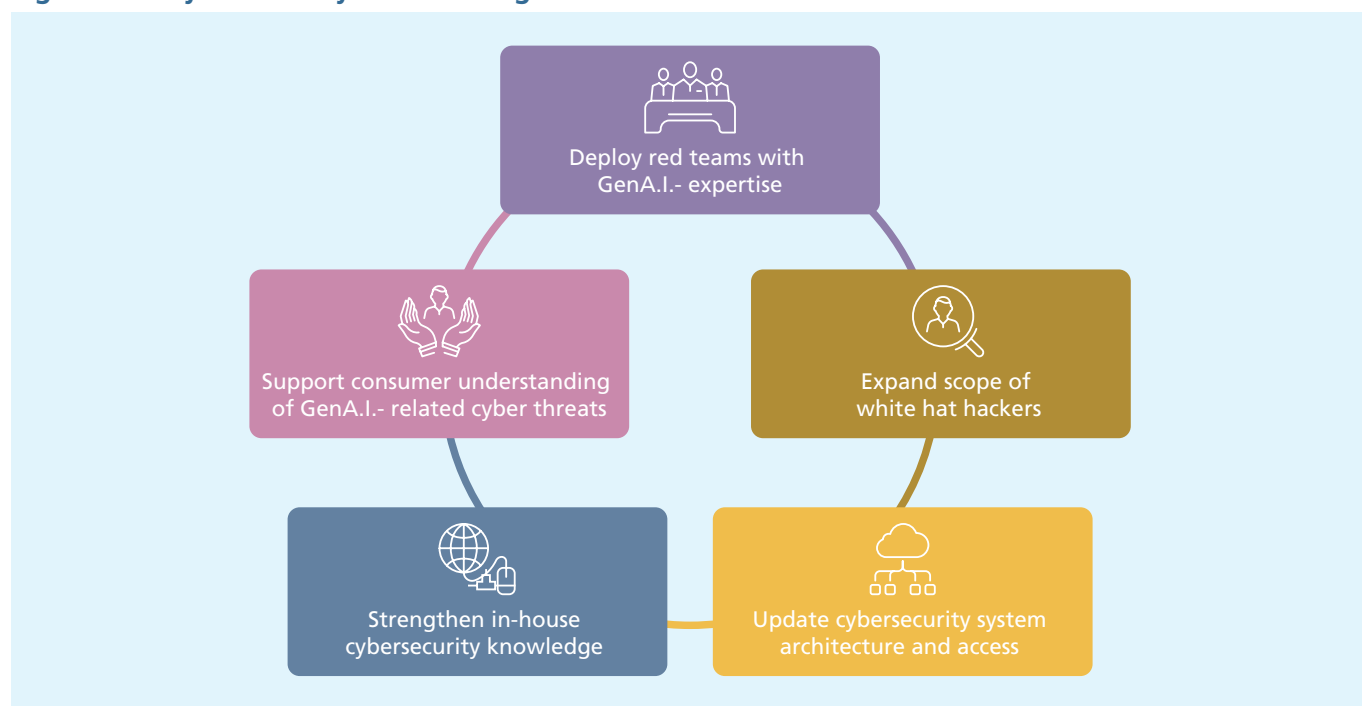
Even previously robust cybersecurity frameworks may require modifications to effectively address new threats as GenA.I. adoption continues (Figure 2.6). Such modifications can be divided into 'hard defence' measures aimed at the digital infrastructure and the arguably equally critical 'soft defence' measures aimed at people and the organisational environment.

Financial institutions keen to build up effective hard defence measures can explore the use of dedicated teams with expertise in GenA.I. tools and techniques ('red teams') that can adequately test the boundaries of the security features of the GenA.I. system, including through appropriate prompt engineering. Ethical security hackers ('white hat hackers') can also be employed to identify and fix GenA.I. vulnerabilities. Moreover, careful design of the cybersecurity system architecture and access controls are warranted,

ideally one that is decentralised with minimal number of linkages between critical systems, alongside a regularly updated incident response plan.

Fundamentally, GenA.I. adopters need to keep up with the development frontier, and cybersecurity best practices need to evolve as new vulnerabilities and threats are discovered. Soft defence measures thus include building up a critical mass of GenA.I. talent and literacy within the organisation, especially in response to emerging cybersecurity threats. In particular, a pipeline of GenA.I. experts will be needed to efficiently identify patches of vulnerabilities and to address constantly evolving threats. An entrenched GenA.I. usage and safety literacy culture within the organisation can also ensure that day-to-day operations are undertaken with a conscious need to defend against GenA.I.-related cybersecurity threats, while promoting an understanding of GenA.I. risks can strengthen consumer protection.

Figure 2.6: Cybersecurity measures against additional GenA.I.-related risks



Source: HKIMR staff compilation.

## 2.4 TALENT STRATEGY IN THE ERA OF GenA.I.

The successful adoption of GenA.I. by financial institutions requires not only adequate digital infrastructure but also a workforce equipped with up-to-date skills to utilise GenA.I. effectively. As GenA.I. adoption broadens, there is also likely to be increasing workplace tension between confidence in the technology and concern around worker displacement<sup>54</sup>. As a result, firms' talent strategies will likely need to change in response to the opportunities and the challenges brought by GenA.I..

Our recent global talent survey<sup>55</sup> found that structural sociodemographic developments and post-pandemic employment trends have contributed to a widening of three gaps in the talent landscape for financial services. First, demographic changes, lower migration, and declining labour participation have produced a **talent gap**, with an insufficient pool of qualified candidates to fill current and future roles. Second, there has been a widening of an industry-wide **skills gap**, with the evolution of skills needed due to societal and technological changes driving a mismatch between the skills that firms require and the skills of current and potential employees<sup>56</sup>. Third, changing employee expectations regarding work-life balance, flexibility, and meaningful work have increased the **performance gap** between actual and desired business outcomes, as financial services firms struggle to retain and motivate talent.

<sup>54</sup> BCG (2024).

<sup>55</sup> Hong Kong Institute for Monetary and Financial Research (2023).

<sup>56</sup> Our previous talent survey revealed that approximately two-fifths of respondents reported that their firms faced intense competition when attempting to attract candidates with emerging skills.

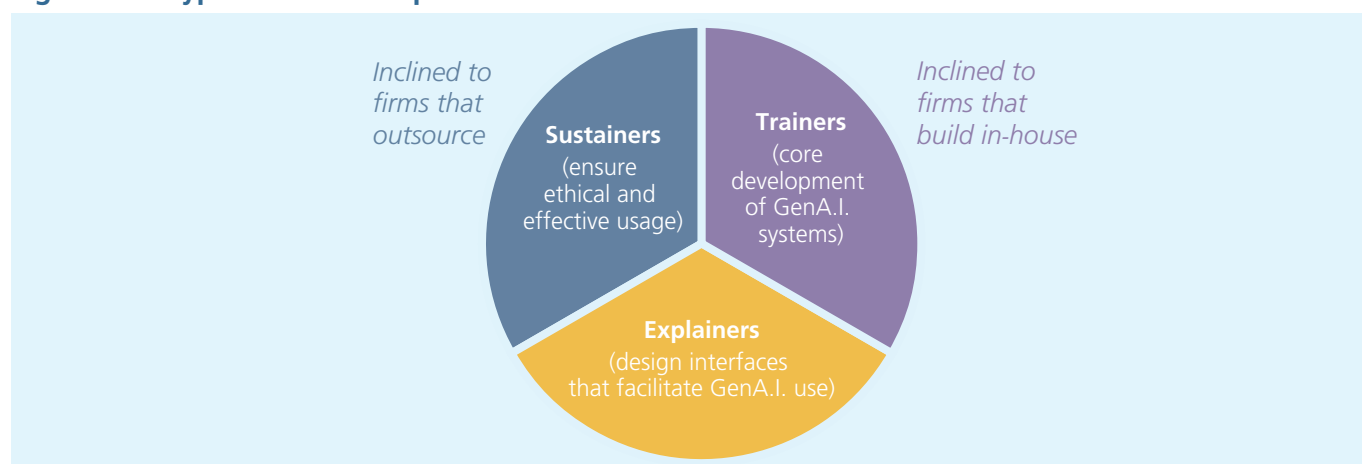


Collectively, these talent gaps may have significant implications for the growth and development of the financial services industry in Hong Kong. In particular, the demand for composite talents who can integrate technical proficiency with effective leadership has grown significantly in recent years. In the Asia-Pacific region, core technical skills such as FinTech, AI and data analytics have become high-priority areas for talent acquisition and training. Against this backdrop, the integration of GenA.I. into the workplace is anticipated to further alter the talent landscape and introduce additional talent challenges and considerations.

### GenA.I.-related talent challenges

As illustrated in Figure 2.7, in the near-term, GenA.I. is likely to further broaden the talent gap in the financial services industry by introducing the need for three new categories of skilled talent<sup>57</sup>: (i) Trainers responsible for the core development of GenA.I. systems, such as the highly-skilled engineers, software developers, and data scientists who build the LLMs behind GenA.I. applications and who often command a high wage premium<sup>58</sup>, or the model developers who can fine-tune pre-trained third-party LLMs for a specific task; (ii) Explainers responsible for the design of interfaces that make GenA.I. accessible, thus serving as a bridge between trainers and end-users; and (iii) Sustainers responsible for ensuring that GenA.I. systems are implemented ethically and effectively.

**Figure 2.7: Types of GenA.I. specialists**



Source: Shine (2023) and HKIMR staff compilation.

At the firm level, the size and composition of the perceived GenA.I.-related talent gap is closely linked to the approach to GenA.I. adoption and implementation, and the degree of outsourcing. For financial institutions preferring to develop GenA.I. applications in-house, the demand for trainers will be especially high, while firms looking to incorporate GenA.I. through partnerships or a combination of internal development and external

solution acquisition are likely to place greater emphasis on addressing talent gaps in explainers and sustainers.

As GenA.I. solutions are computationally expensive and resource intensive to develop, financial institutions' need for trainers should be limited in the near-term. Explainers, being responsible for integrating pre-trained but still highly technical

<sup>57</sup> World Economic Forum (2023).

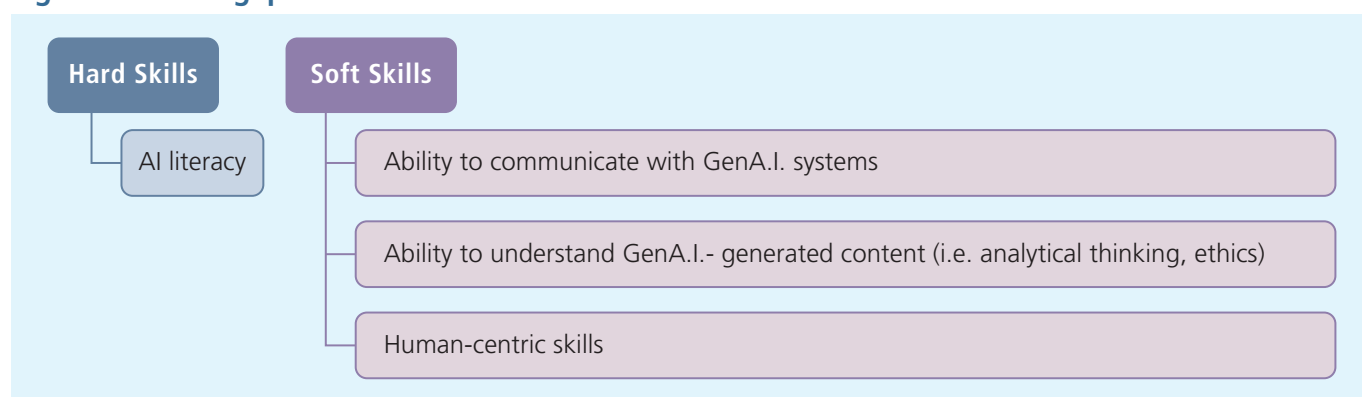
<sup>58</sup> OECD (2023).

GenA.I. models into a company's workflow or system, are likely to be in greater relative demand as GenA.I. adoption picks up. For example, interface and interaction designers as a category of explainers are needed to adapt applications such as personalised GenA.I. assistants to various types of user input, be it text or spoken voice, for improved accessibility. The demand for sustainers, such as ethics and governance specialists, should also become more prevalent as GenA.I. becomes more integrated into business operations, both in

ensuring input data quality and output accuracy during the in-house LLM fine-tuning process, and as gatekeepers in ethics and legal compliance.

In the era of GenA.I., as the financial sector increasingly transforms from a human-centric to a GenA.I. augmented environment, employees will also need to consider how to effectively leverage the opportunities GenA.I. applications can bring. The acquisition of hard technical knowledge and improved AI literacy, as well as soft skills, is required<sup>59</sup> (Figure 2.8).

**Figure 2.8: Skill gaps in the GenA.I. era**



Source: HKIMR staff compilation.

For employees of all levels, a better understanding of AI systems and their merits and limitations, and how different GenA.I. applications can complement human capabilities and support more efficient decision-making, is an increasingly essential skill. The arrival of LLMs has also lowered the importance of skills such as written communication<sup>60</sup>, with a shift in prioritisation towards soft skills that demonstrate human-unique characteristics in an increasingly hybrid GenA.I.-human work environment. Indeed, an interview of HR leaders in Hong Kong found that among the top talent attributes sought by employers were soft skills that are difficult for machines to replicate, such as analytical comprehension and creativity, communication, and adaptability with a growth mind set.

In addition, there is value in employees' ability to communicate with GenA.I. systems through effective prompt engineering. Appropriate comprehension and utilisation of the results generated by GenA.I., and the ability to discern whether the generated output is credible, fair, and representative, are also key. As GenA.I. systems continue to lack the human touch essential for building relationships, understanding emotional context, and handling ethical complexities, it is expected that the effective communication, good interpersonal skills, and strong problem-solving capabilities of employees, especially in the context of dynamic team structures, will be further strengthened. An evolving GenA.I. landscape is also likely to increase the demand on senior management around talent development, staff motivation, and sound and ethical decision-making.

<sup>59</sup> Harvard Business Review (2023a).

<sup>60</sup> World Economic Forum (2024a).

The introduction of GenA.I. into the workplace is expected to bring an overall productivity boost to employees, but GenA.I. may also partially or entirely replace the job duties of some workers, with a number of recent surveys indicating increased employee concerns over potential job losses and AI replacement<sup>61</sup>. Building on the general underlying shift in the workplace towards more fulfilling roles with 'meaningful' impact and better 'work-life balance', an absence of adequate transition and retraining opportunities alongside GenA.I. adoption may lead some employees concerned about prospective job displacement to experience an early loss of purpose and become disengaged. This may impact the effectiveness of GenA.I. integration, with some workers potentially less able or less willing to adapt to the new technology.

### **Implications for financial institutions' talent strategy**

In the future workplace, tasks that involve management oversight, strategic decision-making, interpersonal care, emotional perception and relationship building are likely to remain predominantly human-driven<sup>62</sup>. While GenA.I. may partially automate some aspects of a job role, it can improve workers' ability to perform other tasks more effectively, thereby enhancing overall productivity<sup>63</sup>. Against this backdrop, an 'AI-first' approach in the financial services industry that focuses solely on automation and headcount reduction, rather than GenA.I.'s augmentation potential, could be counterproductive by demotivating employees and leaving them less committed to AI initiatives, ultimately leading to the aforementioned performance gap.

Instead, a balanced talent approach with reskilling and upskilling as part of any firm's core talent strategy can help address talent and skills gaps and build trust, while facilitating more successful GenA.I. adoption. Indeed, training and worker consultation have been found to be associated with better outcomes for workers and performance<sup>64</sup>. Moreover, retaining domain expertise and having more skilled 'human-in the loop' can better enable financial institutions to navigate the hallucination-prone nature of GenA.I., and facilitate a healthier adoption.

Meanwhile, in roles where GenA.I. and human capabilities blend, hard and soft skills are more likely to complement each other, rather than being mutually exclusive. Harnessing AI literacy can significantly boost productivity in the workplace, while human-centric soft skills are essential for the effective and efficient utilisation of GenA.I.. The absence of any one aspect could hinder the full potential of GenA.I. adoption in the workplace. Preserving domain expertise, i.e. in-depth knowledge in specific fields like financial services, should also be prioritised, as such knowledge can help to provide essential checks and balances, especially when the technology is still evolving.

<sup>61</sup> OECD (2023b).

<sup>62</sup> Ernst & Young/EY (2024).

<sup>63</sup> World Economic Forum (2024b).

<sup>64</sup> OECD (2023b).

# Chapter 3

## Insights from Market Participants in Hong Kong

### HIGHLIGHTS:

- The adoption of GenA.I. is progressing steadily across the financial services industry in Hong Kong. 75% of surveyed financial institutions have already implemented or are currently piloting and designing GenA.I. use cases, primarily to enhance productivity and operational efficiency in non-customer facing tasks. This ratio is expected to increase to 87% within the next three to five years.
- However, there are a number of risk management challenges hindering adoption, including concerns regarding model accuracy, data privacy and security, as well as constraints related to talent.
- Financial institutions seek further regulatory guidance on cybersecurity standards, assessment of third-party vendors, accountability, and data localisation and sharing, supported by practical examples and additional facilitation measures.


To support understanding of the state of GenA.I. adoption among financial institutions in Hong Kong, the HKIMR commissioned a survey titled *Financial Services in the Era of Generative Artificial Intelligence: Opportunities and Risk Management*, from October 2024 to January 2025, which gathered the views of market participants on the current state of GenA.I. adoption among local financial institutions, the expected trajectory of GenA.I. development in Hong Kong, and the strategies employed for risk management and talent development. The survey questionnaires were sent to respondents from the banking, insurance, and WAM sectors in Hong Kong. Interviews were also conducted with a diverse group of market participants, including financial institutions and GenA.I. service providers, to explore specific perspectives in greater depth<sup>65</sup>.

### 3.1 STATE OF GenA.I. ADOPTION AMONG FINANCIAL INSTITUTIONS

The adoption of GenA.I. is progressing steadily across the financial services industry in Hong Kong. 75% of the surveyed financial institutions have already implemented at least one GenA.I. use case, or are currently piloting and designing use cases, and exploring potential investment areas. This ratio is expected to increase to 87% within the next three to five years (Figure 3.1).

GenA.I. adoption has been somewhat higher among the larger surveyed financial institutions. Among surveyed firms, 83% of large firms have rolled out at least one GenA.I. use case or are taking steps towards adoption, compared with 63% of small firms. Larger insurers with access to resources for investing in AI initiatives also expressed a keenness to explore GenA.I., which was particularly applicable to insurance-related tasks such as claims processing and fraud detection. A number of larger banks also highlighted the technology's more immediate scope for improving organisational efficiency and internal processes.

Figure 3.1: GenA.I. adoption among the surveyed financial institutions



Source: HKIMR staff compilation based on the GenA.I. Survey.


(1). 'In production' refers to respondent firms that have already rolled out at least one GenA.I. use case to production.

<sup>65</sup> Full details of the Survey can be found in Appendix A.

Our survey found that the primary implementations of GenA.I. in financial services remain largely internal and non-customer facing. The most common GenA.I. use case is the integration of virtual assistants for employees, with customer-facing applications restricted to customer chatbots and enhanced sales and marketing (Figure 3.2).

This was broadly consistent with findings of a separate study<sup>66</sup>, and aligned with the role GenA.I. is anticipated to play for employees, as 75% of the surveyed financial institutions viewed GenA.I. as a tool to enhance productivity and operational efficiency, followed by 53% that viewed GenA.I. as empowerment for innovation and decision-making.

Figure 3.2: Top GenA.I. use cases in front, middle, and back office



Source: HKIMR staff compilation based on the GenA.I. Survey.

The interviews indicated that while most firms recognised the strategic potential of GenA.I., more complex applications necessitated greater trust in the accuracy of the technology, with respondents from the banking sector for instance cautious regarding the use of GenA.I. at its current stage of development in revenue-sensitive and critical business areas. The limitations of current GenA.I. models in replacing certain front-office expertise and a lack of proficiency in interpreting finance-specific terminology and problems also contributed to scepticism among some interviewees about the return on investment from complex models.

Although many small firms expressed enthusiasm for GenA.I.’s potential, finance and human resource limitations, and less advanced existing technology infrastructure, acted as adoption hurdles.

Looking ahead, the emergence of less resource-intensive models and maturing technology, coupled with regulatory engagement, are likely to contribute to the broadening of GenA.I. adoption over time. Indeed, DeepSeek-R1 is already challenging the prevalent view that scaling GenA.I. requires vast computing power and investment.

<sup>66</sup> HKMA (2024). The study found that GenA.I. adoption in the Hong Kong financial services industry was primarily aimed at addressing (i) information overload, (ii) repetitive tasks, and (iii) human errors, with most Hong Kong-based organisations focused on internal-facing use cases aimed at supporting employee empowerment, driven by the perception that the risks associated with internal applications are more manageable and controllable, alongside concerns about whether GenA.I. technology is currently sufficiently advanced to consistently deliver reliable and accurate results directly to customers.




New approaches to language modelling are also improving model accuracy and general performance. The trajectory of these developments should support a broadening of GenA.I. adoption over time.

However, the surveyed financial institutions expected to continue to rely on collaboration with external third-party service providers to customise GenA.I. solutions in the near term,

given the complexity and high costs of propriety development in-house (Figure 3.3). On the low customisation end, this can involve agent-type chatbot models that are additionally overlaid with internal, proprietary, and domain-specific knowledge. On the high customisation end, GenA.I. tools can be adapted to excel in managing complex anti-money laundering (AML) and know-your-customer (KYC) processes, and the deployment of virtual assistants in support of investment advisors.

**Figure 3.3: Main development models for enabling GenA.I.**



Source: HKIMR staff compilation based on the GenA.I. Survey.


## 3.2 RISKS ARISING FROM GenA.I. ADOPTION & MITIGATING MEASURES

### GenA.I.-related risks and challenges

When adopting GenA.I., financial institutions considered model accuracy and performance (highlighted by 95% of surveyed firms), model transparency and explainability (65%), and data privacy and security (64%) as the top three risk-management considerations (64%) (Figure 3.4). This underscores the importance of accuracy and data integrity and security in ensuring the safety and functionality of GenA.I. in an industry

where errors can carry significant financial, regulatory, and reputational costs. Relatedly, specific to the management of GenA.I. model-risk, a lack of talent to interpret and assess risks and complying with regulations also featured as key considerations. This suggests that not only are the surveyed financial institutions concerned about the risks that can come with the adoption of GenA.I., but the management of such risks can also pose significant challenges. Some interview respondents thus noted the importance of human oversight and controls at this stage of GenA.I. development as an integral part of quality assurance processes.

**Figure 3.4: Top risk management considerations when selecting GenA.I. solutions**




Source: HKIMR staff compilation based on the GenA.I. Survey.

Financial institutions were mindful of new methods of cyberattacks that leverage GenA.I. models. In response to a survey question on the three most critical cybersecurity risk that GenA.I. creates or exacerbates, 75% and 60% of survey respondents ranked 'deepfake and artificial identities' and 'phishing attacks', respectively, as among the top such risks (Figure 3.5). A number of interviewees noted that GenA.I. has enabled malicious actors to create more advanced deepfake identities for fraudulent activities, including in claims processing.

Sophisticated and personalised phishing has also been elevated with the assistance of GenA.I. technology. However, despite these concerns, only 16% of the surveyed financial institutions considered cybersecurity risk as a major barrier to GenA.I. adoption, with firms active in utilising and expanding their existing cybersecurity frameworks to mitigate the new cybersecurity challenges posed by GenA.I. technology.

Figure 3.5: Top cybersecurity risks that GenA.I. may create or worsen



Source: HKIMR staff compilation based on the GenA.I. Survey.

Third-party concentration and associated risks also featured among key risk management challenges for firms. For example, some interviewees noted that software providers often integrated GenA.I. into their products without clear disclosure, and that it was difficult to distinguish between results generated by traditional AI from added GenA.I. features. To help manage such risks, many interview respondents emphasised the need for operational resilience, alongside thorough due diligence and robust data security measures. A few interviewees also suggested that AI governance guidelines may need to be extended to technology vendors providing services to the financial services industry, to guard against the incorporation of GenA.I. into existing products such as software portfolios without their financial institutional clients being notified.

In addition, many interviewees expressed some concern around data localisation and the direction of global regulation on this issue. Currently, many GenA.I. service providers host their services outside of Hong Kong and many of the latest GenA.I. solutions are only accessible through public clouds (e.g. OpenA.I.'s ChatGPT, Anthropic's Claude). This already creates some implementation challenges. As one survey respondent shared, their head office has requested that all IT infrastructure and data related to GenA.I. be hosted locally through on-premise data centres or private

clouds, which requires transcribing the data and ensuring compliance around data sharing across jurisdictions.

### Mitigating measures


To strengthen risk management, financial institutions in Hong Kong have made solid first steps towards responsible GenA.I. adoption and development, supported by updated regulatory guidelines. There is a clear prioritisation of transparency and accountability in GenA.I. tools, alongside a strong emphasis on data protection and safeguarding customer information.

In general, overarching governance frameworks have not changed significantly, but survey and interviews indicated that a large share of financial institutions are enhancing existing procedures and guidelines to mitigate new GenA.I.-related risks. Among the surveyed financial institutions already adopting and piloting GenA.I. use cases, around half (skewed towards medium-sized firms) undertook a centralised approach to GenA.I. governance and approval processes (Figure 3.6). This was followed by a federated approach (27%), more frequently adopted by large-sized firms with often diverse operations and needs, where decision-making and implementation of GenA.I. governance and approval processes are made locally, albeit typically with centrally-issued guidelines.

Although both a centralised and a federated approach can facilitate standardisation, a centralised model can enable easier sharing of information and infrastructure across teams, and can promote consistency and coherence in data management, while a federated approach can better enable local teams to make decisions

that take into account local implementation considerations, thereby enhancing operational efficiency and responsiveness. Small-sized firms lacking resources and scale more frequently had no formal GenA.I. governance and approval processes in place or made decisions at the business unit level for small-scale adoption purposes.

Figure 3.6: Operating framework for GenA.I. governance and approval processes




Source: HKIMR staff compilation based on the GenA.I. Survey.

The interviews indicated that larger financial institutions, in part reflecting their existing governance structure, tended to adhere to group-level policy on responsible AI governance across business lines, albeit implementation and specific use case development may be further adapted in line with local regulations. In some cases, to ensure compliance of GenA.I. use both with group-level governance and standards and with local regulations, a multi-layer review process may be utilised, such as a local review process between the IT, data management, and data protection teams, in addition to a group-level review process of GenA.I. use cases and data.

Among the surveyed financial institutions currently adopting and piloting GenA.I. use cases, large-sized firms took the lead in implementing a formal model-risk management framework to assess vulnerabilities throughout the model life-cycle (Figure 3.7). Those that reported no formal assessment of model accuracy and risks was dominated by firms at the piloting and design phase and small- and medium-sized insurers and WAMs. As the technology advances, a desire by financial institutions to exploit GenA.I. for more advanced and customer-facing applications will warrant further enhancement in model-risk management.

Figure 3.7: Approach to assessment of GenA.I. model accuracy and risks

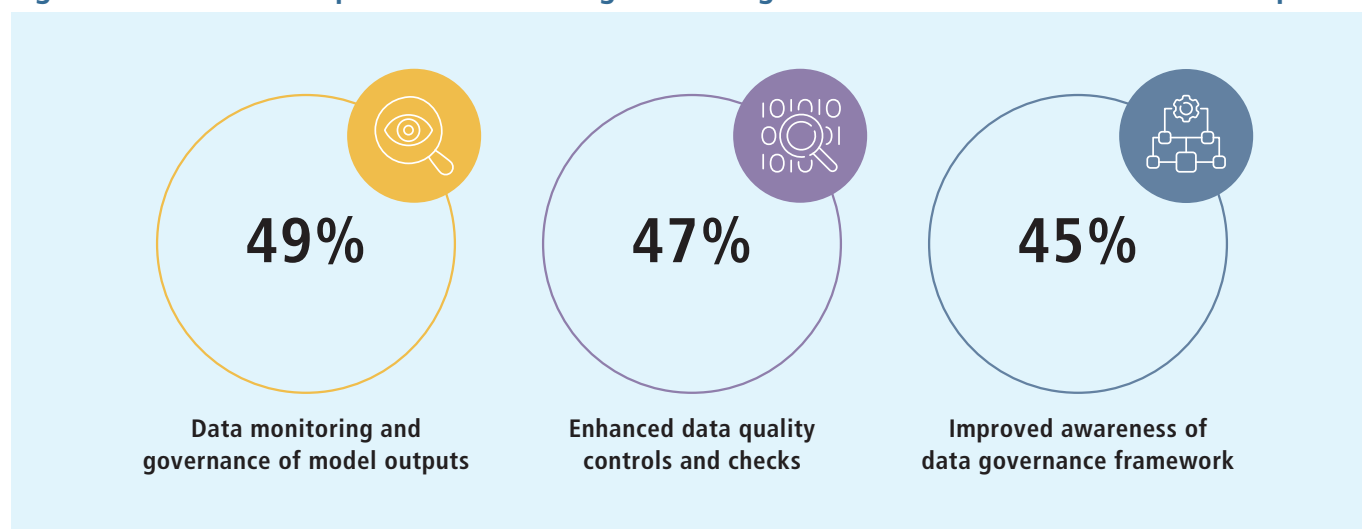


Source: HKIMR staff compilation based on the GenA.I. Survey.

Many financial institutions are already implementing or expanding existing data governance programmes to meet new requirements associated with GenA.I. adoption. A number of interview respondents noted that evolving customer expectations and a growing recognition of responsible AI principles are pointing to a need to prioritise transparency and accountability. The survey results identified several common areas for improvement, notably data monitoring and governance of model outputs, enhanced data quality controls and checks, and improving awareness of the data governance framework (Figure 3.8).

The interviews revealed that numerous firms are also addressing specifically the quality and security of unstructured data and placing a strong emphasis on data protection, especially concerning customer information through safeguards such as the inclusion of data protection and privacy experts in GenA.I. review teams, as well as content filtering to prevent information leakage. These priorities are aligned with those highlighted by regulators with respect to consumer protection, including in Hong Kong.

Figure 3.8: Areas for improvement to strengthen data governance and facilitate GenA.I. adoption



Source: HKIMR staff compilation based on the GenA.I. Survey.

Measures to mitigate cybersecurity risks stemming from the adoption of GenA.I. indicated some sectoral variation in approach, but in general pointed to an active and comprehensive approach in addressing GenA.I.-related cybersecurity challenges. Insurance firms tended to focus on practical strategies such as security awareness training to educate employees (Figure 3.9), followed by content filtering to prevent leakage of sensitive information. WAM firms tended to prioritise oversight from their board and senior management. Banks, in general placed roughly equal importance across all risk management measures, including regular security assessments to identify vulnerabilities and system-level prompt engineering.

The survey and interviews indicated that a 'human-in-the-loop' approach was considered essential, especially at this stage of GenA.I. technology maturity, and most firms are increasingly incorporating 'human-in-the-loop' considerations into their AI governance frameworks for quality assurance purposes. However, the interviewees also highlighted challenges in effectively integrating human oversight in a highly automated GenA.I. environment that can produce vast volume of outputs that humans cannot process in a short time. There were also challenges around accountability, with many interviewees concerned that senior managers accountable for decisions made by GenA.I. applications may lack the requisite technical expertise and training.



Figure 3.9: Top two strategies to mitigate GenA.I.-related cybersecurity risks

Insurance	86% Security awareness training to educate employees on the specific risks amplified by GenA.I.	71% Content filtering to prevent leakage of sensitive information
WAM	75% Board and senior management oversight of the cybersecurity strategy and framework	75% Content filtering to prevent leakage of sensitive information
Bank	71% Regular security assessments on GenA.I. solutions to identify vulnerabilities	71% System-level prompt engineering to constrain behaviour of GenA.I. applications

Source: HKIMR staff compilation based on the GenA.I. Survey.

In response to third-party concerns, GenA.I. providers interviewed are increasing the transparency of their products. They are also offering product features such as content filtering to address data privacy and security risks, and IP indemnity clauses to allay concerns of IP infringement. GenA.I. service providers interviewed have also typically established their own responsible AI frameworks based on international standards, as well as introduced a number of risk mitigation measures to facilitate responsible AI adoption. These include adopting the practice of logging, whereby all inputs, the logic behind outputs, and the outputs themselves are systematically logged, allowing users to trace results for monitoring purposes. Some vendors offer features to identify biases in datasets, to ensure that the data does not unfairly favour specific groups. And in support of explainability, some vendors enable models to clarify the reasoning behind outputs, including the sources of answers and the logic that led to those conclusions.

### 3.3 IMPLICATIONS FOR FINANCIAL INSTITUTIONS' TALENT STRATEGIES

A pre-requisite for successful GenA.I. adoption is a high-quality and robust talent pipeline. Technical skills in AI and data analytics were previously identified as among the primary skills gaps in the financial services industry in Hong Kong<sup>67</sup>. Around 80% of the surveyed financial institutions identified technical skills required in the use and development of GenA.I. as some of the top skills gaps faced by the industry. 60% of survey respondents also highlighted compliance skills as a key skills gap in supporting GenA.I. initiatives (Figure 3.10). These figures are broadly in line with another survey of CEOs reported by IBM in 2024, where 62% of respondents reported that a lack of expertise was the main stumbling block in executing their AI and automation strategy to at least a moderate extent<sup>68</sup>.

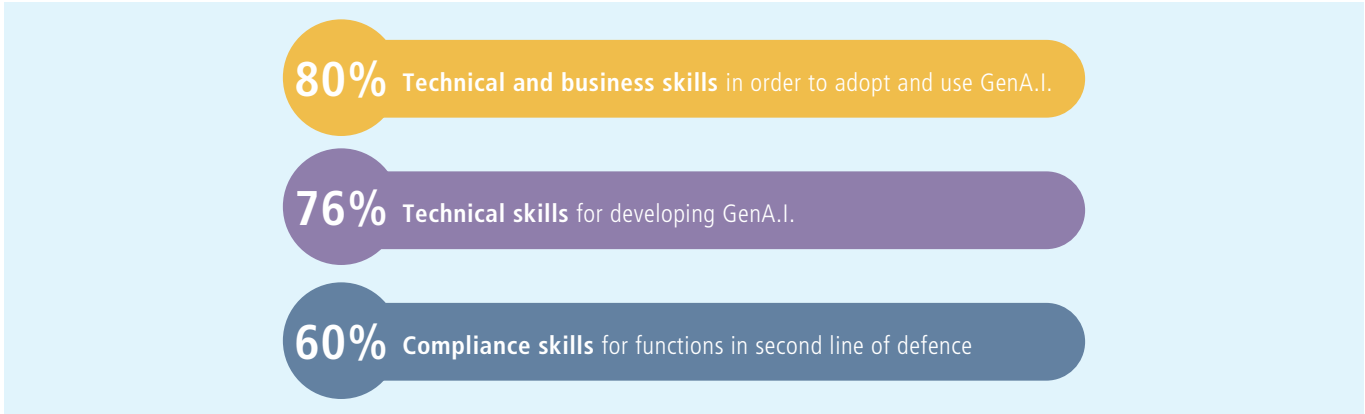
<sup>67</sup> Hong Kong Institute for Monetary and Financial Research (2023).

<sup>68</sup> ERP Today (2024).

The interviews also indicated that financial institutions faced particular challenges in acquiring risk management experts with both GenA.I. and local financial regulation expertise. This pointed to a lack of an adequate pool of talent with a combination of AI expertise, a thorough understanding or experience of financial

regulation, and experience of local financial market operations – skillsets that rarely overlap in the job market or are cultivated in current training programmes. There is also a need for more widespread technical understanding across business areas and seniority to promote an AI culture.

Figure 3.10: Top skills gaps in supporting GenA.I. initiatives




Source: HKIMR staff compilation based on the GenA.I. Survey.

Fully bridging the gap between technical execution and financial sector expertise is anticipated to take time. To help bridge skills gaps the surveyed financial institutions are opting for a combination of upskilling existing employees, hiring new talent, and establishing external partnerships, especially with GenA.I. service providers (Figure 3.11). The interviews found that financial institutions viewed a culture of continuous learning and innovation

as important in safeguarding against a rapidly evolving technology landscape, while some firms also noted the difficulty in justifying near-term talent recruitment costs given the as yet unclear return on investment for GenA.I.. As such, most interviewees placed an emphasis on upskilling existing data and technology teams, particularly in areas of data science, cloud computing and business analytics.

Figure 3.11: Talent strategies to build GenA.I. capabilities



Source: HKIMR staff compilation based on the GenA.I. Survey.

Firm size and resource availability also impacted the comprehensiveness of firms’ response to GenA.I.-related talent challenges and how they assemble teams for GenA.I. adoption. The interviews indicated that large financial institutions typically have a centralised team at the group office along with a local team to support GenA.I. development and implementation. They also typically committed significant resources and are actively assessing their operating models to decide whether to repurpose existing teams or create dedicated IT teams focused on developing GenA.I. use cases and applications. In contrast, small firms tended to leverage their current IT resources for GenA.I. adoption.

Despite these differences, a preferred and shared strategy among the surveyed financial institutions of all sizes is the formation of collaborative teams composed of both business and technology professionals to facilitate GenA.I. adoption. These teams work closely with GenA.I. service providers on application implementation projects, fostering effective partnerships across the industry. The survey and interview findings are summarised in Figure 3.12.

Figure 3.12: Workforce formations to facilitate GenA.I. adoption

	Larger Financial Institutions	Smaller Financial Institutions
GenA.I. Team Structure	<b>Centralised</b> team at group office with local teams to manage GenA.I. initiative	Typically leverage <b>existing IT</b> operating models and engage ad-hoc personnel to manage GenA.I. initiative
Resources	<b>Significant</b> resources allocated	<b>Limited</b> resources influence team structure
Operating Model	<b>Created dedicated</b> IT team or <b>repurpose</b> existing teams for GenA.I. use cases	Utilise <b>existing</b> resources for GenA.I. adoption
Top Workforce Formation Model to Facilitate GenA.I. Adoption	71% Preferred a <b>collaborative team</b> comprising both business and technology professionals <sup>1</sup>	

Source: HKIMR staff compilation based on the GenA.I. Survey.

(1) 71% of small and large firms preferred a collaborative team comprising both business and technology professionals, compared with 58% of firms across all sizes who shared this preference.

### 3.4 VIEWS ON THE REGULATORY TRAJECTORY

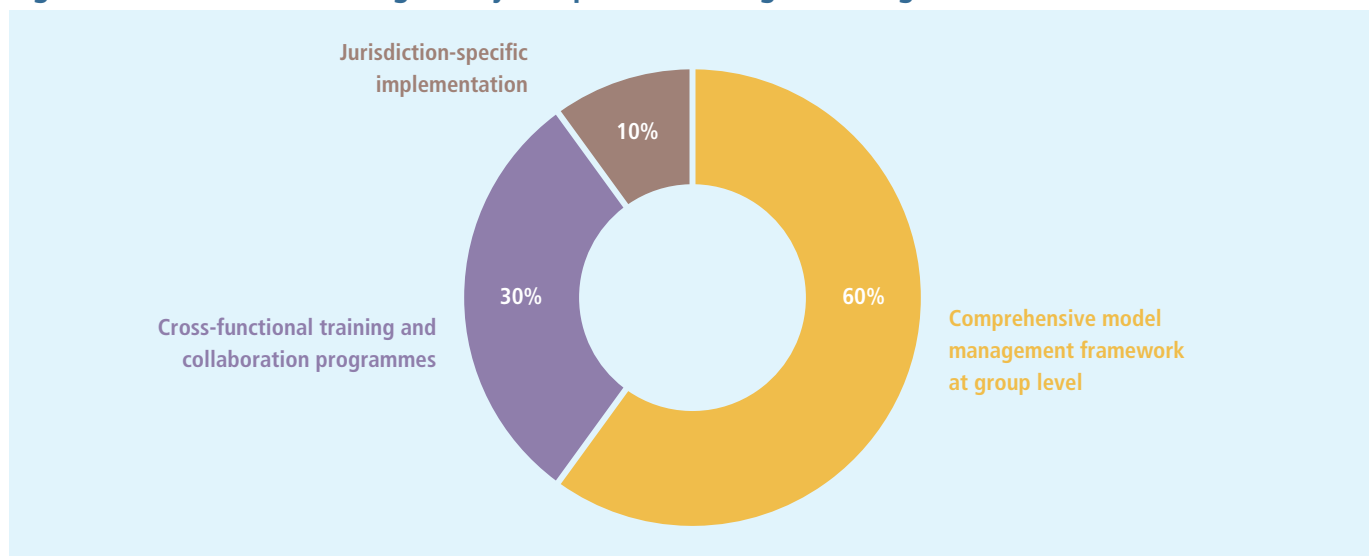
Global financial institutions that operate in Hong Kong tend to align their Hong Kong AI governance policies with international standards while adapting to local regulation. As a result, group offices often play a crucial role for larger financial institutions in defining overarching adoption principles. Indeed, in response to the survey question on how organisations address the operational challenges in aligning with GenA.I. regulation across different jurisdictions, the majority of large-sized financial institutions currently undertaking GenA.I. adoption preferred a comprehensive model management framework at the group level to ensure compliance with the most rigorous regulatory standards (Figure 3.13). This figure is comparable to that for medium-sized firms (63%), which also tended to undertake jurisdiction-specific implementation to meet local regulatory requirements.

The interviews also suggested that for larger firms, a group-level framework, usually set to meet the most stringent regulatory and compliance requirements and leveraging the group's best practices, helps to reduce costs of operating across jurisdictions. Nevertheless, cross-jurisdiction

regulatory divergence pose challenges, and respondents from firms with substantial global presence emphasised the importance of cross-jurisdiction regulatory harmonisation in reducing the costs of GenA.I. implementation. For example, one interviewee noted that due to the rapid iteration of both technology and regulation, their firm must conduct AI governance exercises on a quarterly basis to ensure compliance across the institution as a whole. Chapter 4 will expound further on these issues.

Interview respondents from both the financial services industry and GenA.I. service providers appreciated recent regulatory efforts in releasing guidelines and regulations pertaining to GenA.I. globally and in Hong Kong. Locally, the various facilitation initiatives introduced, such as the Hong Kong Monetary Authority (HKMA) and Cyberport GenA.I. Sandbox and the Cyberport AI Supercomputing Centre, and promotional events organised by regulators and the government aimed at fostering GenA.I. adoption, have also improved the local infrastructure for LLM training and deployment, and have aided firms in acquiring hardware and talent. For instance, firms have used the GenA.I. Sandbox as a way to test and better understand the regulatory boundary before moving to

**Figure 3.13: Cross-border regulatory compliance strategies of large-sized firms**



Source: HKIMR staff compilation based on the GenA.I. Survey.

production and deployment. Nonetheless, there remains a number of issues on which more comprehensive guidance could be beneficial in encouraging experimentation and early prototyping.

In particular, the promotion of guidelines on GenA.I. usage and data governance (selected by 91% of the surveyed financial institutions) and cybersecurity standards (56%), were two priority areas that regulators can help develop to further facilitate GenA.I. adoption. Forty percent of the surveyed firms highlighted the promotion of security assessments of third-party GenA.I. service providers as a priority regulatory area, with a desire among financial institutions for more standardised regulatory guidelines around the due diligence process and on-boarding of GenA.I. service providers. Interviewees also asked for further guidance on accountability between developers and end-users, especially on issues concerning data quality and data ownership as well as more detailed guidance on data localisation, residency, sharing, and traceability that is supported by practical examples, such as the implementation of data masking in model training and the integration of privacy-enhancing technologies.

In terms of facilitation and access, ongoing investment in GenA.I. hardware with wide accessibility was seen to be essential for supporting the continued adoption of GenA.I. across the financial services industry. Some interviewees cited

InvestLM as an example of government effort to train an industry-wide model that would be helpful for smaller institutions. Indeed, in response to a question about the areas where cooperative gains among the financial services industry, GenA.I. providers, and regulators are most likely to occur within the next one to two years, 75% of respondents anticipated that it would be in the development of more advanced use cases and applications, with between 50% and 60% of respondents also highlighting the development of talent, enhanced public understanding, and infrastructure.

More specifically, a number of interviewees urged the government to allocate more resources to the development of GenA.I. applications that can help resolve common industry-wide issues such as fraud detection and AML, with targeted funding or cross-industry working groups seen as potentially helpful platforms in supporting these endeavours. Others highlighted the importance of government-sponsored initiatives supporting the development of a more robust talent pool in AI and data, as well as an expansion of the deployment of the GenA.I. Sandbox, such as to encompass the insurance and WAM sectors. More generally, there was support for initiatives aimed at addressing hurdles that could slow adoption or contribute to uneven adoption and variations in risk management and talent strategies.

# Chapter 4

## The Role of Regulation

### HIGHLIGHTS:

- In the context of financial services, a robust regulatory framework in support of responsible GenA.I. adoption must address finance-specific issues, with clear guiding principles in mitigating unknown unknowns, and be enforceable.
- Across Hong Kong and other major jurisdictions, regulators have begun enhancing the AI regulations that were in place pre-GenA.I. to address new risks in ways that also cater to jurisdiction-specific needs. In Hong Kong, a strong regulatory groundwork is in place despite the technology being at its nascent stage, with the scope for further development as the technology matures.
- In the near-term, multilateral policy platforms present a viable avenue for building international consensus around appropriate supervision and enforcement mechanisms. In the longer-term, active cooperation between regulators, industry and developers through iterative dialogues can enhance both the quality and practicality of policy conclusions, as well as addressing specific GenA.I. pain points.



## 4.1 GLOBAL TRENDS IN AI REGULATION

In recognition of the transformative potential of AI across a wide spectrum of industries, the regulatory and policy landscape of AI regulation has been a dynamic process in jurisdictions worldwide. Early AI regulations were often subsumed within general data privacy and safety regulations, reflecting limited use cases and adoption before the advent of deep learning. However, as AI capabilities and the breadth of AI adoption have evolved, especially since the 2010s, the approach and focus of AI regulation have also shifted in tandem. Today, AI regulation is deemed necessary both to foster AI innovation and to ensure AI's responsible adoption and the mitigation of associated risks.

### AI regulation in the era of deep learning

The first major wave of AI regulation took place in the era of deep learning, when commercially accessible predictive AI systems began to support industry adoption. As use cases tended to be specific, with well-defined bounds around data input, algorithmic processing, and output at this stage, regulation mainly focused on data governance, bias reduction, algorithmic transparency, and strong restrictions on the use and storage of personal identifiable information<sup>69</sup>. These are reflected in early high-level frameworks, including several key circulars issued by the HKMA in 2019 that outlined high-level guiding principles on risk management for the use of AI in banking

and for consumer protection in respect of use of big data analytics and AI (BDAI)<sup>70</sup>. In some cases, high-level guiding principles were supplemented with additional legal requirements around consent, such as the comprehensive EU GDPR<sup>71</sup> that included strict transparency requirements around how obtained data will be processed, as well as requiring explicit consent from individuals prior to data processing.

### AI regulation in the era of GenA.I.

Rapid advancements in GenA.I. since 2022 are posing new challenges for AI regulation, as some familiar AI risks become more salient and novel GenA.I.-specific risks like hallucination emerge. Notably, cutting-edge GenA.I. models used training data sets that in some cases were a thousand to a million times larger than those of AI models from the deep learning era<sup>72</sup>. Governing what goes into these models, as well as the legality and ethics surrounding their inclusion, has thus become inherently more difficult. Consumer protection risks have also evolved, as reflected by timely regulatory circulars that underscore the responsibility of financial institutions for adverse impacts arising from GenA.I. use, as well as in ensuring adequate risk mitigation and due diligence. As GenA.I. technology spreads beyond simple chatbots into real-world commercial solutions, including in finance, the potential ramifications are naturally larger.

With an overarching aim of ensuring responsible GenA.I. adoption, regulators globally have begun re-examining their regulatory frameworks

<sup>69</sup> For instance, the EU's 2018 GDPR includes sections on data privacy requirements, supervisory oversight, governance structure, right to erasure (to remove oneself from a database where data has been collected), supervisory structure, restrictions on data transfers, as well as remedial requirements for data breaches. The HKMA's 2019 BDAI regulatory principles also further require consumer education on any AI's approach to processing consumer data.

<sup>70</sup> Hong Kong Monetary Authority (2019a).

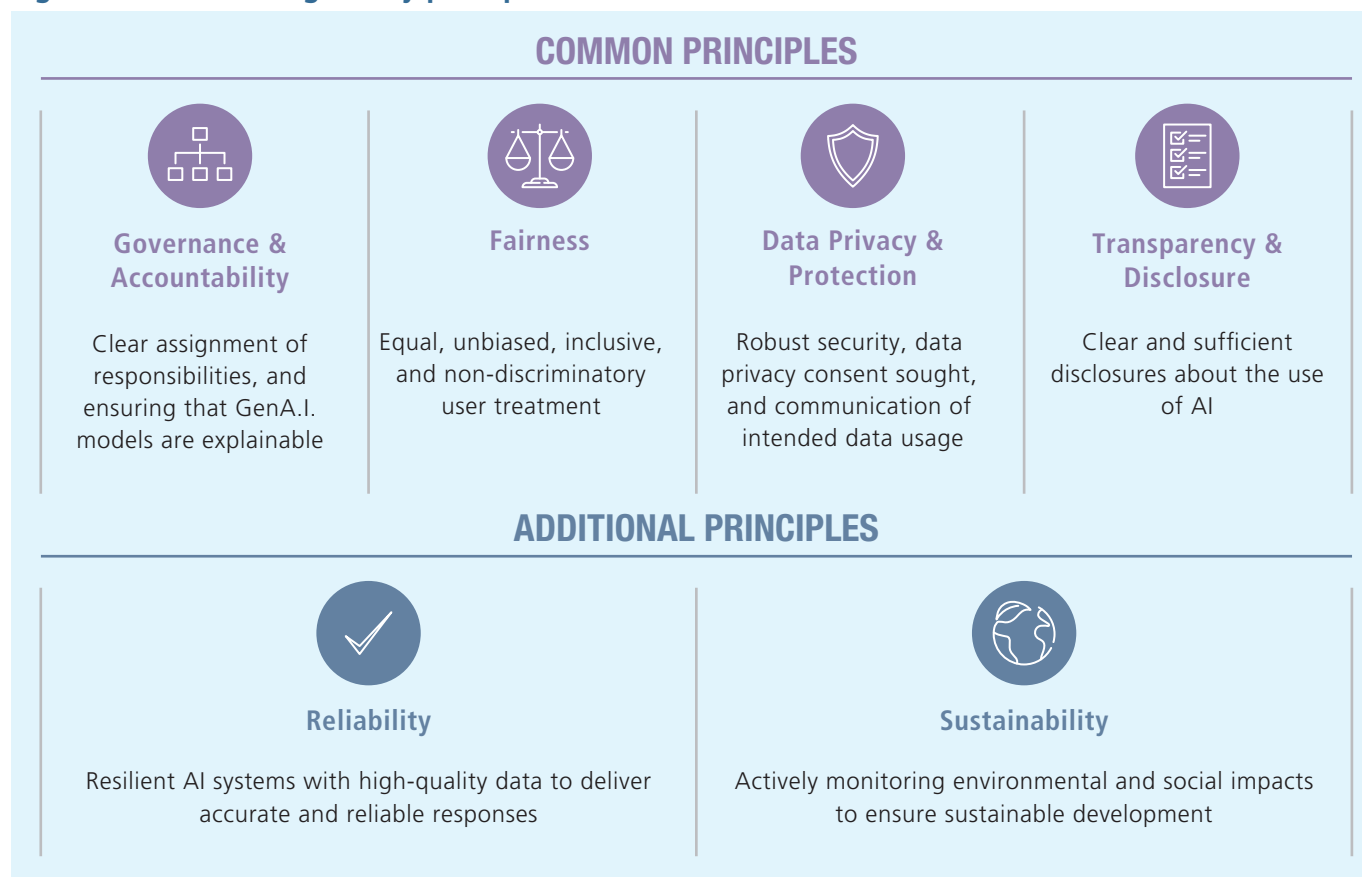
<sup>71</sup> Regulation (EU) 2016/679 (General Data Protection Regulation; GDPR).

<sup>72</sup> Epoch AI (2024; updated 2025).

governing the development and adoption of AI technologies. For now, such re-examination has not led to a wholesale overhaul of existing AI regulatory frameworks. Rather, many jurisdictions are refining existing regulations to suit the realities of GenA.I., in some cases supplementing these with additional principles to reflect specific policy

priorities (Figure 4.1)<sup>73</sup> and to manage specific risks<sup>74</sup>. A number of jurisdictions are also updating supervisory guidelines in response to concerns over human safety, and in anticipation of widening GenA.I. adoption<sup>75</sup>.

**Figure 4.1: GenA.I. regulatory principles**



Source: HKMA research.

<sup>73</sup> Hong Kong Monetary Authority (2024a).

<sup>74</sup> For example, see the *Trustworthy AI Framework* in Deloitte AI Institute (2020), which described AI systems that are (i) private (with respect to safeguarding) user privacy, (ii) transparent and explainable, (iii) fair and impartial, (iv) responsible, (v) accountable, (vi) robust and reliable, and (vii) safe and secure.

<sup>75</sup> For instance, in Hong Kong, the HKMA has provided updates to the regulatory principles for pre-GenA.I. BDAI use cases to address GenA.I.-related risks. The FSTB's dual-track approach essentially enables innovation when GenA.I. remains nascent without losing track of key risks that GenA.I. adoption poses. In Singapore, MAS' multi-phased Project MindForge adopted a public-private cooperative approach to develop GenA.I. solution alongside risk management. The US has also followed up the October 2023 AI executive order with a memorandum in October 2024, aimed in part at fostering AI safety.

## Characteristics of recent AI regulatory approaches

Across jurisdictions that have introduced regulatory refinements or updated their supervisory guidelines, current AI regulatory approaches can be broadly characterised along three spectrums: (1) the degree of codification; (2) the scope of industry coverage; and (3) the scope of sandboxing and facilitation. Although the approaches are diverse in form, these are commonly guided by an end goal of facilitating responsible adoption, and balancing between innovation and safety. Figure 4.2 summarises the three main characteristics of recent AI regulation introduced in a number of jurisdictions.

First, the **degree of codification**. At one end of the spectrum are principles-based frameworks that provide high-level operational guidelines. This more 'dynamic' approach allows regulators and industry participants more early flexibility to experiment with the mitigation of new risks as these emerge. On the other end of the spectrum are rules-based AI regulatory frameworks, such as the EU's AI Act<sup>76</sup> that entered into force in August 2024, which tend to have well-defined scopes of risk tolerance, prohibitions and penalties. Such an approach allows regulators to more quickly identify and remediate violations, but comes with less flexibility, especially if evolving technology warrants significant updates in future.

Second, the **scope of industry coverage**. Regulatory concerns around data privacy and model-related hallucination and output instability are still largely enveloped in regulatory frameworks intended for general coverage, such as in the EU and the United States. However, in some jurisdictions, additional regulatory guidance has been issued for specific industries or at specific developers and models that could give rise to systemic risks. For example, Singapore's National AI Strategy 2.0<sup>77</sup> is supplemented by the Monetary Authority of

Singapore's Veritas initiative<sup>78</sup> for the financial services industry that reflects finance-specific considerations around fair and equitable credit access and accountability. In a nod to the concept of 'too-big-to-fail', the EU and the United States have also introduced more stringent regulatory standards for the largest GenA.I. foundational models, addressing concerns that more widely adopted models (e.g. by major banks) can have more far-reaching real-world consequences if they 'fail' as a result of hallucination or model manipulation.

Third, the **scope of sandboxing and facilitation**. In recognition of the transformative potential and evolving nature of GenA.I., regulators in a number of jurisdictions have introduced sandboxes as the primary platform to facilitate technological innovation under a fail-safe environment and to promote responsible innovation. Through sandboxing, both financial regulators and the industry can iteratively uncover new risks and refine corresponding mitigation measures and the appropriate policy mix in lockstep. These sandbox environments range from being fully controlled (i.e. all testing is done without customer exposure) to real-world (i.e. testing is done with full customer exposure, and the environment is made available to the mass market). Sandboxes may be further accompanied by a broad range of innovation facilitation measures, including hackathons, financial grants, access to costly physical and digital infrastructure, and access to GenA.I. expertise.

### Hong Kong's experience

At present, Hong Kong largely adopts a principles-based regulatory framework, stemming from the Office of the Privacy Commissioner for Personal Data (PCPD)'s Guidance on Ethical Development and Use of AI. The PCPD has since

<sup>76</sup> Regulation (EU) 2024/1689 (Artificial Intelligence Act).

<sup>77</sup> Smart Nation Singapore (2023).

<sup>78</sup> Monetary Authority of Singapore (2023).

2021 advocated seven ethical principles for the development and adoption of AI, including fairness, accountability, transparency, reliability, and that AI should be beneficial<sup>79</sup>.

Finance-specific regulations and guidelines are also in place. The Financial Services and the Treasury Bureau (FSTB) has recently released a set of guidelines to promote responsible application of AI in the financial services industry<sup>80</sup>. The HKMA has issued further guidance centred on consumer protection<sup>81</sup> and in support of banks' management of prudential risks associated with AI adoption, in addition to the 2019 BDAI guiding principles. The Securities and Futures Commission (SFC) has released a circular on the risks and opportunities of GenA.I. for the securities sector<sup>82</sup>. The Insurance Authority has indicated immediate plans to implement a robust but flexible regulatory

framework to spur the fair, transparent and ethical use of AI in the insurance sector<sup>83</sup>. The Mandatory Provident Fund Schemes Authority will keep monitoring the adoption of GenA.I. in the mandatory provident fund sector and issue further or updated guidance if necessary, depending on market and regulatory development<sup>84</sup>.

Hong Kong's approach to sandboxing and facilitation has also been multi-pronged. In addition to the GenA.I. sandbox for banks<sup>85</sup>, Hong Kong's 2024–25 Budget has earmarked a HKD3 billion subsidy scheme for AI developers to access the Cyberport AI supercomputing centre for AI development over three years<sup>86</sup>, and the HKAI Lab is running a 12-month accelerator programme that provides GenA.I. start-up founders with access to frontier GenA.I. expertise globally and within the region.

**Figure 4.2: Characteristics of recent AI regulatory approaches**



Source: HKIMR staff compilation.

<sup>79</sup> Office of the Privacy Commissioner for Personal Data (2021).

<sup>80</sup> Financial Services and Treasury Bureau (2024).

<sup>81</sup> Hong Kong Monetary Authority (2024).

<sup>82</sup> Securities and Futures Commission (2024).

<sup>83</sup> Insurance Authority (2024).

<sup>84</sup> Financial Services and Treasury Bureau (2024).

<sup>85</sup> Hong Kong Monetary Authority (2024b).

<sup>86</sup> HKSAR Government (2024).

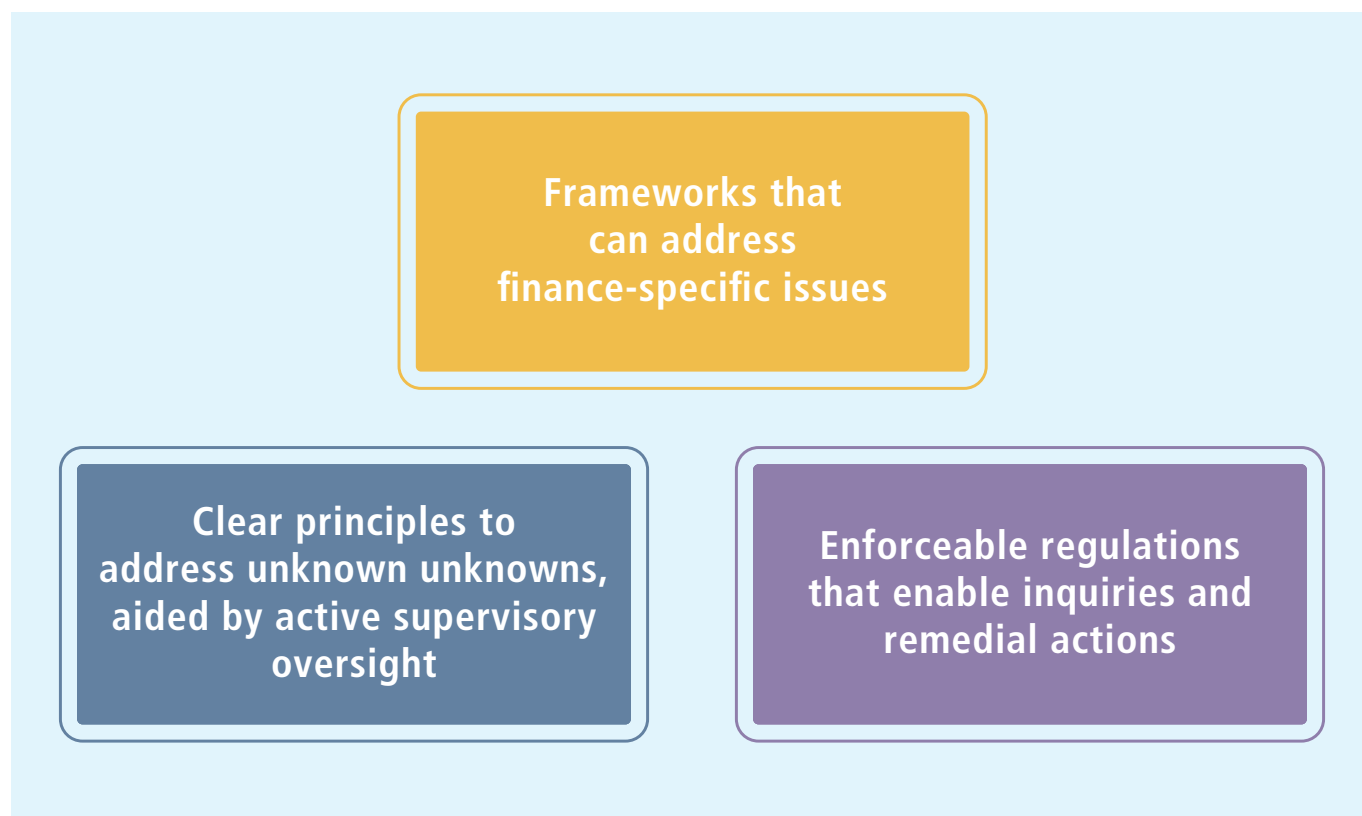
## 4.2 TOWARDS RESPONSIBLE GenA.I. ADOPTION IN THE FINANCIAL SERVICES INDUSTRY

Regulatory actions around AI and GenA.I. in most jurisdictions thus far remain fairly broad and dynamic, as policy aims to strike a balance between mitigating risks and encouraging innovation. However, themes such as accountability, model transparency, data privacy, and human safety have been increasingly highlighted in recent discussions on AI<sup>87</sup>, with human safety a top emerging area of concern<sup>88</sup>. Industry experts have also emphasised the importance of AI systems being fair and unbiased,

which poses the challenge of how to pre-build these characteristics into GenA.I. models' tuning parameters<sup>89</sup>.

For the financial services industry, a robust regulatory framework for responsible GenA.I. adoption should arguably include several key factors (Figure 4.3). First, given the importance of a stable financial system for the efficient allocation of resources and for economic growth, the financial services industry is likely to require regulation to address finance-specific issues. Second, clear guiding principles can play a key role in addressing unknown unknowns, aided by active supervisory oversight as the technology evolves. Third, these measures need to be enforceable.

**Figure 4.3: Key factors underpinning responsible GenA.I. adoption**



Source: HKIMR staff compilation.

<sup>87</sup> Al-kfairy et al (2024).

<sup>88</sup> Hagendorff, T. (2024).

<sup>89</sup> See for instance Deloitte AI Institute (2020).

## Regulatory adequacy in addressing finance-specific issues

Globally, the regulation of GenA.I. in the financial services industry tends to ride on broader overarching AI frameworks with a focus on principles-based guidelines rather than legally binding legislation, with only some jurisdictions, chiefly the EU, moving ahead with a comprehensive rules-based framework. For now, financial regulators' attention regarding the use of GenA.I. in the financial services industry remains largely focused on consumer protection. As GenA.I. adoption broadens, regulatory considerations over market access and market concentration could also increase if only the largest financial services industry players are able to develop and implement GenA.I. use cases. Ensuring that the finance workforce can adapt to the GenA.I. era in a gradual manner will also be important in jurisdictions with a large financial sector.

In Hong Kong, financial regulators have either taken directed steps to address finance-specific GenA.I. risks and emerging concerns, or are actively monitoring the adoption of GenA.I. by financial institutions and stand ready to issue further or updated guidance if necessary. For instance, to ensure consumer protection, the HKMA issued further guidelines in August 2024 on the use of GenA.I. by banks in customer-facing applications<sup>90</sup>, in addition to the 2019 BDAI guiding principles, on governance and accountability, fairness, transparency and disclosure, and data privacy and protection. The HKMA has also provided guidance to help banks manage prudential risks associated with AI adoption more generally. The recent HKMA report on GenA.I.<sup>91</sup> also highlighted the importance of additional AI regulatory principles in finance, namely ensuring that AI systems are resilient and use high-quality data, and managing and monitoring the environmental and social footprints of AI. The FSTB's recent guidelines for responsible adoption of AI also recognised the importance of talent development, and highlighted the HKMA's efforts to nurture AI-related talent in banking.

## Regulatory adequacy in addressing unknown unknowns

While rules-based frameworks tend to guard well against known risks, broad high-level principles can arguably better guide the conduct of developers and industry users in managing unknown unknowns. For instance, the United States' blueprint for an AI Bill of Rights released in October 2022 covered elements of human safety, protection from algorithmic discrimination, and built-in data privacy, which were subsequently reflected in later executive orders. Similar trends have also been observed in the United Kingdom, where principles covering transparency and explainability of AI models, fairness of outcomes and clear lines of accountability across the AI life-cycle serve as the basis for longer-term regulation and innovation. Financial institutions' responsibilities maps<sup>92</sup>, derived from microprudential regulation, can also be used to incentivise responsible decisions related to GenA.I. applications.

In Hong Kong, the use of sandboxes has historically provided scope for the benefits and risks of new technology in financial services and the potential unknown unknowns to be studied within a controlled environment, while allowing regulators and the industry to iteratively build and refine the appropriate safeguards. The inaugural cohort of the GenA.I. Sandbox incorporates 15 use cases from 10 banks and 4 technology partners, selected from 40 proposals, with technical trials expected to continue through mid-2025<sup>93</sup>. The GenA.I. Sandbox, at present, focuses on GenA.I. risk management, anti-fraud measures, and customer experience, with the explored use cases ranging from augmenting credit assessment and fraud detection through automated processing of unstructured data to GenA.I.-powered chatbots handling personalised and complex enquiries.

<sup>90</sup> Hong Kong Monetary Authority (2024).

<sup>91</sup> Hong Kong Monetary Authority (2024a).

<sup>92</sup> Responsibilities maps detail the roles and responsibilities of various stakeholders within individual financial institutions. They ensure clarity in governance, compliance and risk management by clear identification and assignment of accountability when and if failures occur. The UK FCA, for instance, requires regulated firms to maintain such documents.

<sup>93</sup> Hong Kong Monetary Authority (2024b).



In mitigating emerging and novel risks, the HKMA's recent circular on consumer protection further emphasised a 'human-in-the-loop' approach to prevent hallucination and polymorphic errors from spilling over to real-world use, as well as to ensure that decision makers are held accountable for responsible use of GenA.I., while the SFC's recent circular underscored the need for extra risk mitigation in high-risk GenA.I. use cases<sup>94</sup>.

### Regulatory adequacy in ensuring enforceability

To be enforceable, regulation needs to be empowered by the respective jurisdictions. Only a handful of jurisdictions explicitly enable this. For example, the EU's AI Office has enforcement powers to fine violators of the AI act for a substantial EUR35 million or 7% of global turnover<sup>95</sup>. This follows on the EU's track record of enforcing large fines on global technology firms for data-related violations under the pre-existing GDPR. Given growing interest in GenA.I. adoption by the financial services industry, a clear roadmap towards enforceability in the medium-term when GenA.I. development and adoption gains sufficient momentum may be helpful.

While the various guidelines issued by Hong Kong regulators are not law, legal experts opine that they provide a sufficiently clear ground for enquiries, hence action, to be pursued in instances of breaches<sup>96</sup>. Separate data protections for GenA.I. have also been introduced by the PCPD to manage unique GenA.I.-specific data risks, especially in the training process<sup>97</sup>, as well as to ensure data privacy and security when GenA.I. tools are used at the workplace<sup>98</sup>. This gradual scoping of regulation is part and parcel of the iterative process by which authorities and the industry learn the best ways to effectively regulate the technology in the long run without stifling the current dynamism of GenA.I..

Overall, a strong regulatory framework is in place in Hong Kong, and regulation should continue to maintain a sustainable balance between supporting innovation and GenA.I. adoption, while keeping a keen eye on emerging risks.

## 4.3 REGULATORY HARMONISATION

### Challenges for effective regulation across borders

Global financial institutions, such as those that operate in Hong Kong, a premier international financial centre, typically undertake substantial cross-border business activities. This makes managing different regulatory requirements across jurisdiction inevitable and potentially a challenge both from the perspective of financial institutions needing to be GenA.I. compliant in multiple jurisdictions and from the perspective of financial regulators.

For financial regulators, growing cross-jurisdictional divergence in GenA.I.-related financial regulation could potentially lead to the formation of pockets of systemic risks (Chart 4.4). For instance, industry users could simply choose to base their GenA.I. operations and toolkits in jurisdictions with lower or more liberal regulatory standards. A recent multilateral roundtable on GenA.I. in finance noted the need for national regulators to constantly assess domestic regulatory capabilities to enable effective global coordination, as well as the risk, of regulatory divergence between various sectors of the financial services industry, e.g. different stringencies between banking and asset management<sup>99</sup>. Indeed, one of the minimum characteristics that the World Economic Forum has proposed for AI regulation is to bridge regulatory differences<sup>100</sup>. The EU's AI Act also emphasises that EU-level sandbox arrangements should be supplemented by national-level initiatives<sup>101</sup>.

<sup>94</sup> Securities and Futures Commission (2024).

<sup>95</sup> Article 99 of Regulation (EU) 2024/1689 (Artificial Intelligence Act).

<sup>96</sup> Mayer Brown (2024).

<sup>97</sup> Office of the Privacy Commissioner for Personal Data (2024).

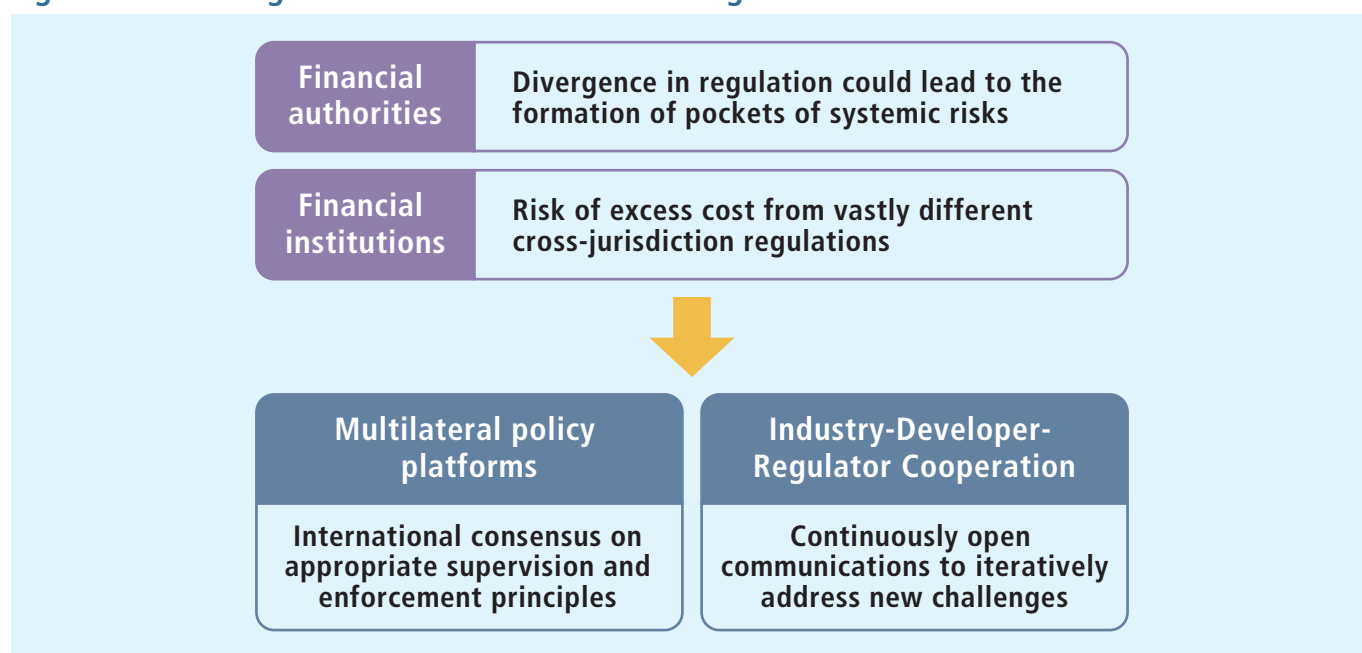
<sup>98</sup> Office of the Privacy Commissioner for Personal Data (2025).

<sup>99</sup> OECD-FSB Roundtable on Artificial Intelligence (AI) in France (2024).

<sup>100</sup> World Economic Forum (2023b).

<sup>101</sup> Article 57 of Regulation (EU) 2024/1689 (Artificial Intelligence Act).

Figure 4.4: Challenges and measures for effective regulation across borders



Source: HKIMR staff compilation.

Financial institutions, meanwhile, could incur sizeable operational costs when managing compliance across multiple jurisdictions with different regulatory requirements (Figure 4.4). Multinational financial institutions seeking to adopt GenA.I. may find it necessary to set up duplicating GenA.I. units across multiple jurisdictions. Thus, any substantive efforts at further harmonising the regulatory standards imposed on GenA.I. development and adoption stand to benefit global financial services industry players in the long run, especially in making the process and outcomes of compliance more targeted.

### The role of multilateral policy platforms

A low-hanging fruit is to capitalise on existing multilateral policy platforms and global standard-setting bodies (SSBs) (Figure 4.4), and build towards an international consensus of appropriate supervision and enforcement mechanisms

that can also reduce the scope for regulatory arbitrage. Multilateral policy platforms and SSBs (Figure 4.5) can offer a workable baseline that is broadly agreed on by major jurisdictions, hence ensuring that localised frameworks are in principle aligned with those elsewhere. Moreover, national authorities can leverage such cross-border platforms to develop policy insights, as well as policy expertise in the GenA.I. domain. A working parallel is the global minimum tax negotiated by the OECD to mitigate tax competition and tax avoidance by large multinational enterprises. Such a process, while not immediate, can help ensure that the frameworks or policy principles agreed upon have the necessary buy-in from major AI jurisdictions, and are therefore effective.

Figure 4.5: Selected multilateral policy platforms and recent AI/GenA.I.-related initiatives

Bank for International Settlements (BIS)	The BIS promotes global monetary and financial stability through international cooperation. The BIS Innovation Hub currently engages in several AI-related projects, such as Project Aurora, which uses AI to combat money laundering. Beyond research on economic implications of AI and the role of regulation, both independently and through its hosted committees including the BCBS, the BIS has been discussing AI and GenA.I.-related issues and policy considerations with national authorities through various working groups and high-level bimonthly meetings.
Basel Committee on Banking Supervision (BCBS)	The BCBS is the primary global standard-setter for the prudential regulation of banks. The BCBS recently published a report on the ongoing digitalisation of finance and the implications for banks and supervisors, which highlighted several GenA.I.-related risks and potential mitigating measures.
Financial Stability Board (FSB)	The FSB monitors and makes recommendations about the global financial system. The financial stability implications of AI are part of the priority areas in FSB's 2025 work programme, with a publication on the vulnerabilities associated with the use of AI in finance planned by end-2025.
International Association of Insurance Supervisors (IAIS)	The IAIS is the global standard setter for supervision of the insurance sector. The IAIS' workplan for 2025-26 noted AI as a strategic theme in global insurance supervision, which builds on a public consultation on managing AI-related risks amongst insurers conducted in 2024-25 and a thematic review of AI-related guidance across jurisdictions in 2023.
International Monetary Fund (IMF)	The IMF works to achieve sustainable growth and prosperity for member countries by supporting economic policies that promote financial stability and monetary cooperation. The IMF is investigating AI's broader impact on economies and societies by gathering global knowledge through surveillance activities, and by convening key actors to share successful policy responses, foster international consensus and harmonise regulations.
International Organization of Securities Commissions (IOSCO)	The IOSCO is recognised as the global standard setter for financial markets regulation. IOSCO's 2025 work programme includes a report on current and near-term AI use cases, as well as issues, risks and challenges that authorities face when considering policy responses, which builds on its 2021 guidance for the wealth and asset management sector using AI. IOSCO further plans to develop recommendations to its member authorities in addressing AI-related risks.
Organisation for Economic Co-operation and Development (OECD)	The OECD works to establish evidence-based international standards and to find solutions to social, economic, and environmental challenges. The OECD produced the first intergovernmental AI standard, first adopted in 2019 and updated in 2024, to reflect recent innovations in GenA.I.. The OECD has also been facilitating discussions on common GenA.I.-related policy priorities as part of the G7 Hiroshima AI Process, and promotes responsible AI use and development through the Global Partnership on AI (GPAI).

Source: HKIMR staff compilation.

**The role of active cooperation between developers, industry users and regulators**

Regulating GenA.I. effectively is a high-frequency iterative process, given the technology's nascent state and continued innovation. Regulations both domestically and internationally can benefit from the participation of industry users and GenA.I. developers in the same policy dialogue to enhance both quality and practicality of policy conclusions

at these platforms (Figure 4.4). This is particularly key for critical sectors with a long tradition of iterative regulation, such as finance. Cooperation can occur on two levels: general regulatory and development issues that cut across industries; and specialised dialogues to address the needs and pain points of the financial services industry. The latter may encompass the various financial services industry regulators, financial institutions, and GenA.I. providers.

# Chapter 5

## Considerations for Fostering Responsible GenA.I. adoption in Hong Kong

### HIGHLIGHTS:

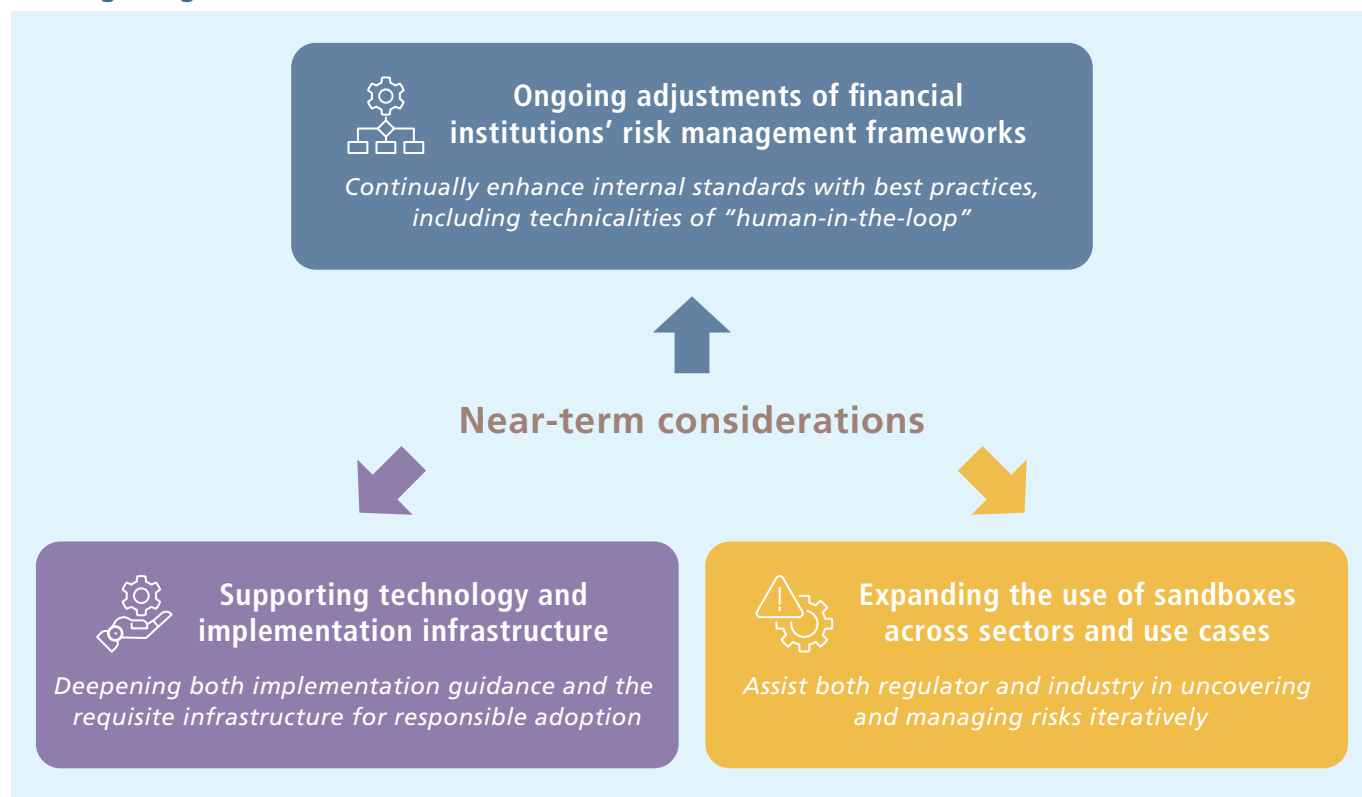
- Financial institutions are adopting GenA.I. responsibly, supported by up-to-date guidelines. In the near term, enhancements of risk management frameworks should continue, supported by an expansion of access to sandbox and facilitation initiatives, additional technical implementation guidelines and supportive implementation infrastructure.
- Ongoing reskilling and upskilling initiatives can bridge GenA.I. talent gaps and mitigate job displacement risks. Financial institutions can leverage collaborative (T-shaped) team structures and external partnerships, and adapt and expand existing technical skills infrastructure to meet growing GenA.I. demands.
- Increased tripartite collaboration between regulators, industry and developers, greater cross-jurisdiction regulatory harmonisation, and continued development of physical AI infrastructure can contribute to a level playing field over the long run.

## 5.1 NEAR-TERM SUPPORTIVE MEASURES

Financial institutions in Hong Kong that are actively piloting or adopting GenA.I. applications have made solid first steps towards responsible GenA.I. adoption and development, with the incorporation of GenA.I.-related risk considerations into internal risk management and governance frameworks,

and supported by updated regulatory guidelines. As the technology matures and as GenA.I. use cases increase in complexity, further adjustments and enhancements to firms' internal governance controls and in the general adoption environment are arguably warranted to ensure that adoption can continue in a safe and responsible manner (Figure 5.1).

**Figure 5.1: Near-term considerations for supporting responsible GenA.I. adoption and development in Hong Kong**



Source: HKIMR staff compilation.

### Ongoing adjustments of financial institutions' risk management frameworks

Financial institutions should continually review and enhance existing internal standards with best practices globally, to ensure that emerging risks are comprehensively addressed<sup>102</sup>. The development of comprehensive risk management strategies

can be supported by additional guidance aimed specifically at the financial services industry and at addressing specific sectoral risks and concerns, such as the six principles highlighted in HKMA's recent research<sup>103</sup> aimed at promoting responsible GenA.I. adoption while safeguarding end-user interests.

<sup>102</sup> Financial Stability Board (2024a).

<sup>103</sup> Hong Kong Monetary Authority (2024a). The six principles are: Governance and accountability, Fairness, Data privacy and protection, Transparency and disclosure, Reliability, and Sustainability.

A number of interview respondents also highlighted the practical operation of ‘human-in-the-loop’ as a particular challenge in the implementation of GenA.I. models. The AI regulatory principles issued by the HKMA already emphasised this issue, while the BCBS recently underscored the importance of human judgement alongside appropriate risk management and governance structures in supporting responsible AI adoption and financial decision-making. This suggests urgency in further developing the principles and technical implementation details behind ‘human-in-the-loop’ to more effectively identify and mitigate unpredictable behaviour exhibited by GenA.I. systems, such as hallucination and output errors due to probabilistic model responses. Ultimately, even if GenA.I. takes the entire product end-to-end (from raw input to analytics to advisory), the final decision and, hence, responsibility should lie with an authorised individual or team of individuals empowered to enforce or override the GenA.I. model’s decisions. This also facilitates a clear assignment of responsibilities and liabilities in GenA.I. use cases, such that innovation is balanced with the appropriate risk-reward matrix.

### Expanding the use of sandboxes across sector and use cases

An expansion in access to GenA.I. sandboxes to facilitate the testing of a broader set of use cases and sector-specific issues could help support broader adoption. The interviews with Hong Kong financial institutions indicated a strong demand from a sizeable share of insurers and WAM firms for sectoral-focused sandboxes. As current GenA.I. adoption tend to be concentrated among larger financial institutions, an expansion of access to GenA.I. sandboxes may also help support a level-playing field, promote competition, and reduce concerns around market concentration and market integrity.

### Supporting technology and implementation infrastructure

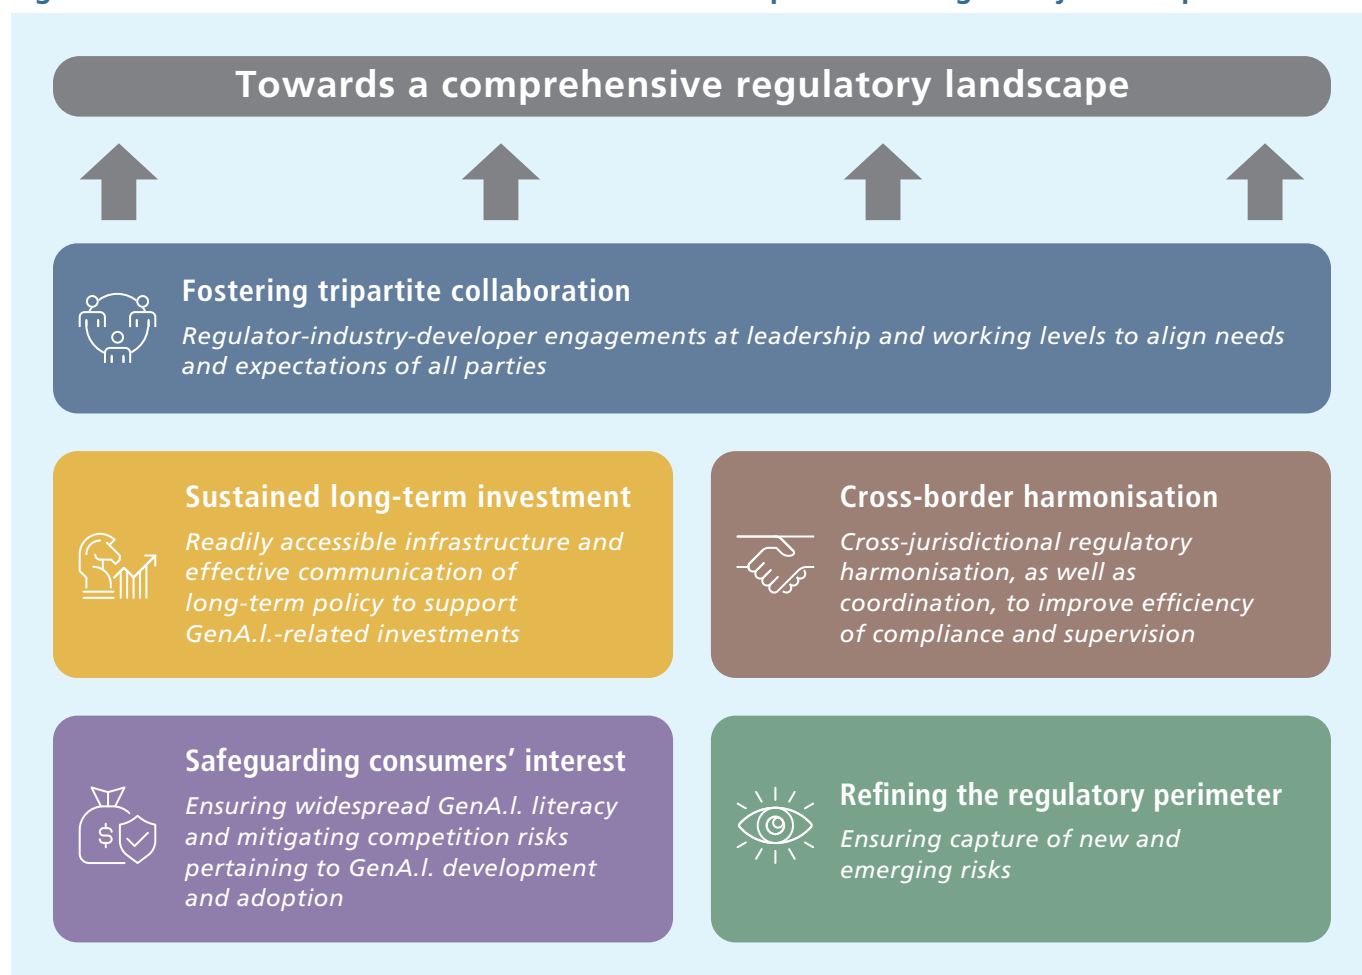
GenA.I. adoption and development requires substantial upfront capital investment in technology and implementation infrastructure. The role of public funding in areas that are less likely to receive direct private sector investment, such as data and computing infrastructure, is also crucial in supporting GenA.I. development and adoption needs<sup>104</sup>. For instance, data localisation requirements may require mature within-jurisdiction cloud infrastructure or high-powered on-premise servers that are difficult to implement, even for the latest baseline GenA.I. use cases and especially for smaller financial institutions. The facilitation of more locally developed GenA.I. tools such as HKGAI V1, including through private investments and capital markets, can also spur innovation and development in the local ecosystem.

## 5.2 MEDIUM-TERM CONSIDERATIONS TOWARDS A COMPREHENSIVE REGULATORY LANDSCAPE

Putting in place a comprehensive regulatory landscape over the medium-term with adequate harmonisation is key to facilitating responsible GenA.I. adoption. Achieving this is likely to require a multi-pronged approach along five pillars – tripartite collaboration and engagement, cross-border harmonisation, sustained long-term investment, refining the regulatory perimeter, and safeguarding consumer interests (Figure 5.2).

<sup>104</sup> Brollio, F., Dabla-Norris, E., de Mooij, R., et al (2024).

Figure 5.2: Medium-term considerations towards a comprehensive regulatory landscape



Source: HKIMR staff compilation.

### Fostering tripartite collaboration through multi-tiered engagement

There are practical benefits to nurturing an ecosystem where three communities – regulators, industry players, and GenA.I. developers – widely and actively engage in a tripartite setting across levels and sectors. Leadership-level conversations can serve to align the high-level needs and expectations of all three communities. More informal and iterative working groups are equally important for ensuring open, practical, and constructive discussions on the regulatory trajectory. For example, a dedicated agency that

mobilises the private sector, academia, and other stakeholders could help track AI developments and use, as well as develop a broader accountability framework for firms in the AI value chain<sup>105</sup>. Announced in the 2025/26 budget, Hong Kong's new AI Research and Development Institute, backed by HKD1 billion of initial investments, could be well suited to serving this role. In the long term, a sustained culture of cooperation, also including academia, could be useful in cultivating continuous GenA.I. research and development that is market- and policy-relevant.

<sup>105</sup> Brollio, F., Dabla-Norris, E., de Mooij, R., et al (2024).



### Fostering cross-border regulation harmonisation

Cross-border regulatory cooperation will be essential, especially as financial institutions in Hong Kong tend to have a strong global presence. As current jurisdiction-specific regulations pertaining to GenA.I. tend to diverge, some degree of harmonisation in regulatory principles may have to be considered, subject to jurisdiction-specific needs. The FSB also recently proposed international and cross-sectoral coordination, including via the FSB itself, to enhance regulatory and supervisory capabilities for overseeing policy frameworks related to AI in the industry. Cross-jurisdictional harmonisation may improve the efficiency of compliance among GenA.I. users in the industry, especially those with large global presence. Existing governance committees that are armed with both business and technology experts, with guidelines and oversight made at the group level, also provide some blueprint for cross-jurisdiction regulatory harmonisation from the industry's perspective.

### Sustained long-term investment in digital infrastructure

Given an evolving technology and operating backdrop, the continued development of physical infrastructure over the medium-term is needed. Hong Kong's recently launched Cyberport Artificial Intelligence Supercomputing Centre (AISC) – whose computing power could more than double to 3000 petaFLOPs by early 2026, equivalent to processing 10 billion images per hour<sup>106</sup> – provides one such infrastructure for those in the AI ecosystem, including firms, government departments, academia and research institutes. To facilitate access, a HKD3 billion subsidy scheme for AI developers over three years has also been launched<sup>107</sup>. As demand for GenA.I. rises, other critical physical infrastructure such as cloud hosting servers and advanced GPUs may need to expand in lockstep. Alongside public sector investments, sustained and close communication of long-term policy timelines can anchor the expectations

of interested GenA.I. adopters and developers, allowing sufficient gestation for long-term investments in GenA.I. technology, and its eventual adoption.

### Further refining the regulatory perimeter

The majority of financial institutions in our survey are leveraging external GenA.I. service providers, and intend to continue to do so in the longer term. Such a concentration of AI service providers catering to the financial services industry and the broader economy may necessitate financial stability considerations. Indeed, BIS research raised the question of whether there should be greater oversight of GenA.I. service providers by financial regulators<sup>108</sup>. Beyond the prevalent approach of managing such risks through due diligence requirements expected of financial institutions<sup>109</sup>, more practical discussions may be warranted around how best to incentivise industry users and developers to make safe, ethical and responsible decisions when adopting and/or developing GenA.I. through existing regulation.

### Safeguarding consumer interest as GenA.I. adoption broadens

To safeguard long-term consumer protection, there may be value in expanding the scope of financial literacy programmes targeted at the general public to include GenA.I.-related issues. Such programmes may cover how GenA.I. works in relation to financial products, consumer rights and protection measures offered by law, and mitigating steps that consumers can take to further mitigate risks at the individual level, including the provision of sensitive data. Another key aspect of consumer protection that entails addressing competition risks that arise from high barriers of entry. The FSB has proposed regulators consider the effects of regulatory and supervisory frameworks on a level playing field not only between sectors, but also between established firms and new entrants in the industry. Sandboxes and technical assistance explicitly aimed at smaller

<sup>106</sup> *The Standard* (2024).

<sup>107</sup> *HKSAR government* (2024).

<sup>108</sup> *Bank of International Settlements* (2024).

<sup>109</sup> *Ibid.* (2024).

institutions may be one possibility, alongside antitrust surveillance that explicitly accounts for inequity in GenA.I. access. Similar to the rollout of digital banks, consortium-based explorations may allow smaller players to assume a fair share of risks when adopting GenA.I. commensurate with potential yields.

### 5.3 FOSTERING A ROBUST TALENT PIPELINE FOR THE GenA.I. ERA

The survey and interviews suggest that the largest GenA.I.-related talent gaps in the financial services industry are in the demand for Explainers (those who design interfaces that facilitate GenA.I. use) and for Sustainers (those who ensure ethical and effective usage). Financial institutions are already considering various talent strategies to build GenA.I. capabilities, predominantly through upskilling of existing employees, followed by some hiring of new talent with GenA.I. expertise. This approach is in line with earlier findings<sup>110</sup> that purely relying on external hiring to address skills gaps may not always be feasible as the supply of particular talent could be rare, especially in a rapidly changing technology era. As such, as GenA.I. adoption broadens and picks up pace, well-structured and ongoing upskilling and reskilling initiatives across the industry need to be considered.

#### Upskilling and upgrading team structures to bridge talent gaps

According to the survey, financial institutions identified the technical and business skills necessary to adopt and utilise GenA.I. as the most significant gap in supporting GenA.I. initiatives. As GenA.I. increasingly takes more job duties that were previously performed by humans, the concerns about job replacement are not negligible<sup>111</sup>. A comprehensive upskilling initiative is therefore needed to enable employees to work effectively with GenA.I., while also equipping them

with other essential skill sets, such as domain expertise. In particular, to remain competitive and adapt to continuous change, employees can consider ongoing learning initiatives as a means to mitigate job displacement risks and address the potential performance gap that may arise from job displacement. Management can consider including training and reskilling in business strategies<sup>112</sup>, rather than treating them merely as a training programme. Training and reskilling is considered a holistic initiative that requires the alignment of workforce planning and recruitment. Factors such as clear role definitions and leadership support should help ensure successful integration of reskilled employees in the new roles, thus fostering long-term business growth.

In the context of GenA.I. skills and other essential skill sets, the concept of the T-shaped team structure<sup>113</sup> – geared towards being a ‘finance–technology innovator’ – also becomes relevant (Figure 5.3). Indeed, our survey and interviews revealed that institutions of all sizes already have collaborative teams comprising business and technology members to some extent, albeit with some variations in team structure across different institution sizes. Notably, the size and structure of the T-shaped team can be small, informal, or project-specific in the early stages. Moreover, T-shaped teams are not a one-size-fits-all solution; financial institutions can further develop their own T-shaped team structures that best meet their needs and capabilities.

Successful adoption of new technology in financial services relies on effective business–technology integration. The concept of T-shaped teams<sup>114</sup> brings together professionals from both technology and finance backgrounds, as well as ‘finance–technology innovators’ who serve as a bridge between finance and technology professionals. The collaborative nature of the T-shaped team structure fosters close links and promotes productive communication within

<sup>110</sup> Hong Kong Institute for Monetary and Financial Research (2023).

<sup>111</sup> Harvard Business Review (2024b).

<sup>112</sup> Harvard Business Review (2023c).

<sup>113</sup> Hong Kong Institute for Monetary and Financial Research (2023).

<sup>114</sup> Hong Kong Institute for Monetary and Financial Research (2021) and Cao, Larry (2021).

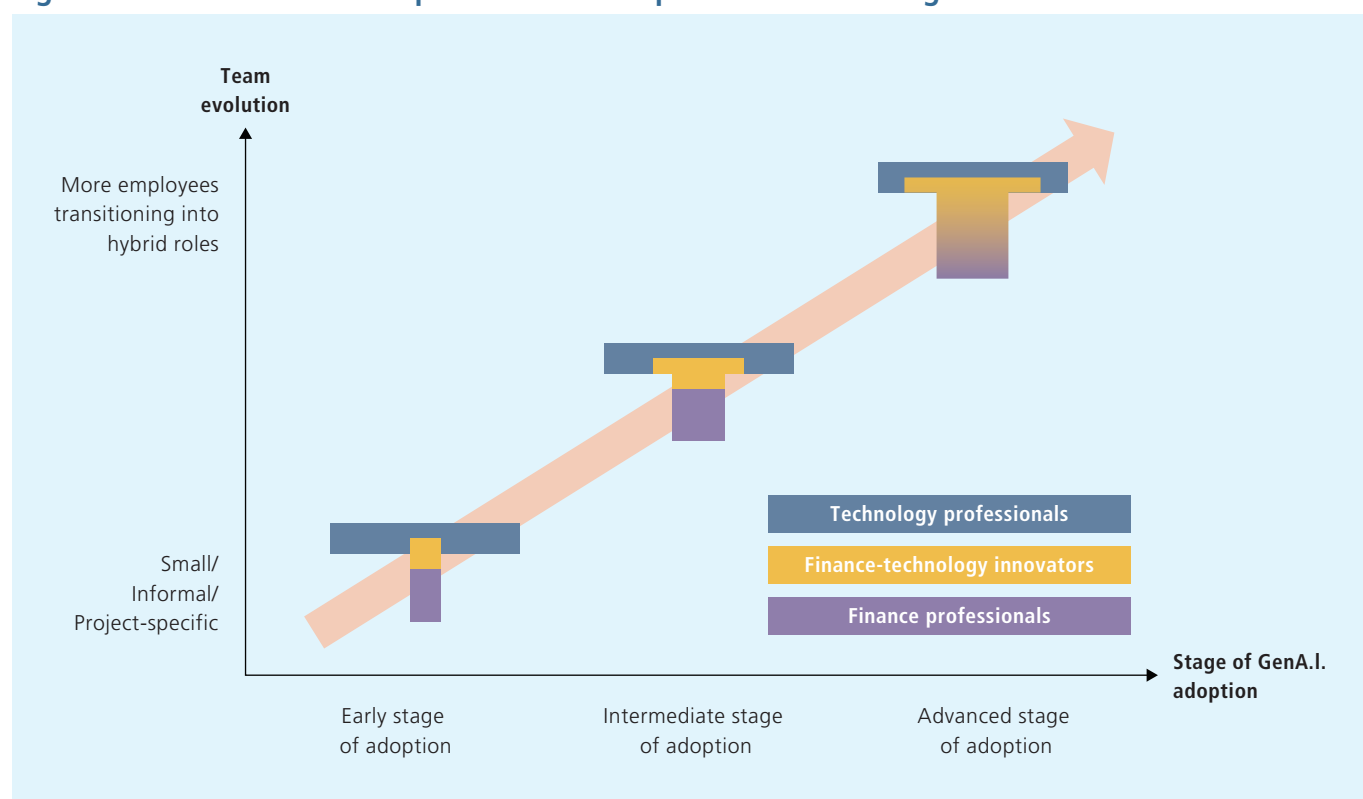
the team through a shared team space and language. In the context of adopting GenA.I. for business use, finance professionals can define business problems, while technology professionals can explain GenA.I. use cases. Meanwhile, finance-technology innovators can translate requirements between the two parties, facilitating effective communication and promoting trust within the team.

In the process of technological transition, the T-shaped team is also likely to evolve. Cross-disciplinary interaction fosters a deeper understanding of each discipline by the other, enabling finance professionals to gain technological insights and technological experts to gain a deeper understanding of financial knowledge and principles. As a result, a mutual

understanding is likely to emerge, blending both types of professionals into ‘finance-technology innovators’. Therefore, we expect the number of innovators to gradually increase (Figure 5.3: orange bars). The T-shaped structure, as a collaborative arrangement, naturally fosters upskilling and reskilling initiatives, as professionals are exposed to new knowledge, tools, and methodologies through their interactions with one another.

The T-shaped team, particularly in the intermediate and advanced stages of adoption, offers opportunities for career advancement, as the value of employees increases when they transition into hybrid roles. Employees with hybrid skill sets are likely to be more adaptable in meeting the evolving market and organisational demands.

**Figure 5.3: Evolution of T-shaped teams in the process of technological transition**



Source: HKIMR staff compilation

### A collaborative path to growth

According to the interviews with financial institutions, there is recognition of the importance of collaboration with entities such as educational institutions, which is beneficial for cultivating a talent pipeline with the necessary skills. Through upskilling and reskilling initiatives, partnering with educational institutions enables companies to access expertise through formats such as practical training. The Institutes of Technology in the UK is a case in point of collaboration among colleagues, universities, and businesses, enabling companies to gain a skilled workforce that fosters long-term productivity and innovation growth.

In addition to partnerships with educational institutions, industry collaboration can also be a viable initiative, as it may be challenging for a single organisation to address industry-wide problems alone. By pooling resources through collaboration within or outside the industry, economies of scale can be achieved, reducing costs and facilitating the collective management of common problems.

### Building GenA.I.-related technical skills infrastructure

To suit the specific needs of building a GenA.I. talent pool, Hong Kong's policymakers may consider leveraging the solid general talent infrastructure already in place. For a start, the Hong Kong government has various existing talent attraction schemes, such as the Quality Migrant Admission Scheme (QMAS), which can be leveraged to attract the requisite AI talent for GenA.I. adoption and development in Hong Kong's financial services industry. Other than a more coordinated attraction effort through existing talent policies, policy frameworks covering continuing professional training can be adapted to include key GenA.I.-related skillsets. For instance, the HKMA's Enhanced Competency Framework

can be expanded to include key GenA.I.-related skillsets for banks. Given the ongoing demand for sandboxes, iterations of these programmes may also be tied to certain GenA.I. sandboxes, especially those that require niche technical or business skillsets to implement, hence enabling more expedient access to the appropriate talent.

Beyond the existing infrastructure, policymakers may consider setting up GenA.I.-specific human capital development funds and networks to facilitate speedier workforce formation, especially purpose-built teams encompassing both business and GenA.I. skillsets. In practice, these measures may include academia-industry-government collaboration via apprenticeships, mid-career (re-) training, and organisational consultation. Specific policy needs can also be calibrated through frequent and constructive engagement between the regulators, industry, and GenA.I. service providers. Retraining programmes could also address any potential job displacement coincident to the adoption of GenA.I. in the industry.

# Conclusions

Financial institutions in Hong Kong recognise the strategic potential that GenA.I. can bring to the financial services industry, and have begun to explore practical GenA.I. applications. Adoption is largely through collaboration with third-party service providers, with use cases primarily in lower-risk internal and non-customer facing applications aimed at enhancing productivity and operational efficiency.

To ensure ongoing effective risk management and compliance, financial institutions adopting GenA.I. have begun to enhance their AI data governance policies and standards, introduce measures to mitigate new cybersecurity threats, and incorporate additional controls throughout the model life-cycle. In response to third-party concerns, GenA.I. service providers are also increasing the transparency of their products. However, with the technology still evolving rapidly, a proactive and iterative approach to strengthening financial institutions' risk management and governance frameworks will be key.

Local financial institutions highlighted three talent gaps: technical skills for developing GenA.I. models; combined technical and business skills for facilitating effective adoption within the industry; and compliance and risk management expertise essential for managing GenA.I.-specific risks. While financial institutions have deployed a multi-faceted talent approach with a focus on upskilling staff and augmentation through offshore development centres, the survey and interview findings suggest that financial regulators and policymakers have a role to play in further strengthening the local talent pipeline.

Both financial institutions and GenA.I. service providers appreciated the efforts of the Hong Kong government and regulators to promote responsible GenA.I. adoption and collaboration, such as the release of updated circulars and guidelines on the use of GenA.I. models, the introduction of the GenA.I. Sandbox, and the promotion of industry-wide events that support knowledge exchange. The newly launched Cyberport AI Supercomputing Centre is also anticipated to support innovation. However, at the disaggregated level, the speed of GenA.I. adoption continued to exhibit some variation. GenA.I. adoption has been somewhat higher among larger surveyed financial institutions, as smaller institutions faced more acute resource and talent constraints.

As such, advancing broad-based GenA.I. adoption in Hong Kong's financial services industry will likely require further actions to promote a level playing field and to strengthen the local technology and talent ecosystems. These include broader financial literacy programmes and sandboxes and expanded industry-wide technology infrastructure. Further collaboration among stakeholders, including through targeted funding or cross-industry working groups, may also help to reduce adoption hurdles and support better understanding and solutions of issues such as accountability, data localisation, data sharing and data protection.

Looking ahead, as GenA.I. adoption broadens in Hong Kong's increasingly digitalised financial services industry, further study may be warranted to assess the role of GenA.I. in enhancing the efficiency and accuracy of regulatory compliance through innovative RegTech solutions, in strengthening the potential of SupTech as a transformative force in financial supervision, and in contributing to policy setting and the operational efficiency of financial regulators.

# Appendix A:


## Background of the GenA.I. Survey

The *Financial Services in the Era of GenA.I.: Opportunities and Risk Management* survey (GenA.I. Survey) was commissioned by the HKIMR, and executed by Ernst & Young Advisory Services Limited over October 2024 to January 2025. The GenA.I. survey aimed to understand the current GenA.I. adoption landscape, as well as the risks and considerations behind its adoption in Hong Kong's financial services industry. It also aimed to reveal the various governance and mitigation measures considered by financial institutions in response to GenA.I.-related risks, as well as talent gaps faced by the industry.

55 entities from the banking, insurance, and WAM sectors responded to the survey (Figure A.1)<sup>115</sup>. These institutions were headquartered in various jurisdictions, including Hong Kong, mainland China, other Asian jurisdictions, and the United States. The survey also encompassed a variety of small, medium, and large institutions across all three sectors (Figure A.2).

Our survey respondents constituted 57% of total deposits in the banking sector<sup>116</sup>, and 48% of gross premia in the insurance sector<sup>117</sup>. The WAM sector respondents were SFC-registered legal entities, covering mainland Chinese, Hong Kong and international institutions of varying sizes.

**Figure A.1: Number of survey respondents by sector**




Source: HKIMR staff compilation based on the GenA.I. Survey.

<sup>115</sup> Calculated based on coverage of master groups by sector, such that responses from various subsidiaries within the same sector are combined and counted as one response.

<sup>116</sup> Based on financial statements of banks as of December 2023.

<sup>117</sup> Based on the IA's reported statistics on gross premia as of December 2023.

Figure A.2: Distribution of survey respondents by sector and size of institution



Source: HKIMR staff compilation based on the GenA.I. Survey.

The findings of the GenA.I. Survey were supplemented by 30 in-depth individual interviews, covering representatives from 6 GenA.I. service providers from various jurisdictions, 9 banks, 7 insurers, and 8 WAMs firms. The interviews with GenA.I. service providers focused on collaboration with the financial services industry, responsible GenA.I. adoption, talent development and cross-industry use cases. The interviews with financial institutions leaned towards larger firms that are actively engaged in GenA.I. initiatives to draw out insights complementary to the survey.



# Appendix B:

## Glossary of Technical Terms<sup>118</sup>

Term	Meaning
Adversarial attacks	Attempts to fool an AI model by introducing subtle, malicious changes to input data, such as adversarial images designed to mislead classification systems.
Air-gapped system architecture	A security measure where critical systems are physically isolated from unsecured networks to prevent unauthorised access or data breaches.
AML (Anti-Money Laundering)	A global network of laws, rules, and processes designed to detect and expose funds that have been illegally obtained but presented as lawful income.
Artificial Neural Networks (ANNs)	Computational models inspired by how the brain works, made up of layers of connected nodes (neurones) that help recognise patterns and learn from data.
Backpropagation	A training method for neural networks that reduces errors by calculating how much each connection (weight) contributes to the error and then adjusting those weights step by step, working backward through the network.
BCBS (Basel Committee on Banking Supervision)	The primary global standard setter for the prudential regulation of banks and provides a forum for regular co-operation on banking supervisory matters.
Big Data	Large, complex datasets that require advanced computational techniques to process and analyse, often used as input for AI models.
Black Box Effect	The phenomenon where the internal workings of an AI model are not interpretable or explainable, making it difficult to understand how decisions are made.
Bug bounty	A programme where organisations reward individuals (typically ethical hackers) for identifying and reporting security vulnerabilities in their systems.
Cloud Services	Platforms offering scalable, on-demand computational and storage resources via the internet, enabling AI training and deployment without reliance on local hardware.
Concentration risk	The risk of over-reliance on a single vendor, technology, or service provider, which could lead to operational or security issues in the event of failure or disruption.
DALL-E	An AI model developed by OpenAI that generates images from textual descriptions.

<sup>118</sup> The glossary is sourced from various publicly available information and should not be considered as official definitions.

<b>Data masking</b>	A technique used to hide information by changing its original characters and digits. It is commonly done to meet regulatory and privacy standards.
<b>Data poisoning</b>	A type of adversarial attack where malicious data is intentionally injected into a training dataset to corrupt or manipulate the performance of an AI model.
<b>Data residency</b>	The legal or regulatory requirement to store and process data in a specific geographic location, often to comply with local data protection laws.
<b>Decentralised infrastructure</b>	A system where data, processing, and decision-making are distributed across multiple nodes or devices, reducing reliance on a central authority and enhancing resilience.
<b>Deep Machine Learning</b>	A kind of machine learning that relies on artificial neural networks to understand and learn from data.
<b>Deepfake</b>	AI-generated content, such as images, videos, or audio, that is manipulated to appear real, often used maliciously to spread misinformation or create deceptive media.
<b>Diffusion Models</b>	A type of generative model that produces images.
<b>Endpoint security</b>	Protection measures applied to devices (endpoints) connected to a network, such as computers, smartphones, to safeguard against cyber threats.
<b>Foundation Models</b>	Large-scale pre-trained AI models, such as GPT or BERT, that serve as a base for fine-tuning across a wide range of downstream tasks.
<b>GDPR (General Data Protection Regulation)</b>	Regulates the processing and transfer of personal data belonging to individuals within the EU.
<b>Generative Adversarial Networks (GANs)</b>	A type of neural network consisting of two competing networks (generator and discriminator) to create realistic synthetic data.
<b>Generative Artificial Intelligence (GenA.I.)</b>	A subset of AI focused on generating new content, such as text, images, music, or videos.
<b>GPT (Generative Pre-trained Transformer)</b>	A family of large language models developed by OpenAI, designed to generate human-like text.
<b>GPU</b>	Specialised hardware accelerators that utilise numerous smaller, highly efficient processing cores to process data in parallel, making them particularly effective for tasks such as graphics rendering and complex computations.
<b>Greedy Layer-Wise Pretraining</b>	A technique for deep neural networks where each layer is trained one at a time.

<b>Hallucination</b>	In AI, refers to a phenomenon where a large language model (LLM) detects patterns or objects that do not exist or cannot be perceived by humans, leading to outputs that are illogical or entirely incorrect.
<b>HITL (Human-in-the-Loop)</b>	An approach in AI system development and adoption where humans actively participate in the training, operation, or validation of AI systems to ensure better performance, safety, and ethical alignment.
<b>Internet of Things</b>	A system that connects devices and equipment through the internet, allowing them to communicate, exchange data, and perform tasks automatically.
<b>KYC (Know-Your-Customer)</b>	A standard in the investment industry that requires advisors to verify a client's identity and understand their investment knowledge and financial background.
<b>Large Language Model</b>	A type of generative models that can produce contextually relevant language and code outputs, leveraging vast amounts of training data to understand and generate text.
<b>Long Short-Term Memory (LSTM) networks</b>	LSTM is highly effective for sequence prediction tasks as it captures long-term relationships in data.
<b>Model life cycle</b>	The model life cycle refers to the process a model goes through from creation to retirement. It involves stages such as development, testing, monitoring, regulation, and reporting.
<b>Model validation</b>	A process of ensuring that a model performs as intended and meets its specified objectives.
<b>Multimodal Models</b>	Models designed to process and integrate data from multiple modalities.
<b>Parallel Processing</b>	A computational strategy where multiple processes or tasks are executed simultaneously.
<b>Prompt</b>	A query or request for specific output, serving as an instruction to an application interface powered by an LLM.
<b>Recurrent Neural Networks (RNNs)</b>	A type of deep neural network designed for sequential or time-series data, enabling a machine learning (ML) model to make predictions or draw conclusions based on ordered inputs.
<b>Sandbox</b>	A secure, isolated testing environment that allows users to execute programmes or open files without impacting the underlying application, system, or platform.
<b>SEO/Search Engine Optimisation</b>	A way to improve a website so that it appears higher in search engine results. This helps more people find and visit the site.
<b>StyleGAN</b>	An extended version of the GAN architecture, introduced by NVIDIA.

<b>Transfer Learning</b>	A machine learning technique where knowledge gained from one task or dataset is applied to enhance model performance on a related task or a different dataset.
<b>Transformers</b>	A type of neural network created by Google that processes data using self-attention mechanisms instead of traditional methods. This allows it to handle sequences (like sentences) efficiently and in parallel.
<b>Two-factor authentication</b>	A security method that requires two forms of verification.
<b>Variational Autoencoders (VAEs)</b>	Machine learning models designed to create new data by generating variations of the data they were trained on.
<b>White hat hackers</b>	Ethical hackers who use their skills to identify and fix security vulnerabilities in systems, working to protect organisations and users from malicious attacks.

# Appendix C:

## References

**Al-Kfairy et al. (2024).** Ethical Challenges and Solutions of Generative AI: An Interdisciplinary Perspective. *Informatics*, 11(3), 58. August 9, 2024.

**Alzubaidi, L. et al. (2023).** A Survey on Deep Learning Tools Dealing with Data Scarcity: Definitions, Challenges, Solutions, Tips, and Applications. *J Big Data*, 10, 46, April 14, 2023.

**Anwah, O. & Rosenbaum, E. (2023).** AI Is Now the Biggest Spend for Nearly 50% of Top Tech Executives Across the Economy: CNBC Survey. *CNBC*. June 23, 2023.

**Asia Securities Industry and Financial Markets Association (2024).** Enabling an Efficient Regulatory Environment for AI – Practical Considerations for Generative AI. January 18, 2024.

**Bank of International Settlements (2024).** Regulating AI in the Financial Sector: Recent Developments and Main Challenges. FSI Insights on Policy Implementation, No. 63. December 12, 2024.

**Bank of International Settlements (2024a).** Artificial Intelligence and the Economy: Implications for Central Banks. BIS Annual Economic Report. June 25, 2024.

**Banque de France (2024).** The financial supervisor in the age of AI. Speech. France. June 5, 2024.

**BCG (2024).** AI at Work: Friend and Foe. June 2024.

**Bommasani, R. et al. (2022).** On the Opportunities and Risks of Foundation Models. *arXiv preprint arXiv:2108.07258*. August 16, 2021.

**Brollio, F., Dabla-Norris, E., de Mooij, R., et al (2024).** Broadening the Gains from Generative AI: The Role of Fiscal Policies. *Staff Discussion Notes 2024*, 002 (2024), International Monetary Fund. June 17, 2024.

**Calvano, E. et al. (2020).** Artificial Intelligence, Algorithmic Pricing and Collusion. *American Economic Review*, Vol. 110, No. 10. October 2020.

**Cao, L. (2021).** T-Shaped Teams: Organising to Adopt AI and Big Data at Investment Firms. *CFA Institute Research Reports*. August 30, 2021.

**Deloitte AI Institute (2020).** ‘Trustworthy AI’ is a framework to help manage unique risks. *MIT Technology Review*. March 25, 2020.

**Epoch AI (2024, Updated 2025).** Notable AI Models. June 19, 2024, updated March 28, 2025.

**ERP Today (2024).** IBM-Oracle: Skills Gap Biggest Hurdle for AI Adoption. October 16, 2024.

**European Commission (2020).** AI Watch, Historical Evolution of Artificial Intelligence – Analysis of the Three Main Paradigm Shifts in AI. Luxembourg, Europe. November 20, 2020.

**Ernst & Young/EY (2020).** Understand Model Risk Management for AI and Machine Learning. May 13, 2020.

**Ernst & Young/EY (2024).** How GenA.I. Will Impact the Labor Market. April 15 2024.

**Financial Services and Treasury Bureau (2024).** Policy Statement on Responsible Application of Artificial Intelligence in the Financial Market. January 28, 2024.

**Financial Stability Board (2017).** Artificial Intelligence and Machine Learning in Financial Services. November 1, 2017.

**Financial Stability Board (2024).** The AI Adventure: How Artificial Intelligence May Shape the Economy and the Financial System. July 11, 2024.

**Financial Stability Board (2024a).** The Financial Stability Implications of Artificial Intelligence. November 14, 2024.

**Financial Stability Board (2024b).** The AI Adventure – How Artificial Intelligence May Shape the Economy and the Financial System. July 11, 2024.

**Fuchs, A., Passarella, A., & Conti, M. (2024).** Optimizing Delegation in Collaborative Human-AI Hybrid Teams. *ACM Transactions on Autonomous and Adaptive Systems, Volume 19, Issue 4.* November 16, 2024.

**Fuster, A. et al. (2019).** The Role of Technology in Mortgage Lending. *The Review of Financial Studies*, Volume 32, Issue 5. May, 2019.

**Georges, C. & Pereira, J. (2021).** Market Stability with Machine Learning Agents. *Journal of Economic Dynamics and Control*, Volume 122. January 2021.

**Hagendorff, T. (2024).** Mapping the Ethics of Generative AI: A Comprehensive Scoping Review. *Minds & Machines*, 34, 39. September 17, 2024.

**Harvard Business Review (2015).** Emotion and the Art of Negotiation. December 2015.

**Harvard Business Review (2022).** How Great Leaders Communicate. November 23, 2022.

**Harvard Business Review (2023).** Generative AI Has an Intellectual Property Problem. April 7, 2023.

**Harvard Business Review (2023a).** What Will Working with AI Really Require? June 8, 2023.

**Harvard Business Review (2023b).** The Skills Your Employees Need to Work Effectively with AI. November 3, 2023.

**Harvard Business Review (2023c).** Reskilling in the Age of AI. October 9, 2023.

**Harvard Business Review (2024).** Is Your Company's Data Ready for Generative AI? March 26, 2024.

**Harvard Business Review (2024b).** Research: How Gen AI Is Already Impacting the Labor Market. November 11, 2024.

**HKSAR Government (2024).** Budget Speech – International Innovation and Technology Centre. February 28, 2024.

**Heikkilä, M. (2023).** AI Language Models Are Rife with Different Political Biases. MIT Technology Review. August 7, 2023.

**Hong Kong Institute for Monetary and Financial Research (2020).** Artificial Intelligence in Banking – The Changing Landscape in Compliance and Supervision. *HKIMR Applied Research Report No.2/2020*. Hong Kong, China. August 2020.

**Hong Kong Institute for Monetary and Financial Research (2021).** Artificial Intelligence and Big Data in the Financial Services Industry – A Regional Perspective and Strategies for Talent Development. *HKIMR Applied Research Report No.2/2021*. Hong Kong, China. October 2021.

**Hong Kong Institute for Monetary and Financial Research (2023).** Advancing Talent Development in Financial Services – Emerging Global Trends and Their Impact on Hong Kong. *HKIMR Applied Research Report No.3/2023*. Hong Kong, China. December 2023.

**Hong Kong Monetary Authority (2016).** HKMA-ASTRI Fintech Innovation Hub. September 6, 2016.

**Hong Kong Monetary Authority (2019).** High-Level Principles on Artificial Intelligence. November 1, 2019.

**Hong Kong Monetary Authority (2019a).** Consumer Protection in Respect of Use of Big Data Analytics and Artificial Intelligence by Authorized Institutions. November 5, 2019.

**Hong Kong Monetary Authority (2024).** Consumer Protection in Respect of Use of Generative Artificial Intelligence. August 19, 2024.

**Hong Kong Monetary Authority (2024a).** Generative Artificial Intelligence in the Financial Services Space. September 2024.

**Hong Kong Monetary Authority (2024b).** HKMA Announces Inaugural Cohort of GenA.I. Sandbox. December 19, 2024.

**Hu, Q. et al. (2024).** Characterization of Large Language Model Development in the Datacenter. *arXiv:2403.07648*. December 3, 2024.

**International Data Corporation (2024).** Worldwide AI and Generative AI Spending Guide.

**International Monetary Fund (2023).** IMF Fintech Note on Generative Artificial Intelligence in Finance: Risk Considerations. August 2023.

**International Monetary Fund (2024).** Broadening the Gains from Generative AI: The Role of Fiscal Policies. June 17, 2024.



**Insurance Authority (2024).** Insurance Authority Shares Regulatory Insights at the Hong Kong FinTech Week. October 29, 2024.

**Kaddour, J. et al. (2023).** Challenges and Applications of Large Language Models. *arXiv:2307.10169*. July 19, 2023.

**Kirova et al. (2023).** The Ethics of Artificial Intelligence in the Era of Generative AI. *Journal of Systemics, Cybernetics and Informatics*, 21(4):42-50. December 2023.

**KPMG (2022).** Modern Risk Management for AI Models. July 5, 2022.

**Mayer Brown (2024).** Hong Kong PCPD Issues Model Personal Data Protection AI Framework. June 26, 2024.

**McKinsey & Company (2023).** The Economic Potential of Generative AI: The Next Productivity Frontier. June 14, 2023.

**MIT Management Sloan School (2021).** Tapping the power of unstructured data. February 1, 2011

**MIT Management Sloan School (2024).** The impact of generative AI as a general-purpose technology. August 6, 2024

**MIT Technology Review (2017).** Tencent's New Lab Shows It's Serious about Mastering AI. May 2, 2017.

**Monetary Authority of Singapore (2023).** Veritas Initiative. October 26, 2023.

**Morgan Stanley (2023).** Compounding Through the Hype. July 31, 2023.

**Information Services Department (2024).** AI Subsidy Scheme Opens. October 7, 2024.

**OECD (2021).** Artificial Intelligence, Machine Learning and Big Data in Finance: Opportunities, Challenges and Implications for Policy Makers. August 11, 2021.

**OECD (2023).** Generative Artificial Intelligence in Finance. December 15, 2023.

**OECD (2023a).** OECD Employment Outlook 2023: Artificial Intelligence and the Labour Market. July 11, 2023.

**OECD (2023b).** The Impact of AI on the Workplace: Main Findings from the OECD AI Surveys of Employers and Workers. March 27, 2023.

**OECD-FSB Roundtable on Artificial Intelligence (AI) in France (2024).** Proceedings of the OECD-FSB Roundtable on AI. September 30, 2024.

**Office of the Privacy Commissioner for Personal Data (2021).** Guidance on Ethical Development and Use of AI. Hong Kong, China. August 2021.

**Office of the Privacy Commissioner for Personal Data (2024).** Artificial Intelligence: Model Personal Data Protection Framework. Hong Kong, China. June 11, 2024.

**Office of the Privacy Commissioner for Personal Data (2025).** Checklist on Guidelines for the Use of Generative AI by Employees. Hong Kong, China. March 31, 2025.

**Securiti (2024).** The Evolution of Data Quality: How GenAI Is Setting New Standards. July 9, 2024.

**Securities and Futures Commission (2024).** Circular to Licensed Corporations – Use of Generative AI Language Models. November 12, 2024.

**Shine, I. (2023).** We Often Hear That AI Will Take Our Jobs. But What Jobs Will It Create? World Economic Forum. September 18, 2023.

**Smart Nation Singapore (2023).** National Artificial Intelligence Strategy (NAIS) 2.0. December 4, 2023.

**South China Morning Post (2017).** China's First 'Deep Learning Lab' Intensifies Challenge to US in Artificial Intelligence Race. February 21, 2017.

**Stanford University (2024).** The AI Index 2024 Annual Report. April 2024.

**The CMS Data Protection Group (2024).** GDPR Enforcement Tracker Report. May 15, 2024.

**The Standard (2024).** Cyberport Wins HKD3 Billion AI Handout. February 29, 2024.

**Vaswani, A. et al. (2023).** Attention Is All You Need. *arXiv:1706.03762*. February 8, 2023.

**Weedmark, D. (2022).** The Importance of Machine Learning Model Validation and How It Works. *Domino*. November 15, 2022.

**White House (2024).** READOUT: U.S. Department of the Treasury Hosts Roundtable with Consumer and Climate Groups on Climate-Related Risks and U.S. Insurance Markets. March 19, 2024.

**World Economic Forum (2023).** We Often Hear That AI Will Take Our Jobs. But What Jobs Will It Create? September 18, 2023.

**World Economic Forum (2023a).** Jobs of Tomorrow: Large Language Models and Jobs. September 18, 2023.

**World Economic Forum (2023b).** 3 Essential Features of Global Generative AI Governance. November 8, 2023.

**World Economic Forum (2024).** AI Is Shifting the Workplace Skillset. But Human Skills Still Count. January 21, 2025.

**World Economic Forum (2024a).** New Study Finds AI Is Making Soft Skills More Important in the Workplace. January 31, 2024.

**World Economic Forum (2024b).** Leveraging Generative AI for Job Augmentation and Workforce Productivity: Scenarios, Case Studies and a Framework for Action. November, 2024.

**Wu, X., Duan, R., & Ni, J. (2024).** Unveiling Security, Privacy, and Ethical Concerns of ChatGPT. *Journal of Information and Intelligence*. Volume 2, Issue 2. March 2024.

**Yang, C. H. (2022).** How Artificial Intelligence Technology Affects Productivity and Employment: Firm-Level Evidence from Taiwan. *Research Policy*, Volume 51, Issue 6. July 2022.

**Zhu, B., Mu, N., Jiao, J., & Wagner, D. (2024).** Generative AI Security: Challenges and Countermeasures. *arXiv:2402.12617*. February 2024.

## ABOUT THE HONG KONG ACADEMY OF FINANCE (AOF)

The AoF is set up with full collaboration amongst the Hong Kong Monetary Authority, the Securities and Futures Commission, the Insurance Authority and the Mandatory Provident Fund Schemes Authority. By bringing together the strengths of the industry, the regulatory community, professional bodies and the academia, it aims to serve as (i) a centre of excellence for developing financial leadership; and (ii) a repository of knowledge in monetary and financial research, including applied research.

## ABOUT THE HONG KONG INSTITUTE FOR MONETARY AND FINANCIAL RESEARCH (HKIMR)

The HKIMR is the research arm of the AoF. Its main remit is to conduct research in the fields of monetary policy, banking and finance that are of strategic importance to Hong Kong and the Asia region. The Applied Research studies undertaken by the HKIMR are on topics that are highly relevant to the financial industry and regulators in Hong Kong, and they aim to provide insights on the long-term development strategy and direction of Hong Kong's financial industry.

### CONTACT US

Email: [hkimr@hkma.gov.hk](mailto:hkimr@hkma.gov.hk)

Tel: +852 2597 0765

Website: <https://www.aof.org.hk/research/HKIMR>



This is printed on environmentally friendly paper

