

# ARTIFICIAL INTELLIGENCE

▽

## IN FINANCIAL SERVICES

June 2019





---

# TABLE OF CONTENTS

---

<b>FOREWORD</b>	<b>4</b>
<b>INTRODUCTION</b>	<b>5</b>
<b>WHAT IS AI?</b>	<b>6</b>
AI in the financial services industry	7
<b>PART ONE</b>	<b>8</b>
Building an AI capable organisation	8
AI maturity	8
Owning AI	10
Ethical AI governance and oversight approach options	13
Considerations	14
<b>PART TWO</b>	<b>15</b>
How to implement and scale AI	15
Is AI the right solution to the problem?	16
Analysing the business use case	16
Customer experience can be transformed	17
New business models	18
Cyber security and financial crime	18
First steps into AI	19
Advancing to high-value use cases	20
How to scale AI and demonstrate value	21
Measuring RoI and business value	22
Conclusion	23

# FOREWORD

Very few technologies have captured the popular imagination like Artificial Intelligence (AI). It has become a key feature in science fiction movies and news stories about technology.

But beyond the hype, organisations are looking at how it can transform their businesses, drive efficiencies and help them to serve customers better.

AI has become a focus for consumers, institutions and regulators globally. In the UK, the Office for AI and the AI Council have been established. In Europe, the European Commission has developed 'Ethics guidelines for trustworthy AI'. As new capabilities and models emerge it is important for institutions to understand how they can take advantage of each new development. However, while the pressure to stay ahead has never been greater, it is critical that this is done in a responsible way. Consumers are increasingly mindful of the security of their data and how it is being used, and institutions are aware that new capabilities can also create new potential liabilities.

It is worth recognising that financial services have been at the forefront of using algorithms and data for some time. There have been long established frameworks and governance in place to manage this within firms, and organisations know they have a responsibility to comply with laws and regulations, including GDPR.

The pace of development and capability is accelerating, and, as in other sectors, this technology is enabling financial institutions to deliver better and more resilient services to their customers. This transformation can drive efficiencies in back office functions such as cyber security, risk modelling and call centres. It can also help deliver more intelligent offerings to customers, from helping them to protect their money against attempts at unauthorised access, to providing new ways for them to save.

As a business enabler, the responsibility for AI cannot simply be left to the technology functions of an organisation. Creativity, compassion and broad ethical judgement are critical attributes to harnessing responsible and useful applications for AI and broad governance is therefore needed. Firms will need to think horizontally across business functions, from product development, risk and audit, to finance and technology. They should also consider their governance models and assess how well this maps to a new way of working and look at how they manage their skills and capabilities both currently and into the future.

In an evolving field, we don't claim to have the final answer. This paper aims to introduce readers to a number of considerations that they, and their institutions, may need to ask themselves as they look at how to benefit from AI. How each institution identifies the use cases for AI, and seeks to implement the data strategies, cultural and governance arrangements to take advantage of it will be firm specific. It is clear that AI will be a key feature in how firms operate and deliver services. We hope this paper provides a useful set of practical tools to begin to understand how firms can begin to take advantage of what AI offers.



**Stephen Jones**  
CEO  
UK Finance



**Craig Wellman**  
Director, Financial Services  
Microsoft

# INTRODUCTION

Artificial Intelligence (AI) is the software at the centre of the Fourth Industrial Revolution. Today AI is already a part of our daily lives, as we engage with these systems through various applications including search, recommenders and even customer support.

The World Economic Forum cites that “when compared with previous industrial revolutions, the Fourth is evolving at an exponential rather than a linear pace. Moreover, it is disrupting almost every industry in every country. And the breadth and depth of these changes herald the transformation of entire systems of production, management, and governance<sup>1</sup>”. Financial services is very much an industry at the heart of this revolution.

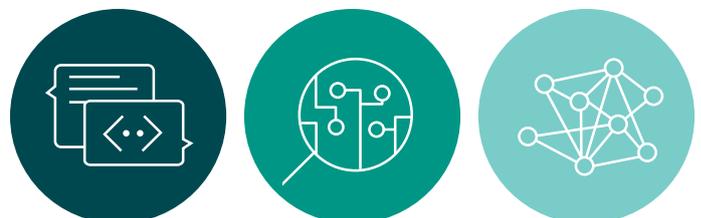
We start this paper by exploring what AI is. While terms such as machine learning, deep learning, neural networks and random forests technologies, all components of AI, are likely to be familiar across organisations, we are focused on the practical exploration of what it means for your organisation. These technologies could augment an organisation’s capabilities in compelling new ways, developing solutions that see, hear, speak, and understand customer needs.

Financial services organisations have a long history of working with data, pioneering the use of analytics to derive insights for creating better, more profitable business models and meeting customer needs. But as we look ahead, the most impactful data-driven solutions will go well beyond analytics and will include AI systems. The paper will bring this alive through examples of AI in market today and share how a firm can assess AI maturity across the broad categories of Cultural, Ethical and Ownership

Firms can often fall into two pitfalls with AI. The first one is to look at AI simply as a tool, and not consider the wider cultural and organisational changes necessary to become a mature AI business. Part One of the paper will look at how, as organisations start to embed AI into core systems, they need to consider the implications of AI that go beyond the technical, including the wider impact on culture, behaviour and governance. In this paper we will also provide a framework for responsible AI and detail what this could mean in practice, given AI is not just about the technology, but the business, social and economic implications.

The second pitfall is to struggle to know how to implement AI, or to scale AI from innovation centres to enterprise-wide solutions. Part Two of this paper will help organisations identify where AI is the right solution, and how to identify the high-value use cases, looking more deeply at analysing the business case. AI maturity in-market means that capabilities need to be considered as a core factor in the business case. Ensuring the fundamentals of an AI environment will directly impact the business case and accelerate your time to market. We conclude the paper by exploring how to scale your capabilities to ensure you drive high value and low-risk outcomes, by looking at how you measure ROI and business impact.

AI has the potential to revolutionise our society, transform every aspect of business and improve economic output. PWC estimates that AI could boost UK GDP by up to ten per cent by 2030<sup>2</sup>. While AI offers huge potential if organisations can mitigate the two pitfalls mentioned above, organisations will need to remember that the implementation of AI must exist in a wider societal context that also considers safety, reliability and fairness, privacy and employment and skills.



1. <https://www.weforum.org/agenda/2016/01/the-fourth-industrial-revolution-what-it-means-and-how-to-respond/>  
2. <https://www.pwc.co.uk/economic-services/assets/ai-uk-report-v2.pdf>

---

# WHAT IS AI ?

Today, one of the most important emerging technologies is artificial intelligence (AI). AI is a set of technologies that enable computers to perceive, learn, reason and assist in decision-making to solve problems in ways that mimic human thinking.

AI systems are already being used in daily life to answer questions, translate languages, optimise power consumption, operate factories, write news stories, colour grade photos, drive vehicles, and diagnose medical issues. AI technologies, much like automation that has come before, allows us to scale our capabilities through machines in an increasingly data-rich world.

By itself, automation is not new. Over the last forty years, the automation of much of the traditional production line has helped productivity in manufacturing increase by 230 per cent<sup>3</sup>. In the past however, automation has largely been confined to highly controlled environments such as a factory floor, or to routine administrative tasks that are easy to describe as an algorithm. What makes the current wave of AI different is its increasing ability to detect patterns from unstructured data and recognise far more nuanced patterns, allowing it to operate in a far wider range of environments.

A surge in accessibility to computing power, due to the rise of cloud computing, in conjunction with the increasing amounts of data being created, has lowered the barriers to entry for both building and using AI technologies. This in turn has allowed organisations large and small to experiment with AI systems.

---

“You can see AI as a tool or as a threat. We use AI to enable our traders to do their jobs better as opposed to looking at things that will replace traders. We focus on how we enhance the employee experience and make it as efficient as possible.” – Head of Strategy, Top Tier European Bank

---

3. Author calculation from ONS data

## AUGMENTED INTELLIGENCE

“AI systems work well when you have lots of data, the task is well-defined and repetitive, and the system is in a closed environment without a lot of unknowns. However, when you move beyond this and either try to adapt a system to a new domain or try to use it in the open world, these algorithms face a lot of challenges.

“People don’t do very much statistical reasoning, we may not be very consistent all the time, and we have

biases. However, we have common sense reasoning, counterfactual reasoning, creativity, and we’re good at adapting to new settings and combining our knowledge from different domains.

“The combination of humans and machines will be more effective, efficient, and reliable than either is on their own.” - Ece Kamar, Senior Researcher at the [Adaptive Systems and Interaction Group](#), Microsoft

## AI in the financial services industry

Within financial services there have been many innovations that have changed traditional banking over time, reimagining the way the industry operates, as well as the nature of jobs.

The financial services industry has a history of using quantitative methods and algorithms to support decision making. These are a foundation of AI systems, and the industry is therefore primed for AI adoption, positioning it at the forefront of adopting and benefiting from AI technologies.

AI can build on human intelligence by recognising patterns and anomalies in large amounts of data, which is key in applications such as anomaly detection (e.g. fraudulent transactions). AI can also scale and automate repetitive tasks in a more predictable way – including complex calculations, for example for determining risk.

One of the most important types of AI is machine learning (ML), algorithmic systems that can recognise patterns and learn without being explicitly programmed. ML has become one of the key AI technologies used by the financial services industry, due to its ability to leverage existing algorithms to improve ever growing amounts of data, thereby creating new capabilities.

AI will become more and more crucial as the data that we create continues to grow. This will lead to a situation where processes that previously did not require AI (e.g. fraud detection), will no longer be able to succeed without AI.

# PART ONE

Avoiding the first pitfall - AI is not just in the domain of technology or just the business

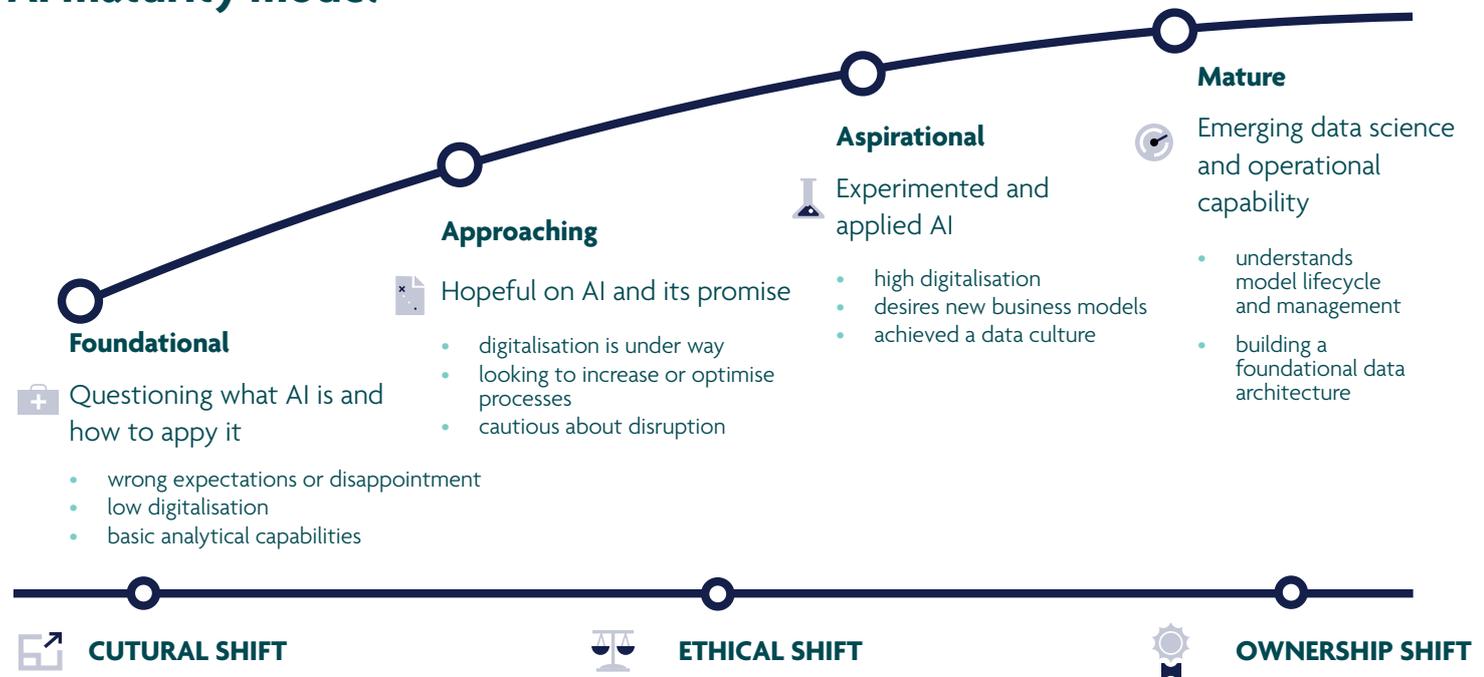
## Building an AI capable organisation

Many organisations across financial services are on a journey to building their AI capabilities. One way to assess your organisation's AI maturity is to consider it as a four-step process, as seen below. Organisations should consider their maturity as a way of avoiding the first pitfall – that of looking at AI simply as a technology tool – and use a model to plan how they move through the stages.

## AI maturity

Underpinning each stage are core concepts associated with areas of focus that contribute directly to AI maturity across the broad change categories of Culture, Ethics and Ownership. Organisations should consider where they sit on this journey, and how they can implement change to progress further.

## AI maturity model



---

**Foundational** – a foundational organisation is seeking to understand the possibilities of AI across broad categories of business scenarios and how these might be applied to bring business benefit. Likewise, foundational organisations are looking to understand how industry peers are using AI to learn how to apply it.

At this level organisations are looking to define a process to help make data-driven decisions. Foundational organisations need to look for experimentation on fast, iterative projects to bring value and build the firm's experience and confidence in AI and begin the cultural shift to that of a data-driven organisation.

**Approaching** – approaching organisations continue to expand their cultural change to data-driven decisions. There is a focus on adopting a data culture and looking to infuse AI into strategic initiatives to drive efficiency or create new business approaches and models. Approaching organisations are ready to learn about owning custom AI solutions through rapid experimentation and quick iterative sprints. They are also beginning to consider how to implement, govern, monitor and improve AI over time and are starting to define their ethical AI standpoint.

**Aspirational** – at all levels aspirational organisations understand that AI is necessary for them to adapt and transform to compete within their industry. They are using AI to improve processes and strive to use data to inform decisions. These organisations are often aware that others are using AI and they are cautious about the disruption AI may bring from competitors inside and outside of their current peer group.

A cultural shift to employee empowerment is often a focus at this level, allowing increased levels of collaboration and idea generation to optimise existing processes and suggest new business models.

**Mature** – mature organisations have embraced a cultural change to include lifelong learning and a growth mindset and are able to naturally run multiple strategic initiatives through rapid, iterative experiments to help translate data-driven insight into action.

Mature organisations look to create digital experiences that are impactful over time and strive to embed their ethical perspectives into all they do – often asking “We know we can do these things with AI, but should we?”

## Owning AI

There are many aspects to consider in 'owning' the AI that supports a mature financial services organisation.

This paper focuses on five key aspects:

1. A culture of fairness and understanding of bias in data models
2. A culture that supports consumer trust and understands reputational impact
3. A core set of AI capabilities and skills
4. Ethical AI governance and oversight
5. International context

### 1. A culture of fairness and understanding of bias in data models

AI systems will inevitably make or support a process where decisions are made about risk levels, product offers or fraudulent activity that will directly or indirectly affect customer experience, options or financial situation. When this happens, we need to ensure AI systems make recommendations without unnecessary bias and do not discriminate on race, gender, religion or other similar factors.

Bias can be introduced at any stage of a system's lifecycle, from inception, through data collection, to modelling, operation and ongoing maintenance and update. To that end, to strive for fairness firms must understand how bias can affect AI systems and ensure they have the tools and principles in place to help them clearly understand the source of bias in order to mitigate this.

AI, and Machine Learning (ML) specifically, use algorithms and statistical techniques to look at patterns and inference in large data sets of 'training data', to create models that then can be used to predict new outcomes. An unintentionally biased training data set may create a model that responds well to the majority of use cases, but which fails to predict adequately edge cases where training data sets are light or unintentionally skewed.

## EXPLAINABILITY OF DECISIONS

As more complex use cases are built to leverage AI, thought and design is needed in the explainability of the decisions or outputs produced by these tools. Full explanations may be needed to reassure customers as to why their credit or mortgage application was not accepted by an AI/ML model. Where firms identify a trade-off between the level of explainability and accuracy, firms will need to consider customer outcomes carefully. Explainability of AI/ML is vital for customer reassurance and increasingly it is required by regulators.

A mature AI organisation, that recognises the need to ensure that AI systems are fair, will insist on processes to identify bias in datasets and ML algorithms. Identification and mitigation of bias in datasets starts with an evaluation of the source of data, understanding how it is organised and testing it to ensure it is representative and achieves fair outcomes for customers.

## WHAT IS BIAS IN AI?

Bias is a feature of a statistical technique or of its results whereby the expected value of the results differs from the true underlying quantitative parameter being estimated. A statistic is biased if it is calculated in such a way that it is systematically different from the population parameter being estimated.

Though statistical mechanisms underpin AI systems, this is not necessarily what we refer to when we talk about bias in the context of AI. Since AI systems learn what they know from training data, when these

datasets inaccurately reflect society, mirror unfair aspects of society or existing institutional prejudices, those biases in the data can be replicated in the resulting AI systems on an even wider scale. At the same time, the lack of diversity or inter-disciplinary practice in the current AI workforce can also make it hard to pick up on problems where they arise.

Source for statistical bias - [https://www.revolvly.com/topic/Bias%20\(statistics\)&item\\_type=topic](https://www.revolvly.com/topic/Bias%20(statistics)&item_type=topic)

An organisation can also identify and mitigate bias in ML algorithms by leveraging tools and techniques that improve the transparency and intelligibility of models. Training employees to understand the meaning and implications of AI, raising awareness for potential bias in results, and ensuring individuality for AI enabled decisions, are all important steps in building fair systems.

Similarly, there should be systems in place to help determine early in the development lifecycle when there isn't sufficient data quality or completeness to be able to continue and build a robust and fair AI system.

What one individual perceives to be fair is not always something others will agree with. Firms must strive to be fair in the creation of their AI models, but equally transparency around how AI models make decisions is critical to allow others to judge and challenge definitions of fairness.

## TRANSPARENCY

As AI increasingly impacts people's lives, firms must provide contextual information about how AI systems operate so that people understand how decisions are made and can more easily identify potential bias, errors and unintended outcomes.

### 2. A culture that supports consumer trust and understands reputational impact

“We have a set of risks around data privacy and security that are forefront. We need to be more aware of consequences, and appropriate steps should be taken.”  
– Head of Business Design, Top Tier UK Bank

AI will become an instrumental way to provide the personalisation that customers increasingly expect. This has implications for trust in AI, which must be earned over time, and importantly must be maintained. The public has entrusted, and continues to entrust, banks with their personal data – with 80 per cent of recent respondents willing to share personal data with banks in return for personalised and convenient services<sup>4</sup>.

## PRIVACY AND SECURITY

Like other cloud technologies, AI systems must comply with privacy laws that regulate data collection, use and storage, and ensure that personal information is used in accordance with privacy standards and protected from theft.

Other studies show that banks continue to be seen as long-term trustees of sensitive customer data – individual wealth, transactional data and often identity – and that customers trust banks to safeguard their data more than organisations in any other sector<sup>5</sup>.

The AI models driving these experiences may well become part of the regular customer touchpoints and be highly visible and frequently used. Alongside an unerring focus on transparency and fairness, firms must also recognise the necessity of privacy and security in establishing ethical AI systems and through that ensure that necessary experimentation, agile, and test and learn approaches do not sacrifice these core operating principles.

Without appropriate testing, governance and control, a rapid growth in AI models could have significant reputational impact and subsequent reduction in consumer trust for individual banks or the industry as a whole.

For continued trust, a cultural AI maturity that ensures the consistent application of transparency, privacy and security principles to support the new challenges AI brings is essential for the long-term benefit this technology has to offer.

4. [https://www.accenture.com/\\_acnmedia/PDF-95/Accenture-2019-Global-Financial-Services-Consumer-Study.pdf](https://www.accenture.com/_acnmedia/PDF-95/Accenture-2019-Global-Financial-Services-Consumer-Study.pdf)

5. <https://www.ncipher.com/about-us/newsroom/news-releases/ncipher-survey-reveals-americans-trust-banks-most-their-personal>

---

“AI still needs to build trust and it needs to earn the trust. It needs to gain trust from the industry. We need to ensure that information shared is reliable. As an example, you go to a specific bank because you trust that bank. With AI, it needs to be the same.” – Vice President of Technology, US Bank

---

### 3. A core set of AI capabilities and skills

Most organisations are looking to digital technologies to help them transform, with AI a key component of that strategy. However, many do not have the data scientists, developers or AI experts they need to build their own solutions.

A growing digital skills gap is not a new challenge for IT. Companies will need to continue to invest in their digital talent pipeline, as well as building their own in-house skills, by giving employees access to training to reskill for the AI era.

Equally, AI is a technology that requires more than just new IT skills. Business leaders need to be able to take full advantage of new capabilities that AI brings. There is a need to understand AI strategy, culture and responsibility, if staff of all levels are to understand the AI art of the possible to lead transformation, optimise operations, deliver enhanced customer value and reduce costs.

### 4. Ethical AI governance and oversight

It's comparatively easy and often compelling to build AI prototypes, bots, ML driven anomaly detection or Robotic Process Automation proof of concepts in innovation labs. As these technologies are built and used in an internal experimental environment, there is more room to explore different solutions at a fast pace. A transition from this experimental state into live production, where these tools will interact with people, requires governance, controls and process.

This transition critically requires parts of the organisation aside from just the technical to engage. This is to ensure that firms ask the right questions to understand and approve AI solutions, as well as create the appropriate levels of controls. These processes should be created and implemented in parallel to provide oversight and guidance, tailored to your organisation's regulatory requirements, culture, guiding principles, and level of engagement with AI.

## PROFESSIONAL PROGRAM IN AI

There are numerous resources available to help individuals grow their knowledge and skills in AI - some of the free to access content is a great place to start to disseminate this knowledge. Last year Microsoft announced its Professional Program in AI, an AI learning track open to the public. The program provides job-ready skills and real-world experience to engineers and others who are looking to improve their skills in AI and data science through a series of online courses that feature hands-on labs and expert instructors. This aims to take aspiring AI engineers from a basic introduction of AI, to mastery of the skills needed to build deep-learning models for AI solutions that exhibit human-like behaviour and intelligence.

The program is part of a broader set of AI training that also includes the enterprise developer-focused AI School, which provides online videos and other assets to help developers build AI skills. That program includes both general educational tools for developers looking to expand AI capabilities and specific guidance on how developers can use different tools and services. In addition the AI Business School, developed to engage business leaders, is aimed at a non-technical audience with the aim of getting executives ready to lead their organisations on a journey of AI transformation.

## ETHICAL AI GOVERNANCE AND OVERSIGHT APPROACH OPTIONS

There are any number of approaches to building a governance system for AI.

Some organisations choose to bring all of their company-wide ethics initiatives and decision making, including those for AI, under a single role of Chief Ethics Officer. This form of governance enables companies to quickly develop policies around ethics while ensuring there is accountability for each decision. However, while the chief ethics officer role can very effectively create policies, without a supporting team they may struggle to implement these.

This can be supplemented with an ethics office by forming a dedicated ethics team from across the company. While ethics offices can be responsible for developing new policies, they are also typically sized to be able to implement policies at scale. With dedicated team members working at all levels of the company, an ethics office can effectively ensure policies and processes are followed in a meaningful way by all

employees and in every stage of AI engagement. Ethics offices also prove adept at building a culture of integrity within an organisation.

Another approach is to bring together a diverse array of experts and senior leaders from within the company to address AI ethics. Ethics committees generally do not have members dedicated solely to ethics but work with different groups within the organisation. This form of governance ensures a diversity of perspectives, often includes opinions from external experts, and commitments from senior leaders. Microsoft has formed the AI and Ethics in Engineering and Research (AETHER) Committee which includes senior leaders from across Microsoft's engineering, research, consulting and legal organisations focusing on the proactive formulation of internal policies, how to respond to issues, consider best practices and provide guiding principles that are used in the development and deployment of AI products and solutions.

AI governance will be needed to serve several purposes, and firms need to consider how they:

- develop and implement policies, standards and best practices
- take responsibility for fostering responsible, ethical and trustworthy AI
- assess and measure individual project/programme risk and overarching enterprise risk
- determine accountability and responsibility for the ongoing AI maintenance lifecycle

### 5. International context

As the use of AI becomes more pervasive and more powerful, it becomes a shared responsibility across banking, financial services and other sectors to engage with AI responsibly for the common good. Each organisation will have its own take on principles, governance, controls, ethics and guidelines to accomplish this, but an understanding of others' perspectives will become critical to ensuring consistency of approach in a multi-party offering.

Enterprises can also take advantage of and contribute to shared initiatives such as Partnership on AI. This is a group of researchers, not-for-profit firms, non-governmental organisations, and companies dedicated to ensuring that AI is developed and utilised in a responsible manner. They work to advance the public's understanding of AI, serve as an open platform for discussion and engagement about AI and its influences, and develop best practices for AI technologies. By sharing this knowledge, organisations such as Partnership on AI strive to help ensure that AI technologies are used in beneficial and responsible ways.

It's likely that changes to existing laws and regulation will expand to engender trust in AI and minimise risks from its use – such as the recently released European Commission Ethics Guidelines for Trustworthy AI.

It will be important for all banks to be aware of these potential changes and contribute to the dialogue on AI advancements and potential legislation to strike the right balances in the rapid progression of digital transformation in banking.

---

## Considerations

As AI technologies persist and support ever more business processes, as well as an assessment of each individual AI project it is critical to consider the wider implications of an organisation's reliance on AI - to assess the shift of responsibility of decision making from humans to machines across business process and to confirm that this is not introducing significant 'rolled-up' business risk. Organisations must consider how to supplement existing governance frameworks, or create new ones, to ensure that the ethics, appropriateness and risk of AI is in balance with the benefits it promises and an organisation's corporate standpoint.

Equally, creating an AI-based system is not the end of a story. AI systems deal in probabilities and most require continuous training, monitoring and evaluation to ensure they remain true to their original intent, else deviation from design goals may result in decay and unintended bias. The ongoing maintenance of many AI systems is required for the continued benefit and ongoing reduction in potential harm. Your organisation needs to be aware and structured to support these tasks.

In order to do that, organisations should assess their operational 'AI Maturity' – looking at attributes of Strategy, Culture and Organisation and Capabilities to identify a current position on a maturity curve and determine how to increase maturity in order to embrace more advanced AI capabilities.

# PART TWO

## Avoiding the second pitfall

### How to implement and scale AI

The second pitfall firms can fall into is to struggle to know how to implement AI or to scale AI from innovation centres to enterprise-wide solutions. To be successful AI requires ongoing collaboration. One of the key challenges that organisations face is adopting AI at scale, currently many AI projects are in the realm of innovation teams and don't always have a clear path to production – a recent MIT Sloan Management review found that 82 per cent of organisations had failed to adopt AI beyond proof of concept or pilot<sup>6</sup>.

This creates the need to bring AI into the business and demystify what it is, and more importantly, what it is not. There is a need to approach AI with an ongoing dialogue between the business and technical functions.

It is important to articulate the ROI that AI could contribute to get the organisation aligned. The drive for AI adoption needs to start at the top of the organisation and filter down to all levels. By focusing on small incremental wins with a clear ROI, while building an AI driven culture, organisations can maximise the opportunity that AI brings. To support firms, this paper introduces AI Ops as an approach technology teams could adopt to manage development. There are also tools your organisation could introduce in technology teams to support data engineers. Firms should speak to their technology teams to understand how these might be applicable for them. In this process it is important to keep the business and technology stakeholders involved at all stages to ensure that iterations are considering evolving needs.

### WHAT IS AI OPS?

In many large organisations, DevOps – development and operations – is a part of IT responsible for building better apps, at a faster pace. This is a key enabler of digital transformation – as many organisations increasingly face external pressures to be more agile. With DevOps, they can improve the speed at which they build and deliver digital solutions.

By rethinking the role of software procedures in an AI world, DevOps for AI is now becoming a key part of operationalising AI at scale, as it allows organisations to manage their AI pipelines in a way that ensures enterprise standards and is a crucial step in moving from experimentation to industrialisation. With this in place, organisations can track and audit which data and libraries were used, how models were trained and how products were defined. In regulated industries, especially where documentation is a necessity, AI Ops is a key step to creating more transparency and accountability in AI systems.

6. <https://s3.amazonaws.com/marketing.mitsmr.com/offers/AI2018/60280-MITSMR-BCG-Report-2018.pdf>

## Is AI the right solution to the problem?

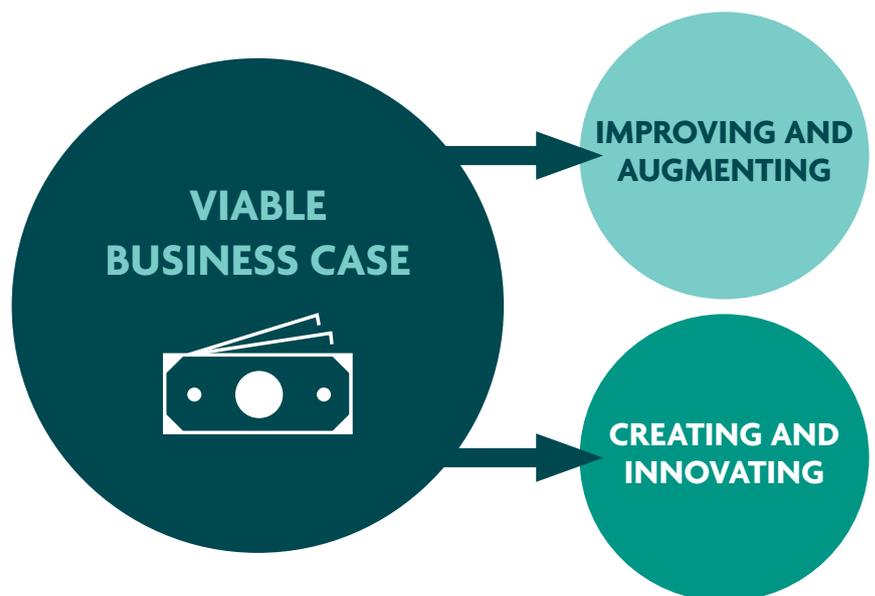
AI will not always be the answer; sometimes the ROI from these systems compared to the traditional ways of doing things does not justify the change. For example, in a scenario where machine learning (ML) algorithms are used to close down alerts for suspicious transactions in financial crime, there is an element of risk that a small number of genuine financial crime transactions could still be approved. In this case, the organisation would need to evaluate if ML is the appropriate tool to use, given the level of risk it is willing to accept. The alternative in this case is to risk fraudulent transactions going through, or genuine transactions being held, potentially resulting in poor customer satisfaction.

There is not a one-size fits all for this assessment and perspectives will differ even within an organisation, but the first key step is to define the problem statements, and then understand for each if AI is the appropriate solution.

## Analysing the business use case

One of the key challenges for organisations is deciding which problems to solve with AI. The technologies that encompass AI include forms of ML, neural networks, natural language processing, speech recognition and so forth. These are often used in conjunction when it comes to creating AI systems. For example, fraud detection might use ML in conjunction with image recognition to scan documents for fraudulent indicators.

The first step is to determine what the priorities are in terms of problems that can be solved with AI. This is centred around where the main value stream lies and articulated in a business case. This initial prioritisation should not come from the technical or business teams alone, but rather be done in collaboration. Some problems that the business could face might not be technically feasible for an organisation in the short to medium term, due to challenges accessing suitable data or available infrastructure. Therefore, technical input is key alongside strategic alignment.



AI can have an impact across all aspects of a financial services business. Some key industry use cases and applications are highlighted below:

## CASE STUDY - MITIGATING RISK THROUGH PREDICTIVE PROJECT OUTCOMES

GainX is a company that uses Artificial Intelligence and Advanced Organisational Design to deliver insights for their customers. Recently, GainX worked with a UK financial services institution (FSI), helping them with their reporting processes for predictive project outcomes across £3 billion of strategic spend.

This FSI is one of the largest banks in the UK and executives were looking to harness the power of their data. They engaged GainX to quickly aggregate

their disparate data and replace historical measures of change with predictive risk management, relevant KPIs, and strategic decision support.

Through this exercise, the customer gathered an understanding of the current risk, where and why projects were likely to fail, where legacy processes were creating losses, and why key talent was highly isolated. As a result, this FSI was able to recover nearly £2 billion in spend at risk.

## Customer experience can be transformed:

- End to end digital onboarding – with better documentation management, firms can rapidly onboard new customers
- Frictionless and hyper-personalised – by delivering great insight, firms can start to personalise and tailor the experience for a customer
- Real-time advice – with great digital assistants, firms can service and advise customers in real time without the need for them to navigate call center dialogues or wait in queues
- Proactive personal financial management – with AI and ML on historical data, firms can begin to provide customers with proactive recommendations on how to budget for monthly spend and large purchases such as holidays.

---

“With data, we are customer centric, not sales driven. It’s about being useful – we can take a short-term cut of something that may cost us profit in the short term, in exchange for the long-term success with the customer. We use basic machine learning, who do you send what message to and at what time and optimise messages with algorithms.” – Chief Commercial Operations Officer, Top Tier UK Bank

---

## New business models:

- As firms start to build marketplaces, firms will need AI to bring the right services/financial products to their customers at the right time, and on the right channels
- Examples of services include loyalty programs, personal shopping, life event planning, insurance, services provided by local small businesses
- Risk management and regulatory compliance
- Firms everywhere are starting to use Robotic Process Automation (RPA) to augment humans and build compliance and auditing into the processes from the ground up, saving huge amounts of manual compliance work
- RPA can also bring together processes to perform risk reporting and stress testing

- Many firms are employing AI to assess risk for credit scoring
- Many startups are providing new sources of data that can be used to assess risk and, in insurance, improve underwriting

## Cyber security and financial crime

- Payments and transaction data can be challenging to scan for fraudulent patterns of transactions. ML can help to identify fraudulent patterns of data to significantly reduce the risk of false positives.
- With new behavioural biometrics solutions, firms can continuously authenticate users and protect customers even if their credentials are compromised.

### CASE STUDY – DISRUPTIVE IMPACT ACROSS RISK SCORING IN BANKING

SparkBeyond empowered a leading global bank to improve the credit risk scoring of the bank's loan products to small and medium enterprises (SMEs). The ability to compliment a client's data with external data sets and test millions of questions on data in seconds or minutes instead of weeks and months enabled the discovery of hidden insights and nuanced patterns behind high/low risk loans. These were explainable (the insights could be explained to human beings; a 'glass box' platform is important for regulators), actionable, and were operationalised into decision support systems across segments of product.

In addition to credit risk scoring, the bank also leveraged the platform for three additional use cases spanning its value chain - including branch location optimisation and fraud detection. Through leveraging SparkBeyond, the bank was able to increase projected net income of 2.3 per cent in the territories in which it was deployed in. First rolled out in the US and Poland, the bank is now extending this globally.

### CASE STUDY – BEHAVIOURAL BIOMETRICS FOR FRAUD DETECTION

BioCatch delivers behavioural biometrics analysing human-device interactions to protect users and data. Banks and other enterprises use BioCatch to significantly reduce online fraud and protect against a variety of cyber threats, without compromising the user experience.

As an example: A new digital (challenger) bank ran a marketing campaign to attract customers and quickly became the target of a massive online account opening fraud campaign. Fraudsters moved money from a compromised account into a new account that they fully controlled, allowing them to quickly withdraw funds without the legitimate user noticing and creating significant liabilities for the new digital bank. For every 100 good applications, there were 900 bad applications, each generating thousands of dollars in losses when fraudsters withdrew their funds and disappeared.

BioCatch, analysed user behaviors related to application familiarity, computer proficiency and data familiarity. The system generated real-time alerts that stopped the fraudsters and allowed the bank to resume normal operations.

## First steps into AI

When starting to outline the initial stages in AI, it is crucial to consider in parallel how these systems will be put into production and maintained. This should be a thread that develops concurrently with AI development. AI being deployed in proof of concepts, and only in proof of concept scenarios, is not a sustainable strategy as these experiments require a road to production in order to reap the ROI and not create a “sunk cost” situation.

Organisations looking to productionise AI products need to create the right frameworks and processes that can support the transition from the experimental phase to enterprise grade.

The risk is that the value of AI will not be realised and the technology will fail to be adopted. MIT Sloan Management Review’s 2018 research report found that most organisations – 82 per cent of those surveyed – had “failed to adopt AI beyond pilot or proof-of-concept projects”<sup>7</sup>. This could be a high sunk cost to the organisation and create AI fatigue, which can in turn hinder an organisation’s appetite for true innovation.

After understanding the value streams and ensuring that there is a plan around path to production, the next step is to determine the risk of change. For example, if the main value stream is threatened by an increasingly competitive market, the need to change might outweigh the inherent risk of introducing new systems. Ultimately, to ensure that AI is implemented successfully both value and risk as a result of change will need to be evaluated.



## AI UNLOCKS OPPORTUNITIES – ESG RATINGS

The era of “ESG-aware” investing, whereby Environmental, Social and Governance (ESG) considerations are a part of investment and risk management decisions, has evolved from niche to normal.

The thinking is that, if ESG factors are ingrained in investors’ thinking, then a triple-win will follow – more valuable companies, more value for end-investors and, ultimately, a more sustainable system. AI is helping the next generation of companies reduce their environmental and social impact by improving efficiency and developing new products. AI applications for sustainability are at an early stage, but the data suggests they can bring significant benefits in the medium term.

ESG Rating calculation comes with its own challenges of data quality and hidden parameters that impact the overall rating or scoring. Many companies are now implementing machine learning, natural language understanding and neural networks to quickly analyse the data that is gathered under various indicators, combined with the adoption of a complementary technology infrastructure.

The ESG Rating Score of any company can be ranked against other companies in the same industry group and can be updated annually.

The low-value, low-risk quadrant is often the first step into AI for organisations. Applications that sit within the low risk space usually contain requirements that tap into pre-existing mechanisms which the organisation can leverage for AI. For example, a quant team that is working on predictions could incorporate AI into their toolkit to improve prediction times. Deploying ML within chatbots to serve customers who prefer to interact with a self-service portal can improve the service, reduce human error and benefit from efficiency savings.

7. <https://s3.amazonaws.com/marketing.mitsmr.com/offers/AI2018/60280-MITSMR-BCG-Report-2018.pdf>

Taking advantage of democratised AI services can be leveraged to reduce administration tasks and further improve operational efficiencies, such as:

- Computer vision, which allows computers to identify, process and understand images such as written documents, forms and paper applications.
- Using AI knowledge bases to provide quick access to frequently asked questions without relying on a single point of contact.
- Improve screening of on boarding customers as part of Know Your Customer (KYC) checks by leveraging access to Big Data and Machine Learning to predict trends which may cause an issue for compliance.

## Advancing to high-value use cases

The high-value quadrants are where organisations strive to focus their resources, as these areas promise the most opportunity for building a competitive strategy based on AI. This matrix is dynamic, as AI matures in organisations, infrastructure is laid, and mechanisms are developed to support this technology on a wider scale.

This growth in foundations will allow a business to progress to the most beneficial applications. For example, by laying down the foundations for a robust modern data platform, organisations can start unlocking the value of the data within them to create new AI systems.

A key part of growing in maturity is creating an ongoing dialogue between the tech and the business, to ensure that the products being designed are fit for purpose.

## TOOLS TO ENABLE YOUR DATA ENGINEERS

Platforming which supports the full end to end AI/ML process is vital for success. Consideration is needed firstly to achieve a full understanding of the business problem AI/ML will address. A platform which supports access to correct, clean, clearly understood and catalogued data is of fundamental importance.

Platforms which allow data engineers to use the tools and languages they are familiar with is important. Modern Big Data/ML tools (such as Azure DataBricks) support the languages data engineers and data scientists are familiar with meaning that both teams can collaborate and work in a shared environment.

Data scientists are reliant on a data pipeline in the shape required to perform data exploration using domain knowledge and statistical analysis to support the machine learning modeling process across feature engineering,

Model Training and Model Evaluation. A modern 'best of breed' (ie: top of the range in its category) AI/ML platform supports ML Model Management across version control and deployment to be consumed within web applications, chat bots and IoT devices.

## How to scale AI and demonstrate value

Deploying advanced AI solutions successfully in an organisation requires a shift from a deterministic view of execution to a probabilistic approach. However, this doesn't mean that the discipline required to measure the outcomes changes. AI solutions should inherently be considered as experiments with a focus towards solving a certain problem (classification, recommendation etc.), with a concurrent plan of how to scale if successful.

To facilitate this, firms should consider how they plan to run execution cycles with production bound experiments. Firms should specifically consider:

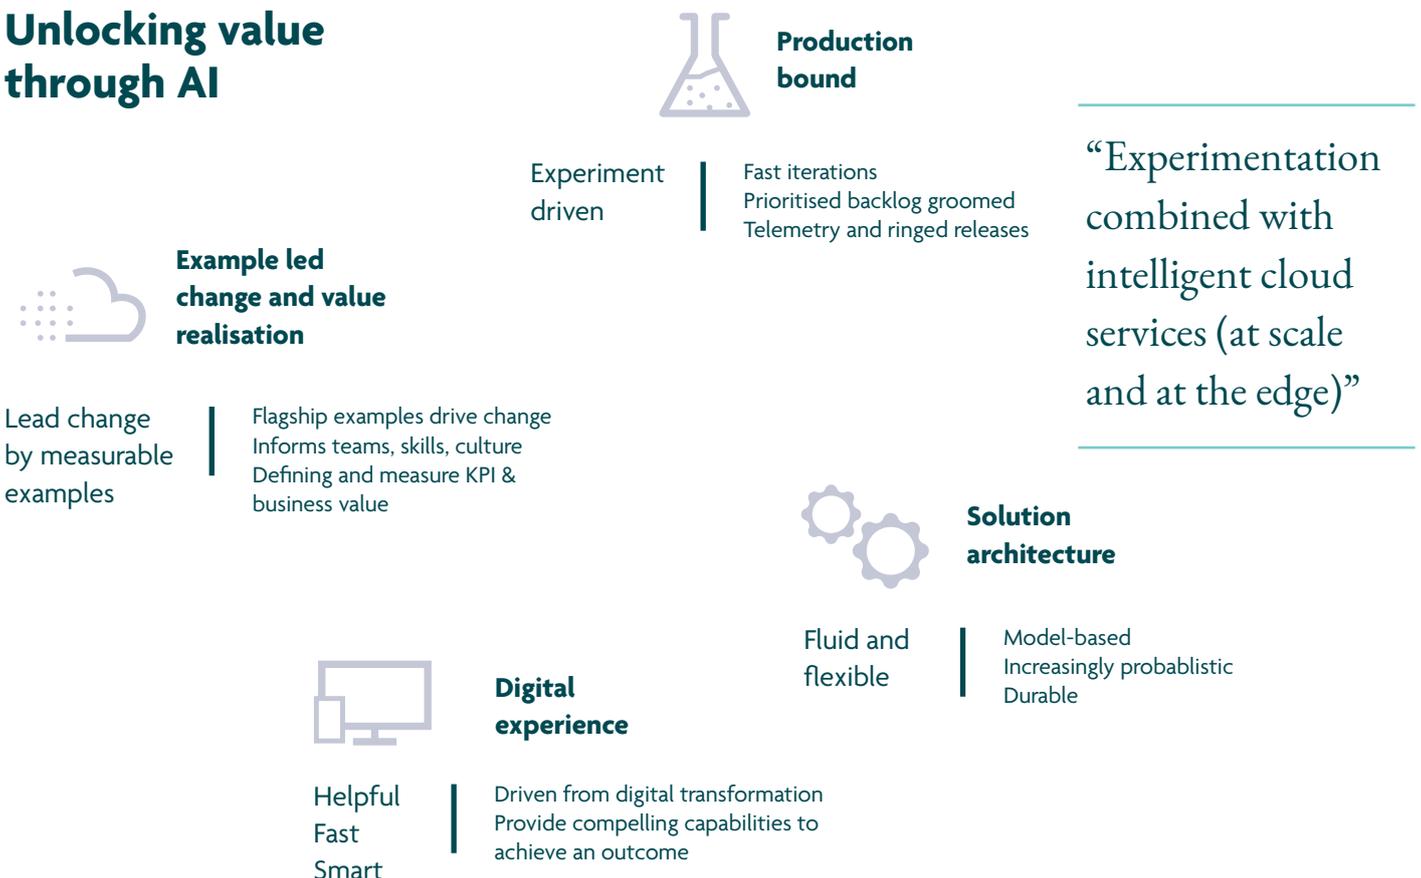
**1. Solution architecture** – Leverage existing model and template-based solutions which have been tested and are durable to adapt to an organisation's needs/ requirements.

**2. Digital experience** – Consider digital transformation as a vision for the 'To-Be' state. In some scenarios, making the current processes more efficient is not the answer. Consider doing different things (transformational) rather than doing the same things differently (incremental).

**3. Example led change and value realisation** – Consider flagship examples to build momentum for change. This is critical to defining the measures of success through KPI improvement and aligning these experiments with business value.

**4. Production bound** – One of the biggest risks of deploying AI at scale is not considering its path to production. This often stems from the fact that experiments are not designed with the future commercial roll out of the technology in mind. Fast iterations with a prioritised backlog of features and telemetry to monitor the performance enables robust progress in order to achieve the business outcomes.

## Unlocking value through AI



---

## Measuring RoI and business value – what gets measured, gets done

---

“Even with relatively small data sets, we’re seeing a 10-20 per cent uplift in the predictive power of some credit risk models when we apply advanced analytics techniques. This gets slightly reduced in production as we have to overlay the models with some additional constraints. Once we actually operationalise this we’ll have 1. Better fraud detection and 2. Better ability to extend credit to people where traditional models would say no. Scaling this out across our portfolio will take some time but it will mean better lending decisions, better customer service and therefore better products.” - Cloud and Big Data Lead, Top Tier Bank

---

Prioritising experiments based on measurable business impact is important to the long-term focus on AI in an organisation. Not all experiments are likely to succeed but measuring the impact, allows the organisation to pivot and fail fast from experiments which are unlikely to create value, while focusing on use cases that will.

In its most elementary form, the measurement of RoI has two parts:

1. ‘The Measure’ – the aspect of business which needs measuring. Good examples are efficiency in customer request processing, revenue growth or cost avoidance etc.
2. ‘The Metric’ – the specific mechanism used to assess the measure of success. Examples are cost associated in dealing with a customer request in a bank call centre, product mix ratio, penalties avoided quarter-on-quarter through improved compliance.

These measures and metrics should have a baseline from before AI systems start operating, and then be used to measure their viability in the face of business impact.

Equally, measuring the impact also reinforces the positive message of the business impact created through AI. This builds the case for wider adoption across the business for solutions leveraging AI.

## Conclusion

The financial services industry has a rich history of analysis and decision making drawn from large data sets. With the wider availability, cost and ease of use of AI solutions, there is now an opportunity for viable and useful access to data driven-decision making at scale. However, with the opportunity comes the need to use AI responsibly and to consider the broader reaching implications of this transformational technology's use.

In this paper, we've explored the different stages of maturity in the AI journey and introduced a maturity model to support this process. As part of this journey organisations will need to consider how they grow and adapt their strategy, culture and organisation in parallel with technical capabilities. This is key for firms to holistically support the successful adoption of AI in a responsible manner, ensuring accountability for AI systems.

Scaling AI ambitions successfully past first initiatives brings new considerations. As AI projects typically require greater upfront investment than traditional solutions, and increased ongoing maintenance, understanding the medium to long-term ROI becomes an essential tool in understanding the right time and place to deploy an AI system – or not.

To avoid the first pitfall – looking at AI simply as a technology tool, and not considering the wider cultural and organisational changes necessary to become a mature AI business. To avoid the second pitfall – failing

to implement and scale AI – success requires ongoing collaboration. AI skills and expertise should not just be in the domain of technology teams; instead, the opportunities, implications and responsibilities should be shared across an organisation.

Across organisational change and AI implementation, organisations should also consider how their governance arrangements enable the use of AI. The governance layer needs to be appropriate to a firm's level of AI maturity and the level of new risk introduced but also must be built to be rapidly scalable as needs demand. The introduction of this governance is the time firms should step back and consider the broader implications of AI and its appropriate use. This paper has introduced four ethical principles of **fairness, privacy and security, transparency and accountability** as a starting guide for firms to consider.

Building a solid foundation now for governance for AI will stand organisations in good stead, both to manage programmes and spend effectively, and deliver high quality services to customers. Skilled and experienced AI resource is still scarce so expanding awareness of AI at both business and technical levels is vital to taking AI from experiment to broadscale production usage.

As AI is an evolving field, this paper cannot answer every question. It is clear however that by having an organisation-wide approach, and the tools to implement AI in the most appropriate areas, firms can take advantage of AI while still delivering resilient, responsible and high-quality services to their customers. It is hoped the practical tools and approaches in this paper will help firms begin this process and position them well for the future.

---

## CONTRIBUTORS



### **Jonathan Middleton, Manager, Digital and Technology Policy, UK Finance**

Jonathan Middleton joined UK Finance in 2018 as Manager for Digital and Technology Policy. He oversees work on a range of topics including innovation, fintech, disruption, regtech and AI.

Prior to joining UK Finance, Jonathan was Head of Technology Policy at the Government Digital Service. Jonathan is a graduate of Oxford University and the University of London. He is also a Fellow of the Royal Asiatic Society.



### **Walter McCahon – Manager, Data Policy, UK Finance**

Walter McCahon oversees policy work on privacy and data protection issues at UK Finance, particularly the General Data Protection Regulation and Data Protection Act. Before coming to UK Finance, he worked for the BBA in various policy areas, including privacy, data protection and Open Banking. He previously worked for the New Zealand Banker's Association and New Zealand Ministry of Economic Development on financial sector regulatory reforms.



**Kate Rosenshine - Head of Data & AI Cloud Solution Architecture, UK Financial Services, Microsoft**

Kate currently leads the Data and AI Cloud Solution Architecture team for Financial Services at Microsoft UK, helping organisations shape their data strategies in a scalable and responsible way. Her main focus lies in the intersection between technology and business, how data can shape organisations and AI systems.

Prior to joining Microsoft, she worked at start-ups where she focused on leveraging big data and behaviour analytics to augment decision making. Kate comes from a background in scientific research, specialising in neurobiological genetic engineering. She holds a MSc in Molecular Biology from Bar Ilan University and a MBA from Tel Aviv University.



**Janet Jones - Head of Industry Strategy, UK Financial Services, Microsoft**

Janet currently leads the Industry Strategy for Financial Services at Microsoft UK, ensuring that drivers of change and emerging technological trends across the sector are core to how Microsoft works with financial services organisations, supporting their digital transformation.

Before joining Microsoft in 2018 Janet held roles within commercial banking; most latterly at Lloyds Banking Group and prior to that at Barclays and RBS.



**Hannah Cawthorne – UK Financial Services Industry Lead, Microsoft**

Hannah Cawthorne is the Digital Transformation Solutions lead for Financial Services in the UK. Hannah works closely with both financial institutions, member bodies and financial services solution providers to ensure that Microsoft is assisting our customers transform to meet the needs and opportunities in the digital economy.

Hannah is an experienced leader across change and cross group collaboration and has worked in or with financial services organisations throughout her 20-year career. She also spends time working with school children to help them understand future career opportunities and learn skills to assist them with their future.



**Angus Duncan – Senior Digital Advisor, UK Financial Services, Microsoft**

Angus' current role as Digital Advisor for Microsoft involves helping organisations in the UK financial services sector digitally transform, architecting and delivering change through technology projects. He has a particular interest in Open Banking and Blockchain opportunities.

Angus 20+ year career is based upon an understanding of business drivers and the positive impact technology can have on these.

## ACKNOWLEDGMENTS

**Marina Arnaout** – Microsoft

**Kalai Elumalai** – Microsoft

**Richard Fitzgerald** – Microsoft

**Levente Nagy** – Microsoft

**Kostas Papadopoulos** – Microsoft

**William Maunder-Taylor** – SparkBeyond

**Reecha Mishra** – Microsoft

**Angelique Mohring** – GainX

**E-Yen Tan** – Microsoft

**Rishi Srivastava** – Microsoft

**Frances Zelazny** – BioCatch

This report is intended to provide general information only and is not intended to be comprehensive or to provide legal, regulatory, financial or other advice to any person. Information contained in this report based on public sources has been assumed to be reliable and no representation or undertaking is made or given as to the accuracy, completeness or reliability of this report or the information or views contained in this report. None of UK Finance or any of their respective members, officers, employees or agents shall have any liability to any person arising from or in connection with any use of this report or any information or views contained in this report. © 2019, UK Finance.

