



In partnership with
 OliverWyman

THE IMPACT OF AI IN FINANCIAL SERVICES

Opportunities, Risks and Policy Considerations



REPORT QUALIFICATIONS/ASSUMPTIONS & LIMITING CONDITIONS

Oliver Wyman was commissioned by UK Finance to collaboratively prepare a report on AI in the financial sector. The primary audience for this report includes specialists and senior management from firms, policy makers and regulators.

Oliver Wyman and UK Finance shall not have any liability to any third party in respect of this report or any actions taken or decisions made as a consequence of the results, advice or recommendations set forth herein.

This report does not represent formal or legal advice, which can only be provided by legal counsel and for which you should seek advice of counsel, and whilst anyone is welcome to use this report, it is entirely at their own risk. The opinions expressed herein represent the views of Oliver Wyman and UK Finance, it is strictly for information and has not been approved by a regulatory body. Whilst public information and industry and statistical data are from sources Oliver Wyman and UK Finance deem to be reliable, Oliver Wyman and UK Finance make no representation or warranty as to the accuracy or completeness of such information and has accepted the information without further verification. Accordingly, Oliver Wyman and UK Finance shall not be responsible or liable for any loss, damages or costs arising from the use of this report. Users of this report should ensure that it is suitable for their use (and that appropriate due diligence has been conducted, including in relation to compliance with relevant laws). Therefore, Oliver Wyman and UK Finance take no responsibility for changes in market conditions or laws or regulations and no obligation is assumed to revise this report to reflect changes, events or conditions, which occur subsequent to the date hereof.

FOREWORD



Jana Mackintosh

Managing Director
Payments, Innovation and Resilience
UK Finance

Developments in artificial intelligence (AI) have provoked a mixture of excitement and anxiety among commentators, politicians, policy makers and members of the public. In this report we aim to present a rounded picture.

We illuminate some of the ways in which AI in the financial services sector can bring value to businesses and to consumers. This includes Generative AI — which has captured the public imagination this year — and also ‘traditional’ Predictive AI, which has a more developed position in financial services.

We hope to assist industry players, policy makers and other readers in understanding the state of play in the sector in terms of current uptake and applications, as well as where the technology might be utilized in the future. The report also explores what firms are up to in terms of analysing, understanding and managing AI risks. Firms are innovating and exploring AI use cases but are doing so carefully, conscious of the risks.

The nature of the technology gives rise to novel regulatory and policy challenges, for which best practice is yet to emerge. We intend for this report to contribute to this debate. Although there is broad international consensus over many AI risks, different approaches to tackling them are emerging globally. Ultimately, these will be tested over time, acting as a natural experiment and revealing what measures best mitigate risks, and which are most conducive to responsible innovation and uptake.

UK Finance looks forward to participating actively in this important area of policy development, leveraging the expertise of the sector through its AI Policy Committee and engaging with the wide range of interested actors within and outside of government.

FOREWORD

**Lisa Quest**

Partner, Head of UK and Ireland

Co-Head of the Public Sector and Policy Practice Europe

Oliver Wyman

AI, particularly Generative AI, is transforming industries with its recent breakthroughs, offering exciting possibilities alongside challenges that demand attention. Oliver Wyman is working with a range of our clients on this topic. This resurgence in AI discussions is now a top priority for executives — and the impacts on the industry will be profound. These include innovative propositions, enhanced user experiences, and increased automation to reduce costs and bridge skills gaps.

We are delighted to work with UK Finance to examine the current and future state of AI in the UK financial services sector in this report. Working at the forefront of this topic, in this collaboratively way with the industry would only be possible with the relationship Oliver Wyman has with UK Finance. Drawing on insights from in-depth interviews and a survey of UK Finance members, we see that more than 90 per cent of which have deployed AI, we highlight the immense opportunities and the steps that the industry has already taken. We also identify key discussion areas for safe AI adoption.

The UK is a leader in financial services and renowned for its adoption of technology, while managing to balance safety and innovation in its regulatory ecosystem. The recent AI Safety Summit underlines this leadership. Given the importance of the AI topic, we have worked hard to support UK Finance to set the immediate agenda for financial institutions and regulators to further refine AI regulations during this critical period in technology regulation.

As advisors to the industry, we understand the effort required to adopt new technologies and create value for all stakeholders. With 60 per cent of respondents anticipating significant cost savings from AI, the path to impactful return on investment is long but promising.

Join us on this journey as we explore the transformative potential of AI in financial services. Together, we can navigate the path towards responsible and impactful AI adoption.

TABLE OF CONTENTS

1. Introduction	1
1.1. Summary of this report	1

2. High-level overview of the current landscape	3
2.1. Foundational concepts for this report	3
2.2. How to think about Generative AI differently to other existing AI methods	4
2.3. What is Generative AI good at and where are its limitations?	5

3. AI in financial services	7
3.1. AI models in the financial services industry — overview	7
3.2. Implementation challenges of Predictive and Generative AI	9

4. Unlocking the benefits of Generative AI	11
4.1. Laying the foundation for a sustainable competitive advantage	11
4.2. Roadmap to unlocking Generative AI benefits	13

5. AI risks and mitigation	15
5.1. Known risks	16
5.2. Emerging risks from Generative AI	17
5.3. Mitigations	18
5.4. Spotlight on the risk of AI misuse by bad actors	19

6. Strategic use of AI — case studies	20
6.1. Case study 1 — Marsh McLennan's LenAI	20
6.2. Case Study 2 — Google's Anti-Money Laundering (AML) AI	23

7. Policy and regulatory landscape	24
7.1. Current state of AI regulation	24
7.2. Financial services sector views on AI regulation	27
7.3. Policy considerations and topics for further discussion	27

8. Conclusion and future outlook	32
---	-----------

Endnotes and additional references	33
---	-----------

1. INTRODUCTION

1.1. SUMMARY OF THIS REPORT

We are in the very early phases of a major technological change. To take stock, UK Finance and its members, in collaboration with Oliver Wyman, have undertaken a study of the state of AI adoption, its emerging applications and the risks it poses to financial services.

UK financial institutions see a substantial opportunity in artificial intelligence, with 90 per cent of respondents in our survey already leveraging Predictive AI in back-office functions, yielding tangible benefits. Although Generative AI is relatively new, more than 60 per cent believed it has the potential to deliver significant cost savings and improvements to operational effectiveness. There is an appetite within institutions to harness the potential of this transformative technology, which will necessitate a re-evaluation of business processes, employee skills, and staffing considerations. In addition, organisations will need to address the potential impact of compute-intensive AI systems, which consume significant resources and take up a large amount of space on sustainability targets for supply chains.

As a highly regulated sector, financial institutions are proceeding carefully with their adoption of AI. For now, more than 70 per cent of Generative AI use cases are in the proof of concept or pilot phase. The initial wave of adoption will provide valuable insights, but it is acknowledged that getting a return on investment will be reliant on data quality and seamless integration into existing systems, a process which could take three to five years. Truly transformative applications are still unknown but are likely to stem from Predictive and Generative AI being used together.

The learning curve is steep, however, and numerous unanswered questions remain. While best practice in AI risk was emerging globally, the advent of Generative AI has surfaced additional risks, such as 'hallucinations', and accentuated the challenge of needing to procure models from external providers. Most institutions believe they are well equipped to identify, monitor and mitigate the risks, with 60 per cent already leveraging existing risk management capabilities and adjusting their frameworks to include Generative AI.

There is support for the UK's flexible approach to AI regulation, based on principles and outcomes, as compared to prescriptive rules on the application of the technology. However, according to our survey, 65 per cent of respondents consider uncertainty regarding the direction of regulation as a top concern for the adoption of AI in the UK. There are open policy questions about ensuring AI guidance has clear scope, balancing the information needs of firms procuring AI tools against the IP concerns of third-party providers, and the harmonisation of cross-sectoral and cross-jurisdictional regulation. This is particularly relevant in the context of emerging international approaches that may have extra-territorial implications. UK Finance is positioned to be at the forefront of these discussions through its AI Policy Committee. Overall, industry members are eager to actively participate in the policy process and support the development of best practices over time.

The industry should aim for a rapid adoption of AI tools to deliver efficiency, a better customer experience, and a more robust sector. This will require all involved, from senior management to technology and product teams in financial institutions, and their counterparts in regulation and technology to get up to speed quickly on existing and emerging risks to be managed.

Key findings from our survey (23 financial institutions)

State of play

91%

of financial institutions have either narrowly or widely deployed Predictive AI in fraud detection and back-office functions with recorded benefits

>70%

of financial institutions are in the proof of concept or pilot stage for Generative AI use cases

Potential benefits

Only 13%

believe revenue opportunities will be in the top three expected benefits. Benefits are expected to come from productivity improvement and operational effectiveness

>75%

expect the same or higher benefit from Generative AI compared to Predictive AI

Process automation, sales and customer service functions

are areas where Generative AI use is expected to be more prevalent than Predictive AI is today

Risks

>65%

of UK high street banks have taken action to upgrade AI risk management policies to account for Generative AI

>95%

are accounting for AI-related risks within risk frameworks

>70%

treat Generative AI-related risks differently to Predictive AI-related risks

Regulation

65%

consider conflicting rules between different jurisdictions to be among the top three concerns relating to regulation

>80%

believe that a collaboration with UK regulators would be beneficial

Source: UK Finance members survey

1.2. PURPOSE OF THE REPORT AND INTENDED AUDIENCE

This report has a number of audiences, and a range of baseline levels of understanding. To ensure accessibility and clarity, we begin with an introductory overview, outlining the essential concepts necessary to understand the contents of this report.

Table 1: Reading guide

	1. Introduction	2. High level AI overview	3. AI in financial services	4. Unlocking the benefits of Generative AI	5. Risks and mitigations	6. AI use cases	7. Policy and regulation
Executives	Skip if you feel you have a good understanding of AI — particularly Generative AI — already						
Business/function managers							
Risk managers							
Policy makers							
Supervisors							
Technologists							

1.3. METHODOLOGY AND APPROACH

This report is based on a proprietary survey conducted among 23 member organisations of UK Finance, representing various institutions in the UK financial services sector. The survey covered topics such as the adoption and deployment of Predictive and Generative AI, use cases, anticipated benefits and risks, risk management, and views on regulation. Follow-up interviews were conducted with nearly half of the surveyed members, providing additional insights. Regular

discussion forums involving more than 30 members were also held to discuss AI policy and regulation. It is important to note that the views expressed by members, while representative, cannot be attributed conclusively to the entire sector. This report, developed in collaboration with Oliver Wyman, leverages their expertise on AI and intellectual capital from past projects and industry experts.

2. HIGH-LEVEL OVERVIEW OF THE CURRENT LANDSCAPE

Key messages:

- It is important to differentiate between Generative AI and Predictive AI, as there are key differences in the applications for which these technologies are suited. Predictive AI models are more suited to tasks requiring reasoning, pattern recognition, and analysis, while Generative AI is more suited to applications requiring fluency, with its strengths lying in content generation.
- The inherent uncertainty or creativity in the outputs of Generative AI models constitute a deliberate design feature, rather than a flaw. Organisations must pay careful attention to the appropriate applications for Generative AI and select the right model for each context.

2.1. FOUNDATIONAL CONCEPTS FOR THIS REPORT

Artificial intelligence (AI) is being adopted by companies and end users across diverse industries around the world. Much has been written about the history of the technology, its potential and the different ways it could be used.

As a technology category, AI covers many capabilities, from advanced analytics, automation, and predictive intelligence through to more recent generative intelligence.

For this report, we will define AI broadly as the spectrum of tools that includes Predictive AI and Generative AI.

It is important to differentiate between Generative AI and Predictive AI, as there are key differences in how these technologies are used from types of models to user input. **Figure 1** displays two different applications, an analytical use case and a code generation use case.

Figure 1: Predictive AI vs Generative AI: Description and use case comparison

	Predictive AI	Generative AI
Details	<ul style="list-style-type: none"> Advanced analytical techniques relying on different algorithms and large organised datasets 	<ul style="list-style-type: none"> These include large language models and multi-modal models which have the power to generate outputs from — usually very — large bodies of data
Types of uses	<ul style="list-style-type: none"> Analysis of large datasets to forecast potential scenarios and find outliers 	<ul style="list-style-type: none"> Interpretation, classification, manipulation, and generation of language content Generation of content across different data types — a combination of audio, code, images, text, and videos
Use case	<ul style="list-style-type: none"> AI fraud detection 	<ul style="list-style-type: none"> AI supported code generation
Example task	<ul style="list-style-type: none"> Bank wants a faster and more efficient way of identifying fraudulent transactions 	<ul style="list-style-type: none"> Bank wants to write code that can be used to classify digitised banking statements
Model use	<ul style="list-style-type: none"> Machine learning fraud detection model 	<ul style="list-style-type: none"> GitHub Co-Pilot
User input	<ul style="list-style-type: none"> Specific model parameters <i>For example: Banking transaction data</i> (Domain specific and often proprietary datasets) 	<ul style="list-style-type: none"> Free-form prompts (text, image, speech) <i>For example: 'Write function to extract document name'</i>
Process	<ul style="list-style-type: none"> Rules or template-based machine learning approach <i>For example: Random decision forests, supervised learning algorithms</i> 	<ul style="list-style-type: none"> Deep learning, large language model Majority transformer-based <i>For example: Billions of parameters from publicly available datasets combined with coding languages available in public repositories</i>
Process	<ul style="list-style-type: none"> Classification of outliers and potential cases of fraud 	<ul style="list-style-type: none"> Generation of content <i>For example: Code suggestions to help a software engineer answer the initial task</i>

Source: Oliver Wyman analysis

2.2. HOW TO THINK ABOUT GENERATIVE AI DIFFERENTLY TO OTHER EXISTING AI METHODS

The emergence of Generative AI presents financial institutions with a new set of tools, which will create value in new ways.

2.2.1. Data differences

The first critically important distinction lies in the data used to train Predictive AI models and Generative AI models. Predictive analytics usually rely on an organisation’s proprietary and domain data, whereas Generative AI models are trained on a vast corpus of data taken from various public and purchased sources. While it is possible that a firm may be able to fine-tune Generative AI models on proprietary data, the underlying, foundation model is still trained on external data (as it would likely be from a third-party provider). The cost of creating and training a foundation model means it’s unlikely that in the near-term a financial institution will do this itself. This presents data privacy

challenges (further details to be found in Chapter 5 — risks). It also presents an opportunity for financial services to come together to build sector-specific foundation models.

2.2.2. Differences in strengths and weaknesses

Many AI models lack explainability, that is, the extent to which the workings of a model, and the reasons for its outputs, can be understood. Generative AI models are particularly opaque, making it challenging to identify the root cause of errors, predict potential mistakes or explain decisions based on their outputs. Generative AI is optimised to generate probable or realistic-sounding answers rather than providing a calculated ‘right’ answer. As a result, the accuracy of the generated answer is uncertain. Technology and protocols need further development to establish secure ranges of confidence and ensure safe use. On the other hand, Predictive AI models are optimised to be accurate and predictable in what they output, but lack the creativity of Generative AI models.

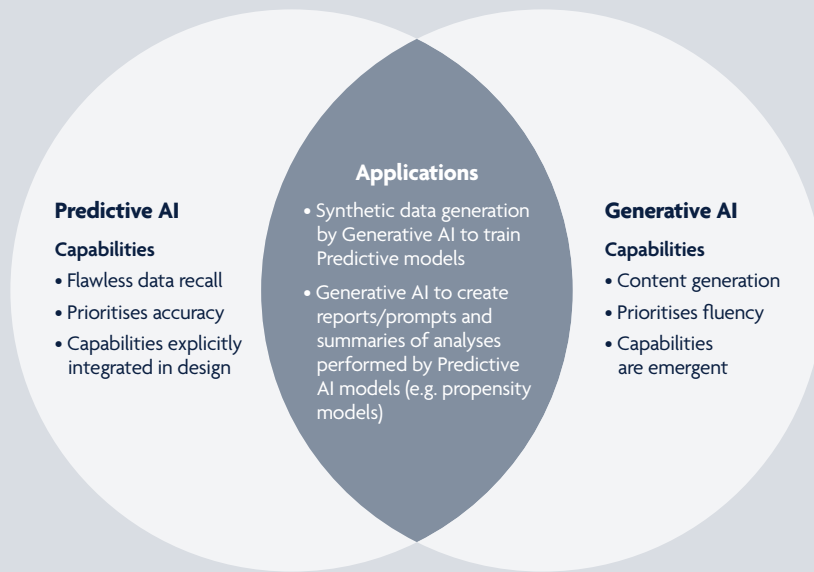
Predictive AI and Generative AI have distinct strengths and weaknesses that must be taken into account when developing use cases and deploying them across institutions. Predictive AI learns by understanding rules and using boundaries to classify data and is used to support data analysis functions such as classification and predictive models. Generative AI’s creativity is a real strength but creates risks when applied against the wrong use cases. Because of these distinctions, organisations need to consider carefully which systems require predictability and in which areas uncertainty and creativity should be valued and encouraged, or can at least be managed effectively.

2.2.3. How Generative AI complements Predictive AI

In financial services, Predictive AI usage is already advanced. The majority of its use cases are well defined and are expected to further develop and continue generating benefits.

There will be a point in which the maturity and strength of the different AI systems could lead to complementary uses when leveraged correctly. For instance, Predictive AI could be used in conjunction with Generative AI for anomaly detection purposes. To fully assess the impact of operating models and business case studies on regulatory compliance and governance, organisations should adopt a holistic approach to AI, thus identifying and enabling these synergies. This is shown in **Figure 2**.

Figure 2: How Predictive and Generative AI compare, differ and reinforce each other



Source: Oliver Wyman analysis

2.3. WHAT IS GENERATIVE AI GOOD AT AND WHERE ARE ITS LIMITATIONS?

Generative AI models have key capabilities that make them good at a specific set of tasks. They are not designed for reasoning, pattern recognition and analytical tasks in the same way Predictive AI models are. The inherent uncertainty or creativity in their outputs is a design feature, not a flaw.

Table 3 lays out some of the key limitations of Generative AI. These should be considered as part of choosing different technologies for different use cases. Recognising this is critical to the design of effective controls and the regulation thereof (see Chapters 5 and 7).

Table 3: Limitations of Generative AI tools

Agency	Knowledge	Reasoning	Predictability
Generative AI tools appear to have agency but are just designed to sound like that	Models have extensive implicit knowledge but are not aware of what they know, or of their own limitations	Sophisticated Generative AI models have learned to generate outputs that look like the product of analytical reasoning but which may not be	Models have tendencies to 'hallucinate' (producing outputs that are factually incorrect but framed with a high level of confidence)
Generative AI models lack goals, desires and the ability to learn in a self-directed manner	Models lack sense of truth or a grounded knowledge base	A model's ability to reason remains 'brittle' and likely to fail, especially when asked to apply new logic and knowledge outside of current training scope	Models can change output dramatically due to small or apparently meaningless changes in model inputs or prompts
	Models cannot recall data perfectly, just its statistical patterns	Opaque logic and processes — making interpretation difficult	Propensity to produce unwanted information can be reduced but 100% removal from a model is likely impossible, and could return if given certain prompts

Source: Oliver Wyman analysis

2.3.1. Summary of efforts in place to address Generative AI limitations

Alongside the development of the models themselves, tools and controls are being rapidly developed to mitigate these limitations or amplify the power of the tool in certain use cases. Broadly, these come in two main pillars: one focused on the technology itself, and one focused on humans. For further details, see Chapter 5.

Technological efforts: Efforts such as prompt engineering, Retrieval Augmented Generation (RAG) systems and guardrails can help financial institutions both get better results and avoid reputational damage.

Human efforts: Companies can manage and adapt their internal processes to ensure that employees can recognise inaccurate outputs and know what to do when they discover them. Examples include wide-scale education and training initiatives, clear governance processes and robust risk management.

3. AI IN FINANCIAL SERVICES

Key messages:

- AI adoption in financial services is increasing, with Predictive AI systems already deployed across various functions, and further growth expected in conjunction with the adoption of Generative AI.
- Generative AI is still in its early stages of deployment but is being explored carefully with technical maturity and customer outcomes in mind.
- The broader implementation of Generative AI in financial services faces challenges such as technical limitations, building a strong innovation foundation and recruiting the appropriate skill set. In addition to the challenge of ensuring alignment with existing legal frameworks, firms must manage the uncertainty about the direction of future AI regulation.
- The financial sector is in the early adoption phase of Generative AI, with an expected mass uptake in key functions, presenting an opportunity for firms to gain a competitive advantage.

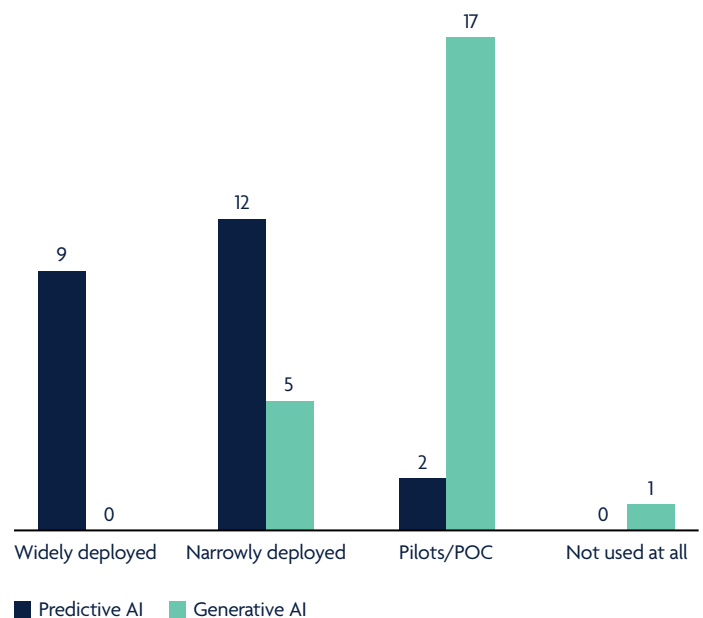
3.1. AI MODELS IN THE FINANCIAL SERVICES INDUSTRY — OVERVIEW

AI adoption in the financial services sector has grown, with Predictive AI being widely used, though Generative AI is still emerging. In this report we will focus primarily — for the sake of definition — on banking and payments, and this chapter will explore the adoption stages of both AI types in the industry. Despite initial adoption concerns, the potential of Generative AI is gaining attention, with institutions considering its integration alongside Predictive AI. As risk management improves, its adoption is expected to rise, potentially becoming a key competitive advantage in finance.

Figure 4 shows data from our survey demonstrating how Generative AI remains in the early stages of narrow or pilot deployment across the financial services sector, as organisations work out how best to use its strengths while managing its risks and limitations.

Figure 4: Survey results on current AI usage (23 financial institutions)

Which best describes your current AI usage at your institution?



Note: Widely deployed (five or more functions), Narrowly deployed (fewer than five functions), Pilots/POC (proof of concept) does not correspond to functions as it is in the conceptualisation/planning phase

Sources: UK Finance members survey, Oliver Wyman analysis

3.1.1. Predictive AI and Generative AI adoption within financial institutions


AI adoption in financial services is accelerating, driven by advances in predictive analytics and machine learning. Many institutions have already deployed Predictive AI systems across a wide range of functions. Despite this broad adoption, most firms surveyed claim that Predictive AI adoption is likely to grow even further in conjunction with the adoption of Generative AI.

Generative AI is yet to be adopted as widely as Predictive AI, according to our survey respondents. This is not surprising, given it is a nascent technology. The highly regulated nature of financial institutions means that a degree of caution will be taken where risks are different.


Due to the widespread vertical application of AI within organisations, AI deployment can be conceptualised based on impact rather than function (front/middle/back office). For the purposes of our report, we have categorised three broad buckets based on impact: growth and retention, productivity and operations, and risk management.

Figure 5: Predictive and Generative AI deployment within financial institutions (23 financial institutions)

Where is Predictive AI currently deployed in your institution?



Which areas in your organisation are likely to be the early adopters of Generative AI?



Sources: UK Finance members survey, Oliver Wyman analysis

The most common Predictive AI use cases are within fraud detection, risk modelling, Know Your Customer (KYC) and document authentication. For example, outlier detection tools can identify suspicious transactions by comparing them to past payments. These outlier transactions are either blocked directly or are flagged to clients, who can review them, then approve or reject the payments.

Generative AI is expected to be more prominent in both productivity and operations, as well as growth and retention functions. For Generative AI, applications include assisting with translating code between languages, document search and response within internal policy and procedures, and generating marketing content. For further case studies on the applications of Generative AI, see Chapter 6.

The business case for Generative AI is still uncertain according to our interviews and depends on the use case. Although our survey identified that growth and retention use cases will be popular in due course, at present, firms are in the experimental phase of Generative AI. Focus is on testing proofs of concept in low-risk functions that adhere to risk and compliance frameworks. The applications being tested are mostly limited to productivity and operations and risk management, with 56 per cent of survey responses indicating that they use Generative AI for process automation. Use cases in client-facing functions will be limited until the business case is proven.

3.2. IMPLEMENTATION CHALLENGES OF PREDICTIVE AND GENERATIVE AI

Predictive AI is a relatively mature technology in comparison with Generative AI and as a result the uses and risks of Predictive AI in the financial sector are better understood. According to respondents, the key blockers to the innovation process for Predictive AI are data availability and technological constraints. This view is consistent across all types of financial institutions and their level of Predictive AI adoption.

A common blocker for the use of Predictive AI, particularly among the smaller-sized institutions respondents, was that the benefits would be limited, a clear business case for its use was lacking, or that the organisation is simply not ready to adopt the technology due to limited access to appropriate data or infrastructure.

Table 4: What are the biggest blockers to the innovation process of AI? (23 financial institutions)

Predictive AI	Generative AI
1 Data availability	Concerns regarding data privacy
2 Technical maturity and constraints	Technical maturity and constraints
3 Regulatory uncertainty	Regulatory uncertainty
4 Limited benefits or clear business case	Hallucinations
5 Internal decisioning and alignment	Internal decisioning and alignment Lack of human capabilities and appropriate training

Sources: UK Finance members survey, Oliver Wyman analysis

3.2.1. Implementation challenges of Generative AI

We are still in the early days of implementing Generative AI and uncertainties exist around practical applications despite the hype. While some promising use cases are emerging, the technology's business case is still being defined by individual institutions.

The broader implementation of Generative AI in financial services faces challenges such as technical limitations, data quality, building a strong innovation foundation, recruiting the appropriate technical skill set, and navigating a constantly evolving regulatory environment. This is precisely why this exploratory stage is critical — institutions are experimenting with Generative AI's possibilities to map appropriate adoption while addressing risks proactively. To overcome these barriers, organisations are taking preparatory steps such as updating risk management frameworks, creating AI Centres of Excellence to link colleagues from various functions within the organisation and circulate knowledge, and investing in contained innovation practices (see Chapter 4).

It is important to note that current applications of Generative AI in growth and retention are still at a relatively superficial level (such as chatbots and image generation). As detailed in Chapters 2 and 5, the technological limitations of Generative AI and the associated risks are new challenges and are for now limiting the deployment of Generative AI into client-facing functions. There is concern about negative reputational and financial impacts in the event of an incident. As such, firms surveyed stated that they are taking a cautious 'test and learn' approach to mitigating risks associated with Generative AI models responsibly, before deploying them in more customer-facing applications.

3.2.2. Generative AI for sustainable competitive edge

The rise of Generative AI has led to uncertainty among organisations about how much money and attention they should be spending on this technology. While there is no simple answer, it's important to consider the potential benefits and risks of ignoring its rise. Financial services will adopt Generative AI in key functions, and customers may come to expect higher levels of service from companies that use Generative AI. At this stage, financial institutions that successfully implement Generative AI into their organisations could gain a competitive advantage that is hard to close. Firms that continue to rely on manual processes may experience higher operational costs and greater inefficiencies compared to those which use Generative AI as a tool to reduce their cost base. More details on the steps to fully unlocking Generative AI benefits can be found in Chapter 4.

3.2.3. Key considerations to think about before setting out to adopt Generative AI

The results of our survey and interviews with respondents indicated few clearly defined revenue-related business cases for adopting and deploying Generative AI. Potential customer-facing use cases are still being explored. There was much more evidence of efficiency and cost-related cases. There are several key considerations that firms should take into account before adopting this technology, illustrated in **Table 5**.

Table 5: Key considerations before investing in Generative AI

Categories	Further considerations
High build and deployment costs	<ul style="list-style-type: none"> • Readiness of existing systems for embedding of Generative AI • Dataset maintenance — large datasets of potentially private data need to be built and maintained • Model sourcing — which tool do you buy? • Model training — getting the data and training the model takes time and money • Customised tools and outputs — tools need to be specialised for industry • Employee training — misuse of tools could be a serious issue
Organisational changes	<ul style="list-style-type: none"> • Employees need sufficient training to use tools effectively • Integration into an organisation at scale is difficult and time consuming • Does the model solve the underlying problems or is it just window-dressing?
Data quality, privacy and security	<ul style="list-style-type: none"> • Data quality is vital, poor data means poor/skewed results and reduced accuracy • Data remediation tools are important to avoid bias • Data management is key to avoid potential IP infringement, loss of private data and maintain corporate security. • Key requirement to monitor/check model outputs

Source: Oliver Wyman analysis

4. UNLOCKING THE BENEFITS OF GENERATIVE AI

Key messages:

UK financial institutions are well positioned to implement AI as a 'system solution' within their organisations and unlock the full benefits of Generative AI over the typical three phases of technology adoption:

- **Phase 1:** Limited data quality and nascent capabilities and infrastructures; value will be created through optimising the current state with point solutions
- **Phase 2:** Improved data quality and access leads to the evolution of existing end-to-end journeys and use cases
- **Phase 3:** Mature data and adoption and deep integration within existing systems allow for fundamental shifts in business; large value creation happens in this phase

In their book *Power and Prediction: The Disruptive Economics of Artificial Intelligence* (Harvard Business Review, 2022), academics Ajay Agarwal, Joshua Gans, and Avi Goldfarb claim that current AI solutions are 'point solutions' that address specific pain points. The full potential of AI lies in its ability to become a 'system solution' where productivity gains are holistic and improve organisational productivity through better analysis and decision-making. This means integrating AI into an organisation's overall strategy and operations.

In this section, we will examine the steps that organisations could take over three phases of technology adoption to implement AI as a system solution.

4.1. LAYING THE FOUNDATION FOR A SUSTAINABLE COMPETITIVE ADVANTAGE

Financial institutions understand that the full potential of Generative AI lies in its system adoption across the business, prompting investment in foundational infrastructure for rapid deployment.

Sustaining a competitive edge requires high-quality data, robust infrastructure for organisation-wide Generative AI integration, and skilled maintenance, in addition to existing innovation and control frameworks.

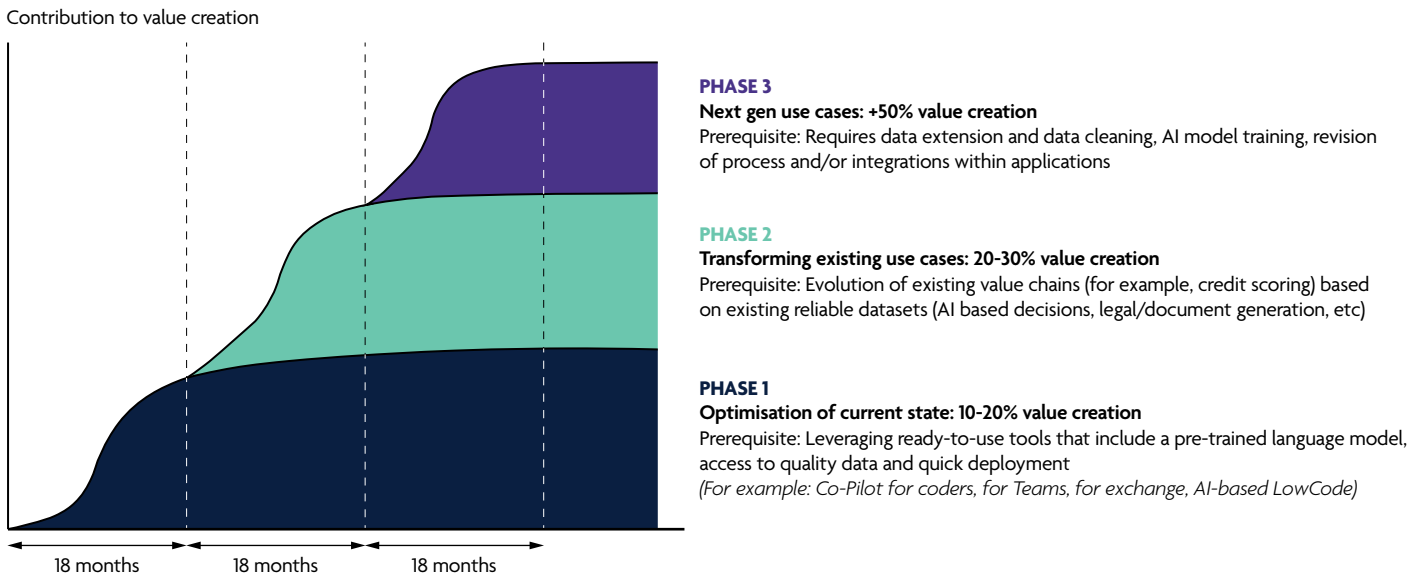
We anticipate the value creation from Generative AI to unfold in three phases, each spanning about 18 months:

1. In phase one, value is created by optimising the current state through point solutions. With limited data and developing capabilities, value creation is modest. Typical technology adoption tells us that this phase will be contributing approximately 20 per cent of the total potential.
2. In phase two, as data quality and accessibility improve, existing processes and applications evolve, generating an additional 20-30 per cent of value. This means realising about half of Generative AI's potential value in total over the next three to five years.
3. Phase three involves mature technology adoption and deep system integration, enabling significant business transformation. Here, most of the value is created, completing the value creation process.

The financial services sector, with its extensive use of Predictive AI, stringent regulatory compliance, and robust internal processes, is ideally suited for safe and effective AI implementation. More than 70 per cent of survey participants are updating their policies to meet

Generative AI's unique needs, with 95 per cent having already made changes to their risk frameworks to account for Predictive AI, showing the industry's proactive approach to responsible AI adoption.

Figure 6: Three-phase model of technology-related value creation



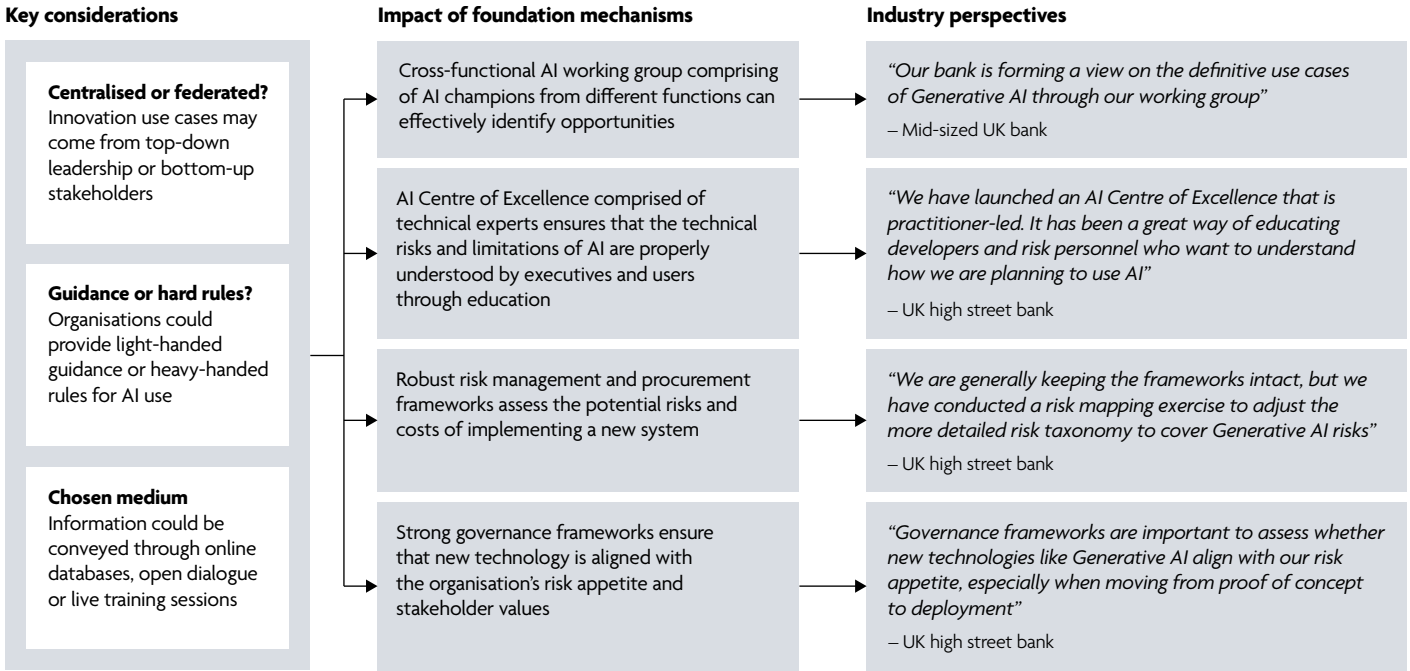
Source: Oliver Wyman analysis

4.1.1. Building core innovation capabilities is crucial for a competitive advantage

As highlighted earlier, creating foundational systems is key to gaining a sustainable competitive edge. However, for this foundation to be effective, it must be tailored to the specific needs of the organisation.

Figure 7 illustrates examples of how surveyed firms have begun to develop this innovation foundation. Most interviewed respondents have prioritised knowledge dissemination by establishing working groups and Centres of Excellence.

Figure 7: Best practice examples of setting up innovation foundations across financial institutions



Sources: UK Finance member interviews, Oliver Wyman analysis


4.2. ROADMAP TO UNLOCKING GENERATIVE AI BENEFITS

4.2.1. Opportunities are developing

Financial institutions are exhibiting caution and thoughtfulness regarding the revenue generation potential of Generative AI, as illustrated in Figure 8. Although there’s potential for Generative AI to boost revenue, companies are wary about its short- to medium-term monetisation prospects. While expecting benefits in customer service from implementing Generative AI, firms also worry about negative client reactions to AI-driven service, fearing it lacks a personal touch.

Figure 8: Perceived benefits of Generative AI within financial institutions (23 financial institutions)

What are your top three perceived benefits for Generative AI?



Source: UK Finance members survey, Oliver Wyman analysis

4.2.2. What are the potential benefits of Generative AI?

Although firms are still largely in the experimental phase of Generative AI, with 73 per cent of survey respondents in the proof of concept stage, they are identifying significant cost and efficiency use cases. Initial efforts are concentrated on internal functions, primarily process automation, with 56 per cent of survey respondents giving this as a use case. In the short term, organisations are leveraging existing productivity tools such as Microsoft Co-Pilot, GitHub Co-Pilot and

GPT and equivalents. In the medium-term, interviewees predict new use cases in growth and retention could lead to gains in product upsell through personalisation at scale. **Table 6** shows where we might see potential profit and loss benefits from the use of Generative AI.

The return on investment of Generative AI will be incremental in the short term but will scale with the number of use cases being employed alongside Predictive AI.

Table 6: Benefits of Generative AI

Financial impact	Functions	Example use cases (non-exhaustive)	Example impact (estimated)
Cost savings	Productivity and operations	<ul style="list-style-type: none"> Remove repetitive tasks Improving the uptime of systems and cost of resolving issues Greater efficiency for summarisation and insights generation Improved knowledge management through document search and retrieval 	<ul style="list-style-type: none"> Speed up code writing through code generation by 25-50% Reduce the time to resolve system downtime incidents by up to 50% Reduce costs in loan underwriting accuracy and document preparation by 5-10% Up to 30% productivity gains across analyst roles by processing information at speed and scale that were not possible before
	Risk management	<ul style="list-style-type: none"> Reduce costs of compliance by automating report preparation (for example, AML reports) Improved accuracy of models through synthetic data creation 	<ul style="list-style-type: none"> Improved fraud detection by contextualising transactions and developing fraud tests and red-flag markers which can save up to 5% of these costs Reduction of legal negotiations periods by 20-30% through summarisation of legal documentations
Revenue generation	Growth and retention	<ul style="list-style-type: none"> Improve automation of onboarding new customer or products Improve prospecting and product offering Generation of marketing materials Increase engagement and brand awareness 	<ul style="list-style-type: none"> 60% faster response time on prioritised client requests Reduction in cost of running customer service centres by 30-45%. Improved product personalisation driving retention and Customer Lifetime Value (CLV) — can lead to increased revenue by 3-5% 60% of new product documentation can be automatically generated, resulting in faster time to market for new products

Sources: The AI Tipping Point (Oliver Wyman & Morgan Stanley, 2023), A framework to assess impact of AI on US Banks (Autonomous, 2023)

5. AI RISKS AND MITIGATION

Key messages:

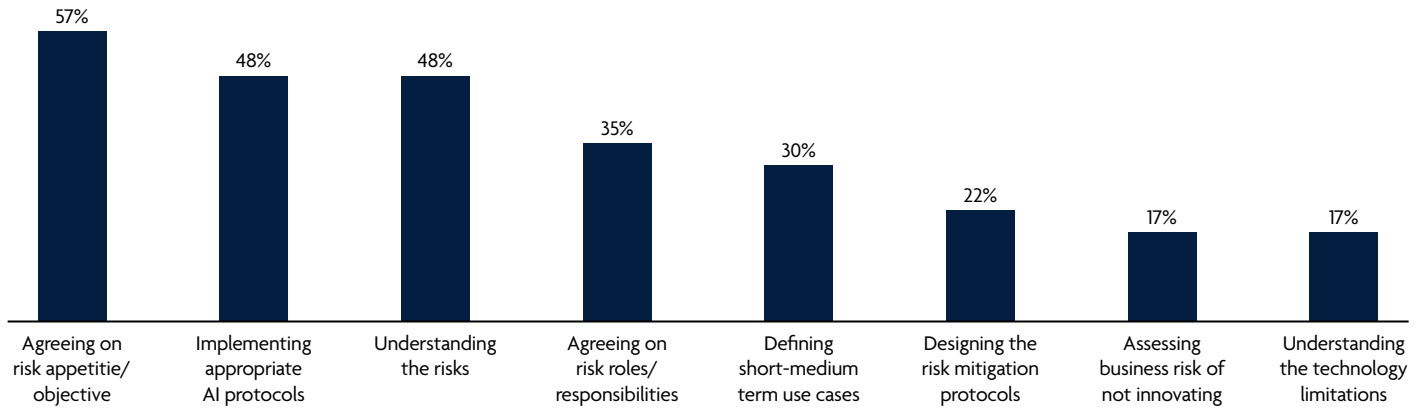
- Generative AI poses emerging risks for financial institutions, stemming from a lack of control over training data and uncertain outputs.
- Financial institutions are concerned about bad consumer outcomes, which may lead to reputational and regulatory risks, as well as risks related to intellectual property, data usage, and privacy breaches.
- To mitigate risks associated with Generative and Predictive AI, financial institutions are investing in educating and upskilling their organisations, establishing robust risk frameworks, and implementing effective vendor governance processes.
- Ongoing regulatory cooperation is critical to ensure that both firms and regulators keep pace with technological advancements and mitigate potential risks.

Financial institutions are cautiously adopting AI, particularly Generative AI, due to the evolving and complex risks involved. 95 per cent of firms surveyed said they account for AI-related risks in their risk framework, and 60 per cent said they have started building an approach to model bias and fairness.

However, with Generative AI's integration, financial institutions face new challenges that test existing risk management approaches. Key concerns include compliance with Financial Conduct Authority (FCA) Consumer Duty guidelines, customer risks, cybersecurity, data privacy, and intellectual property infringement. Potential use by staff outside of control frameworks accentuates these risks.

As regulatory landscapes and AI applications evolve, financial institutions must deepen their understanding of the risks associated with both Predictive and Generative AI to develop robust mitigation strategies and align with regulatory expectations, ensuring the safe and effective use of AI technology. As they develop use cases, the strategic value of Generative AI will emerge, guiding investment and organisation-wide optimisation.

Figure 9: Survey members’ largest challenges in addressing AI risk (23 financial institutions)



Source: UK Finance members survey, Oliver Wyman analysis

5.1. KNOWN RISKS

There are a number of known risks to both Predictive AI and Generative AI. Predictive AI has been around for several years, and financial institutions have taken steps to establish risk protocols associated with its use and application in credit risk management, fraud detection, and other use cases. Although good progress has been made, further deepening of the understanding of risk and enhancement to controls are ongoing. **Table 7** highlights a typical taxonomy of common risks associated with AI systems in general. These risks are common to both Predictive and Generative AI; however, Generative AI introduces a new layer of challenges that needs to be addressed.

Table 7: Risks pertinent to all AI systems

Risk segment	Details
Accountability and oversight	Correct management, policies, lines of responsibility and other governance measures are required in relation to AI systems to prevent unintended, unlawful or detrimental consequences
Transparency and interpretability	The complexity of AI systems can lead to difficulties in understanding and explaining the use, purpose and rationale of automated and AI-assisted decisions, whether in communications to customers, regulators or internal stakeholders
Data privacy	Inappropriate use and handling of private information can lead to data leaks or intrusive analyses being conducted
Bias and fairness	AI systems built using datasets that are inherently biased or otherwise unfair can produce similarly unfair outputs. Bias can also be introduced by AI design choices or by those interpreting the results. Additionally, it should be noted that outputs can be deemed unfair due to the way the data is used rather than any inherent bias (for instance, the courts have determined that factoring gender into motor insurance pricing is illegal discrimination)
Security	AI systems use large volumes of information, which can be lost, accessed without authorisation, damaged or destroyed, or misused for fraud or other economic crimes Anyone with access to company data may be able to inadvertently ‘join the dots’ and draw inferences using AI, which may reveal unexpected sensitive or confidential information

Source: Oliver Wyman analysis

5.2. EMERGING RISKS FROM GENERATIVE AI

5.2.1. Why is Generative AI more difficult to control?

The emergence of Generative AI and its inherent limitations (as explored in Chapter 2) has raised a fundamental question: how can the trustworthiness of AI systems be ensured? This issue has become a shared priority for the private and public sectors.

The additional uncertainty around the trustworthiness of Generative AI models stems from four main factors:

1. **Uncertain outcomes:** Generative AI exhibits unpredictable behaviour, which can undermine performance tests and risk assessments.
2. **Opaque logic and processing:** Although Predictive AI can have low explainability, Generative AI models are particularly opaque in how they generate their outputs, making it difficult to identify root causes of errors and to predict potential mistakes, and meaning that decision logic may also lose transparency.
3. **Lack of accuracy or numeracy:** As outlined in Chapter 2, Generative AI is optimised to work out probable or realistic sounding answers, rather than giving a calculated 'right' answer. As such, it is uncertain that the answer will be an accurate one. Technology and protocols still need to evolve to secure safe ranges of confidence.
4. **Third-party procurement:** These models are typically built and trained by third parties. This adds additional dimensions to manage regarding control and transparency. Firstly, financial institutions are likely to lack control over the foundation models used in Generative AI, as they are developed and hosted externally, limiting their understanding of the models' training and functioning. These challenges may be heightened if an open source model is used. There may also be concentration risk as only a limited number of vendors possess the necessary capabilities and technology to provide Generative AI solutions.

Although organisations can control the use of third-party tools internally, other Generative AI tools remain very accessible. Anyone across an organisation can use some products for free on the internet, regardless of skill or training. Such use is harder for firms to control and can lead to potential misuse.

We identify some examples of Generative AI risk for financial institutions in the following sections.

5.2.2. Bad customer outcomes

Absent effective controls, these characteristics of Generative AI can lead to bad customer outcomes, reputational damage and regulatory compliance risk through:

1. **Discriminatory or biased outcomes:** Training data or system design problems can lead to discriminatory or unfairly biased outputs in any type of AI system. The new complexity for Generative AI is that this technology can be used to produce content. Unfairly biased content can be more subtle and qualitative than statistical bias in Predictive AI use cases, and be potentially harder to test for and monitor.
2. **Unreliable or incorrect outputs:** Generative AI models can hallucinate, which is when models produce outputs that are factually incorrect but framed with a high level of confidence. In other words, hallucinatory outputs are not justified by the data the models were trained on. An example where these hallucinations can lead to negative outcomes for customers is in fraud detection, if the model produces false positives or false negatives based on its assumptions about what forms fraudulent behaviour.

5.2.3. IP, data usage and privacy breaches leading to regulatory and financial risk

1. **Copyright and IP:** Copyrighted text or media may be used as training data, tainting outputs with proprietary or protected extracts. A study found that even a well-aligned model (having gone through processes to ensure that the generated outputs are consistent with the intended goals) is still prone to copyright infringements.¹
2. **Privacy or data security violation:** AI models are vulnerable to data privacy attacks, where private information that was used in training can be extracted from the model by malicious users. A study found that personal or sensitive information can be extracted from a large language model's training data by simply asking the model to provide it.¹ This could pose a security risk to financial institutions — or employees acting on their own initiative — that wish to use internal data as inputs or prompts for Generative AI.

5.2.4. Other risks

While this section has primarily addressed risks related to consumer outcomes, intellectual property, data, and privacy, it is important to acknowledge that the risks associated with AI are not limited to these areas. There are additional risks that extend beyond the scope of this report. These include the risk of doing nothing and missing the benefits of AI, and, societal risks, such as employment effects or the potential impact of compute-intensive AI systems on sustainability targets for supply chains — all of which will drive a need to reassess business processes, employee skills, and staffing considerations.

5.3. MITIGATIONS

The risk associated with AI spans multiple disciplines and necessitates technical, mathematical, legal, compliance and risk expertise. As such, coordination across organisations is required. There is no current consensus on a single best way to mitigate AI risk. However, important mitigation techniques include:

1. Constraining Generative AI use to ‘appropriate applications’:

It’s essential to define appropriate use cases and train models with suitable datasets, as risks are closely tied to specific applications. Human oversight is also crucial, varying from approving every output in sensitive cases like marketing to monitoring performance metrics or assisting users with chatbot tools.

2. Enhancing business awareness: Firms are focusing on educating their organisations on the risks and correct usage of AI. 66 per cent of surveyed institutions have actively engaged employees through town hall meetings, AI working groups, and AI Centres of Excellence. AI skills can be highly specialised and technical or more general, such as the effective use of prompts and awareness of appropriate versus inappropriate tasks for the use of publicly available AI tools.

3. Quality assurance on model outputs: This involves understanding the technical limitations of these models and developing controls that help mitigate the risks associated with AI outputs. Firms are

educating their employees on how to effectively evaluate the ‘quality’ of outputs generated by AI models. In addition, granular testing protocols and operational mitigations should ensure responsible and ethical use of AI and effective risk management throughout the AI lifecycle. Examples of these protocols include:

- **Retrieval Augmented Generation (RAG):** The model is made to use a specific set of documents as the information source, instead of the internet or the prior model knowledge. This can enable an organisation to harness a Generative AI model but with outputs drawn solely from specific, proprietary data.
- **Guardrails:** Tools to help users enforce structure on the output from a model. This could mean preventing the model from producing content that references sensitive or unwanted topics. Example software include Guardrails AI and Nemo-Guardrails.

- 4. Clarifying training data diversity:** Whether a model is developed in-house or supplied by a vendor, firms need to have a clear understanding of the diversity of the dataset on which the model was trained on. No dataset can fully represent all people equally, but users of AI models at least need transparency regarding which segments or minorities are under-represented, in order to consider any potential for bias.
- 5. Robust internal risk framework:** Most firms have robust risk frameworks in place, and adding another layer of risk management could create complexity and be a risk in itself. Instead, adapting existing frameworks to account for AI — and particularly Generative AI — risks is likely to be the best approach. Indeed, nearly all the respondents to our survey stated that their risk frameworks already account for AI-related risks. However, most survey respondents are improving their risk frameworks to account for Generative AI. This includes AI related risk appetite definition, updates to governance structures and review of risk management enablers.

The risks associated with AI use in financial institutions are tangible and multifaceted. While self-regulation can be a useful tool, it should not be relied upon exclusively and there is scope for public sector and regulatory intervention on the topic, which is further explored in Chapter 7.

5.4. SPOTLIGHT ON THE RISK OF AI MISUSE BY BAD ACTORS




The discussion of AI risks in this paper is focused on the risk of harms being caused inadvertently by legitimate firms. However, a further risk that cannot be overlooked is the potential for AI to be used by bad actors to cause harm intentionally. These could be private individuals, criminal groups, ‘hacktivist’ organisations or state-backed entities.

Fraud

Generative AI in particular has the potential to be used by fraudsters for harmful use cases.

A plethora of tools exist that criminals can use in social engineering (the process by which fraudsters manipulate individuals to execute a transaction, provide personal information or take other actions to facilitate a fraud). But Generative AI has the potential to enable social engineering further. Examples can be seen in **Table 8**.

Table 8: Examples of Generative AI use to facilitate fraud

	Description
 Image generation	Image generation is already being used to ‘invent’ people for the purposes of fraud, particularly in romance scams, which make it harder for the victim to identify that a fraud is taking place compared to where the fraudster simply reuses the photograph of a real person
 Deepfake audio	Deepfake audio could be used to impersonate people who are known to a customer in order to convince them to make a payment. It has already been used to facilitate multi-million-dollar frauds against high-net-worth individuals. Over time it may also be possible for such fake audio to defeat voice-based identity verification systems
 AI text generation	AI text generation could facilitate the production of more convincing phishing messages at greater speed than is currently possible manually

The above techniques — in conjunction with real-time deepfake video — could in the future be used to manufacture an entire persona capable of video calls. This could facilitate not only social engineering but it could also be difficult for firms’ identity verification controls to detect, enabling, for example, the opening of fraudulent bank accounts.

Such fraud techniques exist already without Generative AI, but these technologies can make it easier for individuals with limited expertise in committing fraud to produce content more cheaply and quickly.

The current limitations of Generative AI include imperfect human images and restricted access to advanced tools. As technology progresses and costs decrease, AI-enhanced fraud is expected to become more prevalent in the mass market.

Cybersecurity

There is broad acceptance among the cybersecurity community that Generative AI tools have lowered the barriers to entry into some technical attack methods, including:

- **Malware creation and modification:** In July 2023, the US Federal Bureau of Investigation warned that malicious actors were using Generative AI to generate, modify and enhance malware, a task formerly the preserve of highly skilled actors. Modifying malware complicates its detection by antivirus software, as it may not have a recognisable pattern or signature.
- **Vulnerability identification:** Generative AI can help malicious actors build automated vulnerability identification tools.
- **Password cracking:** Malicious actors can utilise Generative AI to create tools for the effective cracking of passwords, such as lists of potential passwords tailored to a specific target.

Mitigating these threats

Potential mitigations to minimise risks include:

- Developing defensive AI tools to better detect cyber risks, fake images, or impersonations and to better identify high-risk suspicious transactions.
- Enhanced information and intelligence sharing between firms and between sectors to gain a broader and clearer view of risks across the economy and facilitate the identification of suspicious activity.
- Introducing further layers of protection in customer authentication, such as applying additional identity verification steps with voice ID.
- Updating public education messaging to increase awareness of new risks and social engineering techniques.
- Further development of guardrails by developers of Generative AI to protect against misuse by users.
- Greater use of privacy enhancing technologies to better protect sensitive information.

Existing collaboration forums between the public and private sectors will need to keep up to date with technological changes, ensuring that the latest typologies and risk information are shared, and best practices developed.

6. STRATEGIC USE OF AI: CASE STUDIES

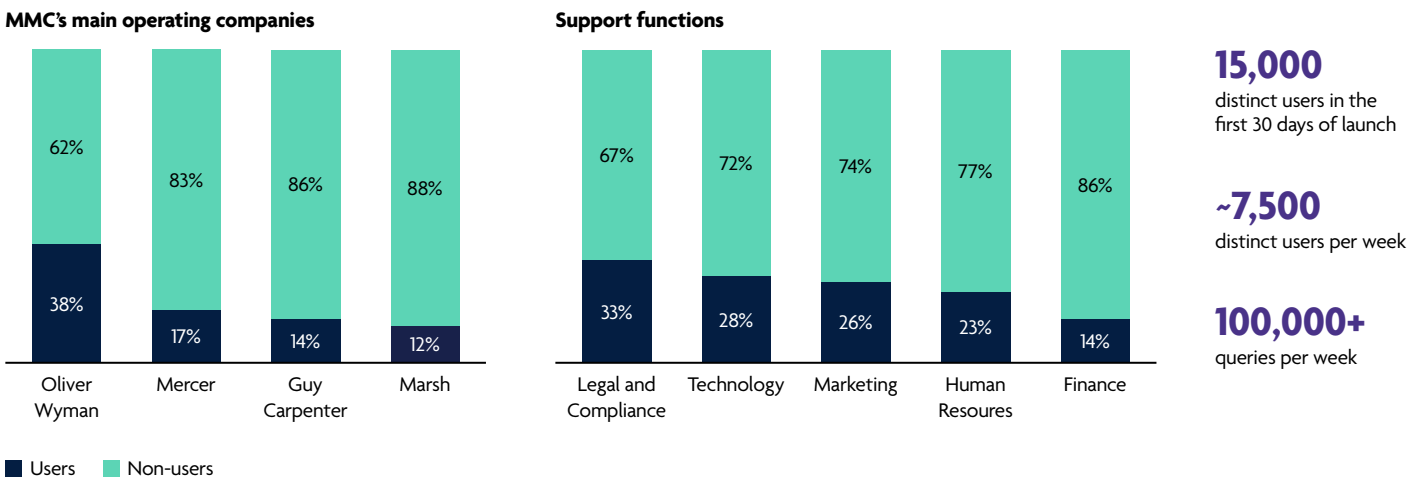
As seen in Chapter 3, the adoption of AI is well underway in the financial services sector. This chapter provides two examples of AI implementation in financial institutions. The first case study involves the deployment of a proprietary Generative AI co-pilot tool across Marsh McLennan (MMC) offices. The second case study examines the collaborative development and implementation of a Predictive AI Anti-Money Laundering (AML) product by Google Cloud and HSBC.

6.1. CASE STUDY 1 — GENERATIVE AI TOOL: MARSH MCLENNAN’S LENAI

In September 2023, MMC officially launched its proprietary Generative AI assistant called LenAI. The tool provides MMC colleagues with all the capabilities of ChatGPT, using GPT 3.5 as its underlying model, while ensuring the organisation’s data and information remain secure within MMC’s cloud environment.

LenAI had a significant adoption within MMC, reaching 15,000 distinct users across MMC business units within the first 30 days of the launch. The extent of LenAI deployment is shown in **Figure 10**. As a result of this large-scale deployment, numerous proofs of concept were launched, and approximately 200 feasible and scalable use cases were identified, including coding assistance, document summarisation, and supplementing brainstorming processes.

Figure 10: Extent of LenAI adoption across MMC business units during first 30 days of launch



Sources: MMC, Oliver Wyman analysis

MMC is still in the early stages of LenAI deployment. As such, the full cost-saving and productivity benefits of LenAI are continuously evolving and are not yet defined. However, as shown in **Table 9**, the preliminary benefits of LenAI in selected use cases show promising

potential. In addition, early users across MMC are feeling optimistic about LenAI’s usefulness, which will further drive the integration of the system with their day-to-day processes, as shown in **Figure 11**.

Table 9: Early stage financial and productivity benefits from selected LenAI use cases

Use case	Description	Size of task	Stage of deployment	Benefits experienced
Executive compensation Global Disclosure Database	To extract executive compensation data from unstructured lengthy documents	50 fields with over 10,000 unstructured documents	MVP	<ul style="list-style-type: none"> • 88% accuracy • Estimated \$450,000 saving over three years
Health Census POC	To upload census information and transform it into Mercer’s taxonomy	29 fields, with seven different languages	POC	<ul style="list-style-type: none"> • 95% accuracy
New Zealand Invoices	To reduce the processing time of incoming invoices by extracting all invoice information	10,000 invoices in the same invoice template	Production	<ul style="list-style-type: none"> • Over 100 hours saved per year

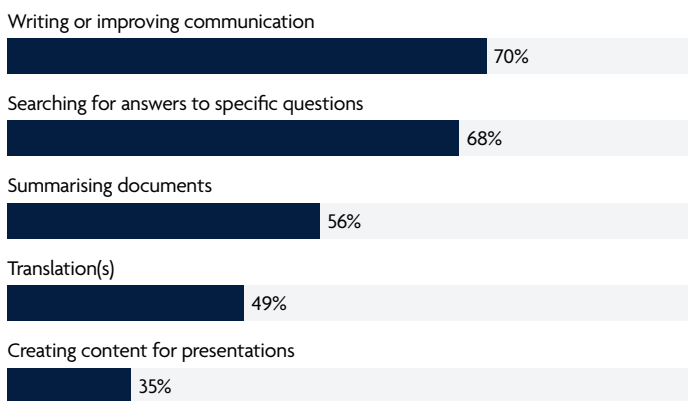
Source: MMC, Oliver Wyman analysis

Figure 11: Survey results measuring adoption and usefulness among early users of LenAI

LenAI’s adoption and usefulness



Top 5 uses of LenAI in day-to-day activities



Source: MMC, Oliver Wyman analysis

The decision to develop LenAI was prompted by the recognition of the potential benefits of large language models after the introduction of ChatGPT in November 2022. As MMC initiated an evaluation of the business case for a widely used Generative AI tool, various technology




providers were considered. However, it was decided to leverage Microsoft’s OpenAI API to create a cost-effective and secure large language model solution within MMC’s private cloud environment.

The successful early adoption of LenAI within MMC was the result of a strategy that prioritised a quick and broad roll-out in low-risk applications (and with explicit guidance to quality check outputs), instead of following a linear process of identifying specific use cases and conducting narrow testing. This approach promoted widespread experimentation with the technology as a way to unlock immediate productivity gains and explore the potential for further use cases. A critical enabler for this roll-out strategy included the implementation of firm-wide educational measures on the correct use of LenAI (for instance prompt engineering training) and risk mitigation techniques (such as training on model limitations and methods of checking outputs).

MMC is still in the early stages of deployment and development of LenAI. Further adoption across the organisation is expected as users get more comfortable with harnessing the technology. 79 per cent of LenAI early users across MMC state that they are discovering more ways to use LenAI through experimentation. Additionally, MMC is continuing to track the development of current and future use cases through user surveys and continuous colleague feedback.

To date, the expected current and future use cases can be segmented into three categories: information retrieval, information summarisation, and work product creation. A non-exhaustive list of illustrative current and future use cases in these categories is illustrated in **Table 10**.

Table 10: Use cases and capabilities of LenAI (non-exhaustive)

Capabilities	Current use cases (non-exhaustive)	Future use cases (non-exhaustive)
 Information retrieval	<ul style="list-style-type: none"> • Search documents and answer questions based on uploaded files • Extract quotes and commentary from interview transcripts 	<ul style="list-style-type: none"> • Have a history of search and the ability to save specific conversations • Translate documents
 Information summarisation	<ul style="list-style-type: none"> • Summarise documents and presentations for meetings • Summarise meeting transcripts • Clean up transcripts to create detailed interview notes • Summarise differences between versions of documents (for example, reinsurance contracts) 	<ul style="list-style-type: none"> • Larger capacity to read through documents, such as larger PDFs • Read and analyse articles available online and provide accurate answers with sources
 Work product creation	<ul style="list-style-type: none"> • Code generation • Supplement brainstorming processes • Speed up broad research such as gathering basic business information • Proofreading and rewriting emails and other documents for improved communication • Conducting a pre-mortem analysis for a current proposal or project • Analysing support tickets to find resolution steps • Produce interviews/surveys 	<ul style="list-style-type: none"> • Image generation • Flow chart creation • Generate table responses • Create PowerPoint presentations or Word documents from extracted information • Analyse survey submissions

Source: MMC, Oliver Wyman analysis

Aside from expanding the number of use cases, MMC is also working to improve LenAI's core capabilities to include a wider range of readable file types, personalised recommendations, more advanced natural language processing capabilities, prompt sharing functionality, and the creation of prompt libraries organised by department.

6.2. CASE STUDY 2 — GOOGLE CLOUD’S AND HSBC’S ANTI-MONEY LAUNDERING (AML) AI

The growing scale and complexity of financial crime compliance poses a challenge for traditional rules-based transaction monitoring systems.

In 2021, HSBC partnered with Google Cloud to introduce a cutting-edge AML dynamic risk assessment (DRA) system. Powered by Google Cloud’s AML AI, this advanced solution is trained on HSBC’s production data and undergoes rigorous validation testing. By analysing live and historical data, including transactional patterns, network behaviour, and Know Your Customer (KYC) information, the system generates risk scores for groups of retail and commercial customers. This enables the identification of financial crime cases and streamlines the investigation workflow. The tool is designed to adapt to changes in the underlying data, resulting in increasingly accurate outcomes. Leveraging cloud technology, the solution reduced overhead costs and complexity while ensuring the bank’s customer data remains encrypted and protected.

The DRA represents a significant advancement in AI tools for financial crime detection, being more effective and efficient than traditional rule-based monitoring systems. In the UK market, the bank performs AML analysis on approximately eight billion transactions across 63 million accounts monthly. Table 11 summarises the benefits of the DRA in comparison to traditional systems.

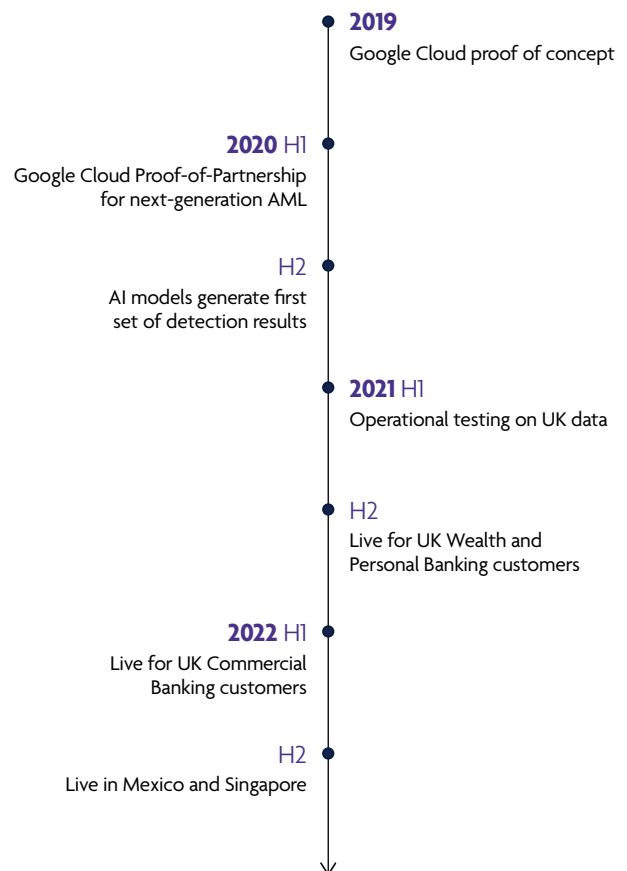
Table 11: Benefits of using the DRA, powered by Google Cloud’s AML AI

Benefit	Description	Results
Speed	DRA significantly improves the speed of data analysis and result generation	<ul style="list-style-type: none"> Reduced batch analysis cycle time from ~30 days to two to three days Results generated faster, in under 12 hours
Accuracy	DRA outperforms traditional systems by leveraging transaction flows and other parameters to detect complex typologies	<ul style="list-style-type: none"> The bank is able to detect two to four times more ‘true positive’ risk, versus a traditional system Able to identify new typologies of suspicious behaviour (for instance, misuses of business loans)
Efficiency	DRA generates significantly fewer alerts than traditional systems, reducing the level of ‘noise’ from false positive Suspicious Activity Reports	<ul style="list-style-type: none"> Alert volumes decreased by more than 60%, reducing wasted investigator time. Further improvements in recent months have been observed as the model learns from newly detected suspicious activity Enhanced customer experience by minimising the need to engage with customers on false positive alerts

Concerns about the governance of model risk, and about explainability, have limited uptake of fully AI-based transaction monitoring among some institutions. Google Cloud’s AML AI solution built in strong governance by providing compliance functions with auditable and explainable risk scores to support regulatory compliance. Privacy and data security concerns were addressed by having data stored in HSBC’s Google Cloud project environment, encrypted with its encryption keys while at rest and in transit.

The proof of concept phase for this solution began in 2019 and reached production within the UK in the second half of 2021. The DRA is currently live in the UK, Mexico, and Singapore. To qualify for the full launch of DRA in new markets, operational testing is carried out to verify that the model outperforms the incumbent AML system using a variety of key metrics. A timeline of these efforts is shown in Figure 13. Most recently, Google Cloud launched AML AI in June 2023 to enable other banks to adopt a similar programme in three to nine months, building on the experience during the partnership between HSBC and Google Cloud.

Figure 13: Timeline of DRA project launch



Source: Cloud-Based Financial Crime Detection at Scale (Celent, 2023)

Sources: Cloud-Based Financial Crime Detection at Scale (Celent, 2023), Google Cloud Launches AI-Powered Anti Money Laundering Product for Financial Institutions (Google Cloud, 2023)

7. POLICY AND REGULATORY LANDSCAPE

Key messages:

- The financial services industry in the UK has expressed a preference for an outcomes-focused, principles-based regulatory approach to AI risks. This approach is likely to have the flexibility to accommodate the varying use cases of AI technology across different industries, while still ensuring that AI risks are addressed effectively.
- Although the UK's AI Whitepaper sets out a provisional approach to AI regulation, several policy questions require further consideration. In particular:
 - How will authorities ensure the scope of AI guidance is clear?
 - How to ensure that firms deploying AI have access to the information they need from AI providers, while respecting intellectual property concerns?
 - How will potential tensions between regulator's expectations be managed?
- Against the backdrop of different approaches emerging in both the United States (US) and the European Union (EU) that have potential extra-territorial implications, sustained international cooperation is crucial to drive compatibility of regulation where possible.
- Collaboration between industry and authorities can help build and disseminate best practices in AI risk mitigation.

7.1. CURRENT STATE OF AI REGULATION

The regulatory landscape for AI is nascent and still being debated. The EU moved first with detailed draft legislation, while an ambitious US executive order set in motion a series of responses from the private sector, government agencies and departments across all sectors. In contrast, the UK has adopted a gradual approach to AI regulation, with consultation ongoing.



There is wide interest in harmonising policies and promoting international collaboration between regulators. This is particularly relevant, given the EU AI Act has extraterritorial reach, as may eventual US regulation, standards and guidance. They are therefore likely to have a material impact on UK businesses.

7.1.1. Contrasting the approaches of the UK and the EU

The UK has provisionally chosen a principles-based approach to regulating AI. The UK government released its AI Whitepaper in March 2023, which — while still subject to finalisation — outlines a decentralised approach that is intended to be pro-innovation.

The EU has taken a cross-sector application-based approach towards the development of regulations for AI, which is currently the most advanced legislative process globally. The EU AI Act², proposed by the European Commission in April 2021, is intended to complement other EU and member state laws, such as the General Data Protection Regulation (GDPR) and the EU AI Liability Directive³, which is a non-contractual civil liability mechanism aimed at removing barriers to redress when harm has been caused by AI systems.

Table 12: A comparison of the UK's AI Whitepaper and the EU's draft AI Act

Design element	 UK (AI Whitepaper)	 EU (draft AI Act)
Regulator	<ul style="list-style-type: none"> • Sectoral and coordination function 	<ul style="list-style-type: none"> • New cross-economy AI authority and rules
Legal instrument	<ul style="list-style-type: none"> • Central government guidance to regulators, with potential statutory requirement to have due regard to the guidance 	<ul style="list-style-type: none"> • Primary legislation, with supporting guidance
Rules	<ul style="list-style-type: none"> • Outcomes focus encouraged • Regulatory guidance from existing regulators • Regulators to determine targeted use cases and applications 	<ul style="list-style-type: none"> • Governance requirements set in legislation according to risk category • Four effective risk categories: <ul style="list-style-type: none"> – Prohibited use cases, for instance social scoring – High risk use cases, such as creditworthiness assessment or recruitment — ex ante conformity assessment, monitoring and reporting – Applications featuring human-like interaction, content generation or biometrics — transparency obligations – Other AI — no requirements • Key concepts and thresholds such as adequate interpretability, fairness and transparency are not defined, being dependent on the application and context
AI definition	<ul style="list-style-type: none"> • No hard definition but focus on 'autonomous' and 'adaptable' systems 	<ul style="list-style-type: none"> • Explicit definition
Role of standards	<ul style="list-style-type: none"> • Encourages use of standards 	<ul style="list-style-type: none"> • 'Safe harbour' standards to be developed
Approach to Generative AI	<ul style="list-style-type: none"> • Intention signalled to clarify IP law. • 'Deepfake risk' noted • Generative AI noted as a potential pilot for a cross-sectoral regulatory sandbox 	<ul style="list-style-type: none"> • Requirements for foundation models have been proposed, including risk management, design principles, disclosure of copyrighted material used
Penalties	<ul style="list-style-type: none"> • No new penalties 	<ul style="list-style-type: none"> • Potential fine of up to 7% of global turnover for lack of adequate governance

Source: Oliver Wyman analysis

Financial services and AI regulation

In the UK the Bank of England (BoE) updated its model risk management guidance in 2023, though this is not exclusive to AI models, building on long-standing risk management guidance. In October 2022, the FCA and BoE published Discussion Paper 5/22 seeking feedback on safe and responsible AI adoption, including the role of policy and regulation. The discussion paper aimed to surface some of the challenges of regulating AI, including the need for sector-specific definitions of AI, prioritising risks and benefits, and updating existing regulations to support the safe adoption of AI in the financial services sector. A feedback statement (FS 2/23) summarising the views expressed by respondents was published on 26 October 2023, and broadly reinforces many of the points raised in this work.




The EU's AI Act does not distinguish between sectors, though creditworthiness assessment is a high-risk category. Uniquely, prudential regulators are tasked with supervision, rather than the AI Authority in respect of this use case at banks. However, it is unclear for now how the Act will interact with sector-specific supervisors, including under the Single Supervisory Mechanism in financial services.

Liability

In the UK, the Whitepaper noted that liability in AI supply chains is complex and asked for feedback on how best to tackle challenges under the UK's regulatory approach.

In the EU, the draft AI Liability Directive increases the responsibility and potential liability of developers, providers, users, manufacturers, and importers of AI systems. Prior to this directive, fault-based liability rules required individuals to prove negligent or intentionally damaging acts or omissions to seek compensation for damage, which was challenging when it came to AI systems. The directive is intended to simplify the claims process by introducing the presumption of causality and the right of access to evidence.

7.1.2. Comparison of approaches in other jurisdictions

Jurisdiction	Regulatory approach	Papers published	Specific governance/guidance
USA 	<ul style="list-style-type: none"> President Biden's Executive Order requires government agencies and departments to take certain actions on AI risks Covers sector-specific requirements (for instance biological, healthcare, education, defence, critical infrastructure) and 'horizontal' requirements across sharing safety test results, data privacy, advancing equity and civil rights, R&D, cybersecurity, and supporting workers In addition, it orders government bodies to accelerate their upskilling and adoption of AI techniques, and encourages international cooperation 	<ul style="list-style-type: none"> Blueprint for an AI Bill of Rights: released by the White House Office of Science and Technology Policy, outlining principles to guide AI use and potential regulations⁴ National Institute of Standards and Technology framework to better manage risks associated with AI 	<ul style="list-style-type: none"> Securities Exchange Commission proposed rules to prevent the use of 'predictive data analytics', which includes AI, in a way that prioritises the firm's interests over those of its investors⁵ Consumer Financial Protection Bureau has provided guidance for lenders using AI in credit decisions through Circular 2023-03⁶
Singapore 	<ul style="list-style-type: none"> Singapore has decided not to implement economy-wide AI regulations⁷ Specific sectors provide guidance on best practices 	<ul style="list-style-type: none"> The Personal Data Protection Commission (PDPC) released the Model AI Governance Framework in 2020⁸ PDPC and Infocomm Media Development Authority (IMDA) developed AI Verify, an AI governance testing framework and a software toolkit⁹ Monetary Authority of Singapore (MAS) launched the Veritas Initiative, which includes five whitepapers exploring the application of the 'FEAT' responsible AI principles¹⁰ 	<ul style="list-style-type: none"> Veritas is designed to promote the responsible use of AI and data analytics in the financial sector and ensure that the industry remains competitive and innovative A key feature is the co-production of case studies and other products to illuminate good practice with an industry consortium
China 	<ul style="list-style-type: none"> Broad approach that is focused on specific areas of concern, with specific regulations governing the use of AI in specific contexts 	<ul style="list-style-type: none"> In April 2023, the Cyberspace Administration of China (CAC) released measures aimed at regulating Generative AI services in mainland China¹¹ The CAC has released a draft of these measures to the public in an invitation for comments Earlier publications included a 2021 regulation governing recommendation-making algorithms used in online information services 	<ul style="list-style-type: none"> The regulations released by the CAC would broadly apply to financial institutions The approach to AI governance is unique as it centres on specific algorithms. For example, companies may need to lodge multiple filings for the same app, each addressing different algorithms used

7.1.3. Case study: Singapore's MAS Veritas initiative

MAS set up a programme called the Veritas initiative, combining regulatory support, open standards and a sandbox environment to ensure that the existing regulatory regime is evolving in line with technology. As part of the initiative, a consortium of industry players provided open-source code and provide financial institutions with assessment methodologies to comply with the Fairness, Ethics, Accountability, Transparency (FEAT) principles.¹² Additionally, the consortium published a whitepaper detailing best practices from financial institutions that piloted the integration of the methodology. The initiative has been met with positive reviews from both private and public stakeholders, as it effectively allows regulators to learn and understand the technology, enabling them to develop policies and regulations to accommodate, supervise, and control sectoral innovation. Financial institutions have stated that the Veritas initiative showcases how public-private partnerships can lead to more regulatory certainty and positive outcomes for institutions and regulators.

7.2. FINANCIAL SERVICES SECTOR VIEWS ON AI REGULATION

As seen in FS 2/23, UK financial institutions generally support a principles-based and outcome-focused approach to AI regulation that can accommodate the specific use cases and applications of AI technology within different industries. This resonates with surveyed participants, among which there is agreement that any effective regulatory framework must be flexible enough to adapt to the changing landscape of AI technology, particularly as Generative AI use and risk management evolves. There is also agreement that efforts to regulate for AI risks should not duplicate rules that are already in place, noting that sectors such as financial services are already heavily regulated.

A principles-based approach allows for flexibility in how outcomes are achieved. This approach enables financial institutions to optimise customer and business outcomes while still mitigating risks effectively. As such, the financial services industry — on the whole — supports the UK AI Whitepaper's sectoral, risk-based approach focused on guidance and outcomes. Given the criticality of avoiding inconsistent expectations from different authorities, there is support for a central coordination function and development of multi-authority sandboxes.

FS2/23 responses are aligned with our surveyed participants' views that largely favour regulators producing guidance focused on areas where there is uncertainty, rather than an 'AI overlay', which would risk being duplicative of technology-neutral rules.

7.3. POLICY CONSIDERATIONS AND TOPICS FOR FURTHER DISCUSSION

Our survey identified that regulatory uncertainty is one of the top factors slowing AI uptake, both in relation to AI in general and to Generative AI specifically. This uncertainty does not primarily relate to specific rules, per se. Rather, firms feel constrained by uncertainty over what the regulatory framework will look like, how it will operate in practice and how certain key challenges will be solved. Similar themes were apparent in FS 2/23 and also remain unresolved in other jurisdictions. Not all have a clear solution, and statutory and non-statutory options exist.

7.3.1. The definition of AI and clarity as to regulatory scope

The absence of a precise legal or regulatory definition of AI may potentially hinder firms' ability to correctly triage and assess use cases, conduct impact assessments, update risk frameworks, and perform diligence on third-party contracts. It is evident that regulations need to have a clear scope — at a minimum to make clear which systems are not in scope — but this can be provided in different ways.

Statutory option

Defining AI in law could clarify legal compliance for firms, but amending statutes takes time and political effort. There's a risk that this definition may become outdated or misaligned with high-risk systems in practice.

This option seemed not to be favoured in the feedback in FS 2/23.

Non-statutory option

Guidance from central government or regulators could clarify the traits of high-risk AI systems. This approach offers more flexibility and allows regulators to concentrate on evolving areas of risk or uncertainty as technology and best practices advance.

This option could lead to differences between sectors, potentially complicating compliance for firms operating across the economy. (See also 7.3.5.)

7.3.2. Gaps in relation to other sectors

The existing frameworks and regulations that govern UK financial services provide regulators with the necessary authority to set rules and guidance to manage risks in the sector, coupled with strong enforcement and supervision powers. This can include amending rules and setting guidance to account for evolutions in AI risk as envisioned by the current consultation process initiated by the FCA and BoE, and address any potential gaps in AI regulation. However, not all sectors or applications have the same level of scrutiny and oversight. Publicly available Generative AI is perhaps an example: its deployment in a regulated sector like financial services would be subject to the strict sectoral rules that apply. However, use in some other sectors — or by the public — would be subject to less regulatory oversight. It is important to note that any interventions to plug a gap would require a clear gap assessment before implementing new requirements.

Statutory option	Non-statutory option
<p>Legislation could define rules for unregulated, high-risk AI applications and assign a regulatory body. However, premature legislation might inhibit innovation and struggle to keep pace with fast-moving market developments.</p>	<p>Current cross-sector regulators like the Competition and Markets Authority (CMA) and the Information Commissioner's Office (ICO) could effectively manage AI risks in unregulated sectors. Coordination mechanisms like the Digital Regulation Cooperation Forum (DRCF) model involving CMA, ICO, Office of Communications, and FCA could assist. However, if statutory rules are ultimately necessary, they risk being delayed by such efforts, depending on authorities' progress.</p>

7.3.3. Assurance challenges between AI providers and AI deployers

AI systems and algorithms are complex and can lack transparency and explainability, making due diligence by firms procuring and using the products and services more difficult. This is likely to be particularly true of Generative AI products.

Due diligence could be facilitated by having AI providers give detailed product information to firms deploying AI, for example in relation to model design, training and function. However, AI providers are understandably reluctant to disclose source code or other commercial intellectual property. Consequently, firms wishing to deploy AI may be reluctant to do so, due to uncertainty as to how — or whether — it can be used in compliance with regulations. While risk management frameworks are already in place and are being updated by firms, these challenges raise questions about how firms should conduct audits of their third-party providers and procure AI models.

Low transparency by providers regarding their models may also inhibit firms' ability to use multiple providers' products in an interoperable way, and may reduce their ability to substitute one product for another.

While acknowledging that financial institutions are ultimately responsible for their decisions, sourcing models, agreements, and other related matters, it is essential to keep pace with the potential scale of liability that AI models — particularly Generative AI — could create.

These challenges may be heightened if an open source model is used.

The market may produce solutions to these challenges with vendors providing best practice product information over time. Nonetheless, public sector-led options also exist.

Statutory option	Non-statutory option
<p>Statutes could impose specific requirements on AI providers, such as setting minimum standards or mandating information provision to product deployers, akin to the EU AI Act.</p> <p>Assurance processes could also be streamlined by initially presuming provider liability for harms, similar to the EU's AI Liability Directive.</p> <p>However, there is a risk that statutory obligations may not keep pace with industry practices and cause unnecessary costs. Overly burdensome or poorly calibrated obligations on AI developers could also deter investment and innovation in the UK.</p> <p>The US executive order, though different and not yet clearly defined, requires AI system developers to share safety test results. This might impact UK financial services providers.</p>	<p>Regulators could impose transparency obligations on third party providers, either through financial sector-specific rules or by setting technical standards. For instance, SS2/21 on outsourcing and third-party risk management could be revised and broadened for this purpose. However, like statutory approaches, this could deter innovation and investment, and make the financial sector less attractive to AI developers compared to other industries, possibly shifting their focus.</p> <p>Another non-statutory intervention involves public sector bodies coordinating the development of best practices in product documentation for assurance, aligned with emerging technical standards. Collaborative work between AI developers and financial institutions could forge assurance mechanisms for AI models' processes and outcomes. Such an approach would be more adaptable than statutory or regulatory solutions. An example is the AI Assurance programme by the Centre for Data Ethics and Innovation (CDEI).¹³</p>

7.3.4. 'Off limits' AI applications

AI has a wide range of potential applications, some of which will be outside of what society considers ethically acceptable. There may be benefit in clarifying that certain AI applications are simply prohibited. Although there are no obvious financial sector examples, an example of a widely prohibited use in society could be applications of AI for mass surveillance.

Statutory option	Non-statutory option
A list of AI applications could be prohibited by statute, in a similar way to the approach under the EU's AI Act. Conceivably, regulators in the US might similarly prohibit certain use cases or applications, depending on how they implement the new executive order. However, there is a risk that uncertainty over the definitions of prohibited applications could inhibit beneficial innovation and also not be responsive to market developments.	In practice, inappropriate or prohibited AI applications may already be in breach of existing fundamental laws and rights, such as privacy. Regulators could produce guidance on any clear bright lines. This would be more able to adapt to address emerging uncertainties than a statutory approach. However, statute would provide a great degree of certainty.

7.3.5. Harmonisation of regulations and regimes

A key challenge in AI regulation is avoiding a situation where differences in rules and expectations of different regulators create tensions or contradictions. Such a situation would create uncertainty for firms about how to use AI in a compliant way, inhibiting their willingness to innovate or invest. This challenge is particularly acute for firms operating across multiple sectors. This is also a risk for single-sector firms, if there are tensions between the expectations or rules of sectoral regulators and horizontal regulators such as competition, data protection or — if applicable — an AI regulator.

FS 2/23 identified industry concerns over the application of the AI fairness principle. This concern is echoed by surveyed participants, who noted that this is a particularly likely area of future tension.

Fairness is a core requirement under existing regimes, such as FCA rules and the GDPR, as well as being a part of the UK Equality Act anti-discrimination rules enforced by the Equality and Human Rights Commission. Over time, priorities relating to fairness among regulators may diverge. For instance, if the FCA or other sectoral authorities extend their expectations of firms to include progressively more intrusive monitoring of customers, for example to identify and support vulnerable individuals, there is a risk of tension with fairness in the context of data protection rules. Similarly, some regulators may start to expect the use of given fairness metrics, which might be in tension with the prohibition against 'positive discrimination' in the Equality Act.

Other potential areas for tension relate to the principles of transparency and explainability, contestability and redress, and accountability and governance. FCA and BoE expectations in these areas may diverge over time from ICO's interpretation of GDPR requirements, notably in relation to 'automated decision making'.

Statutory option	Non-statutory option
Nominate a single authority to be responsible for all AI regulation, or create a new authority for that purpose. This could ensure that firms only have one source of AI rules, removing the potential for conflict. This could look like the approach under the EU's AI Act.	The UK's AI Whitepaper proposes the creation of a 'central function' to coordinate different regulators and manage cross-sectoral issues. This could work effectively in principle, but considerably more development of the proposal is needed. The central function will also need to contend with the large number of regulators to coordinate and the necessity of maintaining the independence of regulators.
In practice this approach might not be effective. The regulation of the AI might not be readily separated from the regulation of the application to which the AI is applied so tensions between the AI authority and other regulators could arise. For example, there are concerns in the EU that the AI Act's requirements that datasets be representative could conflict with capital requirement regulations.	Recent efforts by the FCA and ICO to provide a common view on tensions between data protection and conduct rules, and the maturing of the DRCF, are promising steps, although other sectors would also need to be covered.

7.3.6. International alignment

Firms operating in multiple jurisdictions, with customers abroad, or with other points of contact with third countries will need to find an approach to complying with AI regulations that meets all requirements. In the UK, the extraterritorial reach of the EU AI Act is particularly relevant, and any extraterritorial obligations from the US over time would be too. Consideration must be given to how rules will interoperate across jurisdictions and which steps should be taken to reduce or avoid fragmented rules.

In addition to AI-specific rules, more generic requirements can also be a point of tension between jurisdictions. For example, laws vary between the UK, EU and US relating to how different patterns of outcomes between protected groups are to be managed. In the UK it is illegal to ‘positively discriminate’, while in the US there is an expectation of proportional representation between different groups (although that may evolve significantly under the ‘Advancing Equity and Civil Rights’ sections of the executive order). In the EU the focus is on minimising misclassification. It is seldom mathematically possible to meet all three tests, meaning that AI fairness tools need to be recalibrated for use in each jurisdiction. Creating wholly distinct regimes in each jurisdiction is likely to be costly and inefficient. A single group-level approach facilitates more cost-efficient and less bureaucratic compliance, facilitating the beneficial uptake of AI. This is more straightforward if there is greater commonality between different regimes globally.

In this domain, whether or not statutory options are pursued, there are valuable international steps that authorities should take to drive alignment internationally.

Statutory option	Non-statutory option
<p>Given the proximity of the EU, and the EU’s position as having the most advanced AI law, one option would be for the UK to put in place its own horizontal AI statute aligning closely with the EU AI Act. This would streamline compliance for firms operating in both the UK and EU.</p> <p>However, this would also mean forgoing the potential benefits of a more flexible, sector-driven regime as is currently the de facto UK framework. It would also potentially put the UK out of step with other AI statutes that might be put in place elsewhere.</p> <p>This would also not address inconsistencies between countries’ rules that have an important impact on AI but do not strictly come from AI-specific regulation, such as variations in anti-discrimination law.</p>	<p>Ongoing discussion in international forums between countries can help draw positions closer over time. This can include building common principles, with acceptance that they may be achieved in different ways.</p> <p>There is clearly an important role for the development of technical standards, such as the ISO 42000 family of AI Standards. Such market-led tools can be applied — when appropriate — by firms in different jurisdictions to help demonstrate regulatory compliance.</p> <p>Regulators and industry should participate in standards-setting processes to inform progress and to themselves learn from the latest AI developments and best practice.</p>

In this domain, whether or not statutory options are pursued, there are valuable steps that authorities should take to drive alignment internationally. There have been several recent positive steps forward, including:

1. The G7 Hiroshima process, which produced guiding principles for AI, alongside an international code of conduct for developers of advanced AI systems, with a commitment to continue collaborative efforts. This is a positive step.
2. The UK’s International AI Safety Summit, which resulted in the Bletchley Declaration, signed by 28 countries including the UK, EU, US and China, and recognising the importance of understanding and mitigating the risks posed by ‘frontier AI’ in order to seize the opportunities presented by the technology. A further series of AI Safety Summits has also been agreed to maintain progress.

7.3.7. Practical steps for regulators to provide certainty and maintain guardrails in a pro-innovation way

There are two key pro-innovation techniques that can help provide certainty to firms: use of sandboxes and public-private collaboration on best practice.

Regulatory sandboxes

Establishing sandboxes is a practical and effective action regulators can take to provide certainty and help identify how to maintain guardrails in a pro-innovation way. Sandboxes provide a controlled environment for testing new technologies. The benefits accrue to the firm directly using the sandbox and also enable regulators to build up their knowledge of the state of the art, identify areas of regulatory uncertainty or tension, and issue good practice notes and reports that can benefit the whole sector.

There is also recognition of the positive impact of a newly announced UK multi-regulatory agency scheme for AI advice, scheduled to launch next year under the management of the DRCF. This initiative will provide tailored support to businesses in meeting regulatory requirements across various sectors while safely innovating with AI.

Public-private collaboration on best practice

Collaboration on how to apply particular regulatory or ethical principles in practice, for example by working through relevant case studies, can encourage the development and sharing of best practice and enable the resolution of areas of uncertainty. A collaborative approach between industry and authorities can iteratively develop best practice and test what works in a practical context.

Such programmes could follow a similar approach to the Veritas initiative in Singapore, as outlined above, as well as harnessing relevant standards setting bodies and their recommendation development processes.

Irrespective of the overall regulatory model, collaboration of this kind can help inform policy development.

Key regulations and acts referenced in this section**UK**

- [UK AI Whitepaper](#)
- [FS 2/23 — Artificial Intelligence and Machine Learning](#)
- [DP 5/22 — Artificial Intelligence and Machine Learning](#)
- [SS1/21 — Operational Resilience](#)
- [SS2/21 — Outsourcing and Third-party Risk Management](#)
- [AI Safety Summit — Capabilities and Risks from Frontier AI](#)

EU

- [EU AI Act](#)
- [EU AI Liability Directive](#)

US

- [Whitehouse — Blueprint for an AI Bill of Rights](#)
- [Executive Order on Safe, Secure, and Trustworthy Artificial Intelligence](#)

Others

- [G7 Leader's Statement on AI](#)
- [China's Generative AI Measures](#)
- [Summary of China's AI Regulations](#)
- [Monetary Authority of Singapore's Veritas Initiative](#)

8. CONCLUSION AND OUTLOOK

Firms and adoption

This report represents a specific moment in time. We anticipate a forthcoming surge of pilots and proofs of concept in the next three to six months, which will further validate the potential of this technology and uncover valuable insights to inform future steps. Executives of financial services institutions have numerous factors to consider when it comes to the adoption of AI, but their primary focus should be on testing the business outcomes promised by this technology and guiding their teams to validate and deliver on those outcomes safely. Responsible value creation must be the focus.

Moving to action, firms should:

- Invest in educating their workforce on the different types of AI and their respective benefits and weaknesses. This should include how to identify use cases, and how to seek funding to test and adopt.
- Identify the way that their firm best adopts technology and develop this model, whether it is a centralised 'centre of excellence', a federated model, or the use of partners and outsourcing.
- Set their governance to enable adoption at scale. Having a senior executive charged with being the centralising driver of adoption can be a successful component. On the side of risk management, moving quickly to establish a clear 'house view' on AI risk appetite and to implement any changes to risk management frameworks and policies can help the business move quickly to adoption.
- Identify their biggest cost drivers — often in technology modernisation and operations — to help validate whether AI tools can help to automate and ultimately create structural cost advantages. The maturity of enterprise data, the technology platform, and skills and capabilities will be the key constraining factors.

Policy and regulation

Authorities are diligently working across various domains to understand and account for AI risks, both at national level and in international forums.

Within the UK we will see the final AI regulatory framework confirmed. Whether or not changes are made to the original proposals, consultations and guidance are likely from the BoE, FCA, ICO, Equality and Human Rights Commission, CMA and others. The cross-sectoral AI advice service will also launch in 2024. And parliamentary interest will no doubt continue.

At the international level over the next year there will be a small AI Safety Summit in Paris followed by a full summit in Seoul, building on the progress made at Bletchley Park. The G7 Hiroshima Process will also continue, building on its regulatory principles and code of conduct.

Given the speed with which the complexity of the issues involved are evolving — and the speed of innovation — it will be essential for the public and private sectors to collaborate to deepen their collective understanding of AI risk, and to identify and disseminate best practice in risk mitigation.

Utilising the insights from broader industry discourse, authorities should concentrate on clarifying their approach and direction, as well as offering early visibility of the expected endgame of regulation.

UK Finance will be engaging actively with this busy agenda, drawing on its new AI Policy Committee, comprised of industry experts on different aspects of AI and AI policy. This is an area with many interested parties; we will be engaging not only with regulators but also with parliamentarians, policy makers, academics and trade bodies. We will work to develop insightful thinking on the opportunities, the risks and the challenges to inform policy development and ultimately deliver better outcomes for consumers and businesses. We will also contribute to key initiatives, such as the Lord Mayor's Ethical AI Initiative and the FCA's AI Sandbox.

As with preceding eras of AI development spanning the past 50 years, we're eagerly anticipating the future and extend our gratitude to all members for their contributions in shaping the content of this report.

ENDNOTES

1. Trustworthy LLMs: a Survey and Guideline for Evaluating Large Language Models' Alignment (Yang et al, 2023)
2. Artificial Intelligence Act (European Parliament, 2023)
3. Artificial Intelligence Liability Directive (European Parliament, 2023)
4. Blueprint for an AI Bill of Rights (The White House, 2022)
5. SEC Proposes New Regulatory Framework for Use of AI by Broker-Dealers and Investment Advisers (Dechert, 2023)
6. CFPB Issues Additional Guidance on Use of AI in Credit Underwriting (Bradley, 2023)
7. Singapore is not looking to regulate A.I. just yet (CNBC, 2023)
8. Singapore's approach to AI governance (Personal Data Protection Commission, 2020)
9. Launch of AI Verify — An AI Governance Testing Framework and Toolkit (Personal Data Protection Commission, 2022)
10. Veritas Initiative (MAS, 2023)
11. China: Generative AI Measures Finalized (Library of Congress, 2023)
12. Veritas Document 1 (MAS, 2021)
13. CDEI portfolio of AI assurance techniques (CDEI, 2023)

ADDITIONAL REFERENCES

- a. Generative Artificial Intelligence in Finance: Risk Considerations (IMF, 2022)
- b. Reinventing insurance with Generative AI (Oliver Wyman, 2023)
- c. Machine learning in UK financial services (Bank of England, 2022)
- d. Artificial Intelligence Public-Private Forum (Bank of England, 2022)
- e. ChatGPT and Other Large Language Models: Banking Edition (Celent, 2023)
- f. The AI Revolution in Banking (Oliver Wyman, 2022)
- g. Artificial Intelligence Applications in Financial Services (Hermes Investment Management, Marsh, Oliver Wyman, Bryan Cave Leighton Paisner, 2019)
- h. Artificial Intelligence Index Report (Stanford University HAI, 2023)
- i. The State of AI in Banking (Forrester, 2023)
- j. How Generative AI will transform CRM (Forrester, 2023)
- k. AI Foundation Models Initial Report (CMA, 2023)

About UK Finance

UK Finance is the collective voice for the banking and finance industry. Representing more than 300 firms across the industry, it seeks to enhance competitiveness, support customers and facilitate innovation. Our primary role is to help our members ensure that the UK retains its position as a global leader in financial services. To do this, we facilitate industry-wide collaboration, provide data and evidence backed representation with policy makers and regulators, and promote the actions necessary to protect the financial system. UK Finance's operational activity enhances members' own services in situations where collective industry action adds value. Our members include both large and small firms, national and regional, domestic and international, corporate and mutual, retail and wholesale, physical and virtual, banks and non-banks. Further information is available at www.ukfinance.org.uk.

UK FINANCE KEY CONTACTS

Jana Mackintosh

Managing Director, Payments, Innovation and Resilience

Walter McCahon

Principal, Privacy and Data Ethics
walter.mccahon@ukfinance.org.uk

Phillip Mind

Director, Digital Technology and Innovation
phillip.mind@ukfinance.org.uk

Sushant Subedi

Analyst, Digital Tech and Cyber

About Oliver Wyman

Oliver Wyman is a global leader in management consulting. With offices in more than 70 cities across 30 countries, Oliver Wyman combines deep industry knowledge with specialised expertise in strategy, operations, risk management, and organisation transformation. The firm has more than 6,000 professionals around the world who work with clients to optimise their business, improve their operations and risk profile, and accelerate their organisational performance to seize the most attractive opportunities. Oliver Wyman is a business of Marsh McLennan [NYSE: MMC]. For more information, visit www.oliverwyman.com. Follow Oliver Wyman on Twitter @OliverWyman.

OLIVER WYMAN KEY CONTACTS

Lisa Quest

Partner, Head of UK and Ireland,
Co-Head of the Public Sector and Policy Practice Europe
lisa.quest@oliverwyman.com

Daive Taliente

Global Chair, Government and Public Institutions Practice
daive.taliente@oliverwyman.com

Adrian Oest

Partner, Retail & Business Banking
adrian.oest@oliverwyman.com

Sian Townson

Partner, Data and Analytics
sian.townson@oliverwyman.com

Tomas Sanchez

Engagement Manager
tomas.sanchez@oliverwyman.com

Lea Chocron

Senior Consultant
lea.chocron@oliverwyman.com

Brendan Baptista

Consultant
brendan.baptista@oliverwyman.com

Piers Butel

Senior Research Analyst
piers.butel@oliverwyman.com