EVIDENT AI OUTCOMES REPORT

Towards a benchmark for AI outcomes in banking

IN PARTNERSHIP WITH

DeepSee

Disclaimer

Evident AI Outcomes Report, October 2023, is published by Evident Insights.

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

EVIDENT AI OUTCOMES REPORT

Contents

Foreword

Since the launch of the Evident AI Index in January 2023, we've spoken with CIOs, CTOs, CDOs and other leaders across fifty of the world's largest and most successful banks. Almost universally, those decision-makers are focused on one thing when it comes to AI: outcomes.

They are considering which use cases to prioritise—from the huge and complex opportunity presented by artificial intelligence. They want to know how those use cases can deliver value for their organisations now and into the future. And with banks looking to accelerate their adoption and reach outcomes through the use of AI, many are also evaluating how they compare to their peers. These questions are difficult to answer; but Evident is on a mission to do just that.

By building an outcomes benchmark, and providing a common framework which will enable banks to evaluate and compare their outcomes across the sector, Evident's ambition is to provide a common language and conceptual framework for the output from AI use cases—in the same way that we do for the inputs such as talent, innovation and leadership. We are not yet at the conclusion of this particular journey, but this report marks the first step.

We have explored and captured the emerging best practice among banks for understanding the AI opportunity; measuring success; ideating use cases; prioritising their development and reaching the delivery stage at which outcomes can be measured. The report is based on a series of interviews and conversations with senior AI leaders between May and September 2023; and we are immensely thankful for their contribution. We are also grateful to DeepSee, for their invaluable insights and collaboration on the content of the report.

ACKNOWLEDGEMENTS
Special thanks to **Elizabeth Dunmore**, **Luke Gbedemah**, and **Tim Gordon** on their contributions to the writing of this report.

Alexandra Mousavizadeh
CEO & Co-founder
Alexandra.mousavizadeh@evidentinsights.com

Annabel Ayles
COO & Co-founder
Annabel.ayles@evidentinsights.com

Executive Summary

1. While banks have experimented with AI use cases for years, and yielded significant results to date, the last year has seen a step-change increase in ambition and investment. Banks are generating more ideas for AI use cases than ever before, face more pressure from leadership demanding ROI, and need to keep up with the rapid pace of AI innovation. Working out how to scale up use cases, deliver value, and orchestrate the AI activities across the company has, in many banks, become the mandate for newly established group AI leadership teams.

2. AI teams in leading banks are focused on building capabilities across five core priority areas:
   → **MAP** current, and potential, AI use cases
   → **MEASURE** the value of those use cases in terms of outcomes
   → **IDEATE** the most relevant and addressable use cases for the bank
   → **PRIORITISE** which uses cases to pursue
   → **OPERATIONALISE** those use cases to deliver results

3. **Map:** Use cases are the building blocks of AI delivery. The exact definition of a use case can vary between banks (is it referring to the individual AI model or the collection of models in a sophisticated business proposition, for example). However, AI leaders in banks must establish a consistent and standardised classification and terminology.

4. **Measure:** Shareholders and senior leaders are increasingly demanding tangible outcomes from AI investments. Banks need a common methodology and process to measure, track and report on the value created by their existing (and future) AI use cases.

5. **Ideate:** The best use cases are intimately tied to business problems, and ChatGPT has led to the proliferation of ideas of AI use cases like never before. Leading AI teams are investing in initiatives to fuel and harness this bank-wide AI ideation, such as increasing AI literacy; embedding AI teams within the business lines; investing in cross-organisation knowledge sharing amongst AI talent; and establishing central use case ideas libraries.

6. **Prioritise:** There may be few limits to the opportunities offered by AI – but delivery resources are always constrained. Banks need a robust, standardised and aligned process to prioritise AI use cases for delivery. This has to cover ROI, operational capacity, risk and governance issues.

7. **Operationalise:** Delivering value from AI at scale requires that foundations be well laid. Banks we've interviewed are focused on investing in foundational AI tools; delivering on long-term data strategy; establishing external partnerships in priority areas where they lack in-house expertise; and ensuring that model validation frameworks are fit-for-purpose, encompassing Generative AI.

8. The race for AI outcomes is only accelerating, and best practices across the sector are beginning to emerge. To that end, we have built an initial list of KPIs that banks can use to assess their progress against the five capability areas explored in this report. In the coming months we will be expanding this list and benchmarking banks' progress against this framework. **If you are a bank and are interested in participating, get in touch.**

CAPABILITY AREA	KPI
Map	Does the bank have a central repository of the organisation's AI/ML use case portfolio?
	Does the bank have a common internal language for defining a 'use case'?
	Does the bank have a common internal language for defining 'artificial intelligence'?
	Total number of use cases
	Distribution of use cases across business lines or functions
Measure	Does the bank use a common measurement framework to assess the ROI of AI use cases?
	Does the bank assess all AI use cases against this ROI framework, at deployment and over time?
	The total revenue uplift from AI use cases
	The total cost reduction (or efficiency gains) as a result of AI use cases
	The total risk reduction (or avoidance) as a result of AI use cases
	The total customer satisfaction improvement from AI use cases
	The total staff satisfaction improvement from AI use cases
Ideate	Is there a formal process at the bank to capture ideas for AI use cases?
	Is there a clear approach to provide support to staff members generating ideas for AI use cases?
	Number of use case ideas generated within a given period
	Proportion of use case ideas generated by technical and non-technical employees
Prioritise	Does the bank have a common (centralised) evaluation framework to prioritise AI use cases?
	Proportion of use case ideas that are approved for POC development
	Time taken for a use case to get approved for POC development
Operationalise	Does the bank have a centralised platform for developing AI use cases which employees across the bank can access?
	Proportion of approved use case ideas that end up in POC
	Proportion of approved use case ideas that end up in deployment
	Time taken for an AI use case to move from approval to POC
	Time taken for an AI use case to move from POC to deployment
	Proportion of use cases that meet initial cost expectations
	Proportion of use cases that meet initial time expectations
	Proportion of uptake among target users of the AI use case

Introduction
The Agenda for AI Leaders

This is an era of discussion and debate. Leaders of business units are being challenged to answer questions ranging from strategy and business process to technology, risk and staffing. The AI debate can end up as a full spectrum recasting of every aspect of their business.

To complicate this they need to be thinking across multiple time horizons. While most use cases are focused on the present, others are focused on things that the bank wants to do but never thought was possible, and further still, things that are new to the world at large.

"Our approach is to look at the opportunity across different time frames; optimize for today, build for tomorrow, and shape the future."
Steve Van Wyk, Global CIO at HSBC

What is clear is that leaders across the banks—guided by those with their hands on the AI steering wheel; CIOs, COOs, Heads of AI and Chief Innovation Officers—need to do three things in parallel and at the same time.

→ **Build a vision to change the bank, not just run it:** for years, banks have focused on AI as a tool for process optimisation. Now they need to rethink what it means to be a bank by imagining new ways of delivering products and services. This is no longer solely about identifying existing processes and improving them, but about rethinking the nature of bank operations. At the heart of this future-facing conversation is what the bank will be as it enters the AI era. Historically businesses like banks have thought of themselves in terms of their physical facade – the buildings that host clients and project their brand image in key markets. The AI bank is an edifice built around data and AI infrastructure. Making this jump will not be an easy one.

FIG 01.
FUTURE ARCHITECTURE OF AI-FIRST BANKS

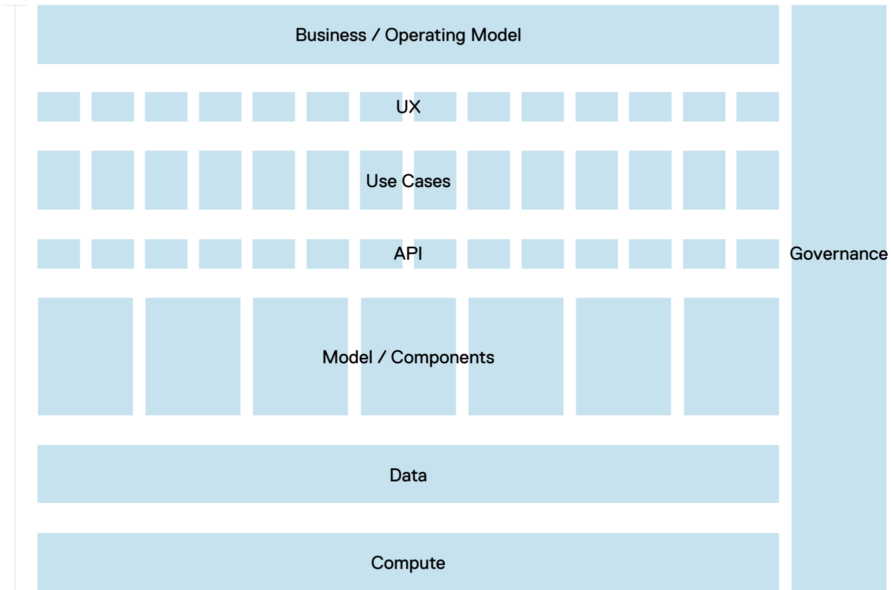

→ **Investing in capabilities that are repeatable and scalable:** leaders need to understand what capabilities, skillsets and investments will be needed to deliver this future. And they need to design the organisation to deliver on the agenda. Balancing this with urgent delivery to a business ever-hungry for new products and transformation is tough. The aim has to be to build the repeatable tasks and tooling that will provide compounded returns. Getting down the experience curve will enable faster deployment and swifter iteration: the key to building the learning organisation required to maximise output of learning machines. Simply put, the more an organisation does something the better they get at it, bringing down lifetime costs and speeding up delivery. AI may occasionally feel like an esoteric new sphere but the lessons of traditional strategic economics remain relevant as ever.

"A use case should contribute to a snowball of information that the models learn to recognise and use, where this rolling momentum means the bank's adoption velocity increases."
Ryan McQueen, Head of Product at DeepSee

"AI leaders in banks are rebuilding their aeroplane as they fly it, through fog and turbulence. They also have a full load of passengers and an inspector from the National Air Authority onboard who expects it to arrive on time."
Tim Gordon, Partner at Best Practice AI

→ **Fund the journey by generating value today:** leaders need to generate (and extract) value today, learning as they go. The market expects both digital transformation and ever-strengthening shareholder returns. There is limited appetite for this to be a trade-off. The economic challenge from AI—that it takes variable cost exercises and recasts them as fixed cost processes—is less stark for banks than it might be for, say, professional services firms. However, the strong likelihood is that aggressive AI roll-out will drive competitive price pressure; whether from new start-ups operating to tighter market segments without legacy costs or simply other legacy banks using AI to scale up their offerings and improve efficiency. Extracting value and demonstrating this publicly will also be key if bank management wants to maintain the support of key stakeholders on what may be a long journey.

The well-planned tortoise, ideally backed by significant central assets and smartly procured external advice and toolkits, should be able to beat the disorganised hare, especially if the latter is busy reinventing the wheel for each product roll-out cycle. The aim is to build and optimise an AI Operating Model, a resilient solution that will evolve over delivery cycles.

THE CHARACTERISTICS OF LEADING BANKS

The nature of this challenge varies for individual banks. For market leaders it is how to accelerate into the transition and capture the compounding impact of internal expertise and tooling. For less AI mature banks, it is assessing which use cases to start with, how to prioritise them and how to build the capabilities necessary for success.

KEY CHARACTERISTICS OF MORE MATURE BANKS:

→ AI is an established item on the CEO agenda
→ AI is embedded as a core component of the bank's strategy
→ Significant investment is being made into AI research and development units
→ Significant investment in standardised data platforms
→ Cloud compute experience and relationships

By contrast, less mature banks were still struggling with articulating bottom-up use case experimentation and innovation. They were solving point to point AI problems, focusing on specific use cases. Rather than an AI strategy, they were focused on using AI to create efficiency against a pre-existing strategy.

KEY CHARACTERISTICS OF LESS MATURE BANKS:

→ Use case successes in individual business units
→ No overarching group AI leadership or coordination
→ No centralised group AI strategy
→ Varying stages of cloud sophistication and data infrastructure

THE IMPACT OF GENERATIVE AI

Generative AI however may yet change this game. There are three ways that this could happen.

→ Firstly, it has changed the nature of the conversation, creating a new dynamic from boards down. A key element of early leadership in AI was simply "*getting it*": that is no longer a relative competitive advantage.
→ Secondly it may upend some of the dynamics of early-mover product advantage. Chatbots, for example, were for many years the quintessential customer-facing use case. No pre-ChatGPT chatbot can any longer be considered state-of-the-art. Investment now will deliver a suite of products that is at the cutting edge.
→ Finally, and perhaps most controversially, LLMs may allow a bank to leapfrog the competitive edge built up by early leaders in technical skills and deep data

sets. Decades-worth of data infrastructure investment may now be replicable in far shorter time frames and lower cost – for example as smaller and smaller data sets are required to fine tune models.

While this may still be up in the air, experience counts. Deeply rooted cultures focused on optimising data-driven business models will still have a competitive edge. But change may be in the air. The race is not yet over. By some measure, it has barely begun.

EXAMPLE | RETHINKING THE ART OF THE POSSIBLE AT HSBC: THE QUEST FOR CONTEXT DATA.

We interviewed David Rice, Global COO, Commercial Banking (CMB) at HSBC, about his vision for AI at the bank, and what has changed in the last year:

"Usually, business strategy has been about defining where you think you can get competitive advantage, and technology has always been thought of as a way of enabling that strategy to be executed. The difference with AI is that it is so pervasive it actually has the potential to change what a bank might be.

It all comes down to this: How will our customers want to interact with a bank in the future, and how does this technology define how they interact? We believe we have to build an AI-centric organisation.

Since the start of time, banks have not been able to grab context data, we only find out what happened after it happened. Take a transaction… We know that it happened, but not why it happened, or how HSBC could have helped. We solved this in the past by having relationship managers, but this is not scalable. If you put a Digital Agent with every CFO / treasurer, you can feed that context data in, identify needs that we couldn't before, and start trying to solve those problems.

AI will not be a competitive advantage, but the proprietary data you can gather using AI will become a compounding benefit. This is the art of the possible. We haven't had to think about what it means to be a bank like this since the advent of the internet, or perhaps the explosion of mobile.

There will be two types of companies: the first will be efficiency-focused. There are huge efficiencies to gain here: 50/60/70/80% efficiency gains in core processes. There will be many organisations that take that efficiency out. But there will be others [the second type] that reinvest that money. If you can change the culture, bring the workforce to learn new skills, better empathy, and customer service because the role has changed, those are the companies that will win."

THIS REPORT

Banks are not sitting on their hands, they have already begun to extract value across AI use cases and capabilities. In the following chapters we explore five key capabilities that are critical in order to deliver AI outcomes at banks. As we move towards a benchmark for AI outcomes we have itemised key performance indicators (KPIs) to help banks assess their strength in each area:

→ Chapter 1: **MAP** current, and potential, AI use cases;
→ Chapter 2: **MEASURE** the value of those use cases in terms of outcomes;
→ Chapter 3: **IDEATE** the most relevant and addressable use cases for a bank;
→ Chapter 4: **PRIORITISE** which uses cases to pursue;
→ Chapter 5: **OPERATIONALISE** those use cases to deliver results.

USE CASES ARE THE BUILDING BLOCKS OF AI ADOPTION

The opportunity presented by artificial intelligence in banking is vast. Through examination of publicly available communications about use cases, academic literature and press releases, Evident has identified the following portfolio of opportunities, grouped into "use case families" which show specific business functions or lines affected by an AI tool or model.

Although this list continues to grow, it is a starting point. It is not exhaustive and new applications for AI in banking are being developed on a constant basis. Unique characteristics in each bank's operating system, regulatory environment, culture and product mix will throw up increasingly bespoke AI solutions and use cases.

The smarter banks have got a handle on their own set of use cases. It is critical that they have evolved a common methodology for identifying, tracking and measuring use cases internally. Leading banks have started to report on the volume of use cases metric.

BANK FUNCTION	USE CASE FAMILY
Credit scoring	Creditworthiness Assessment (Internal) Credit Risk Assessment Determining loan approval probabilities Approve or Decline Assessment Default Prediction Underwriting (Customer recommendations)
Customer acquisition	Personalised Marketing (Text Production) Targeted Marketing Campaigns Secure Digital Onboarding Cross-selling and Upselling (Insight generation)
Customer experience	Client-facing Chatbots Automated Savings & Guidance Cash Flow Forecasting Customer Interactions (Interpretation) Query Resolution (Customer-facing) Query Resolution (Learning)
Customer retention	Lifetime Value Prediction Customer Segmentation Customer Churn Rate Prediction Customer Recovery Recommendations Customer Sentiment Analysis
Trading	Tax Loss Harvesting Advice Portfolio Personalisation Trade Assistance Chatbot Email Classification and Extraction Generating Email Responses Settlement Failure & Anomaly Detection Generating Email Responses Investment Decision-making (Market Trend Analysis) Equity Client Execution Algorithms Portfolio Personalisation ESG Analysis Bespoke Index Creation & Maintenance Algorithmic Trading Market Research & Analysis (Sentiment Analysis) Price Prediction Currency Risk & Hedging Process Margin Reduction Trading Trade Allocation & Affirmation Deal Review (Contract Evaluation) SSI Validation Payment Notification Transaction Validation
Process Automation & Knowledge Management	Back-office Processing (Documentation) Document Ingestion & Summarization Internal-facing Chatbots RPA (Robotic Process Automation Robo-advisors Automation or Support of Human Decisions (Human Resources) Workforce Planning Candidate Sourcing (Human Resources) Low-code Systems Programming Support Personalised Business Proposals Processing publications on banking supervision Workflow Optimisation Claims Processing Automation or Support of Human Decisions (Diagnostics Engines)
Risk Management	Onboarding Validation (Source of Funds) Counterparty Risk Management Know Your Customer (KYC) Perpetual Blacklisting Counter-terrorist Financing (CFT) Anti-money Laundering (AML) Checks De-duplication of Entity Master Records Sanctions Monitoring & Compliance Cross-analysis of Suspicious Activity Reports
Agent-based modelling	Large Scale Market Simulation Non-linear Risk Mapping
Stress Testing	Warning Systems Dynamic Balance Sheet Simulation
Fraud Detection	Real-time Anomaly Detection (Transactions) Payments Fraud Credit Fraud Claims Fraud Account Takeover Detection Fraud reporting Secure Digital Onboarding Threat Actor Mapping (Simulation Modelling) Malware Detection
Synthetic Data	Model Creation Test Environments
Regulatory Compliance	Macroprudential Surveillance AI in KYC & AML Compliance Data Quality assurance Supervisory technology Regulatory technology
Data Architecture	Aggregation Algorithms (Data Enriching) Data Summarising Automatic Data Cleaning Intelligent Data Serving (for production)
Data Capabilities	Communications Classification Term Extraction Reconciliation & Entity Resolution (Data Point Matching) Validation Prediction Audit & Auditor Toolkits Ontology Creation Semantic Search Insight Generation

FRONT OFFICE

MIDDLE OFFICE

BACK OFFICE

FIG 02. MAP OF AI USE CASES

Scale indicative of front office to bank office application

REPORTED USE CASE VOLUMES AT LEADING BANKS

SELECTED BANKS AND THEIR USE CASE REPORTING

BNP Paribas "More than 700 use cases have been rolled out with the intention to create significant value by 2025 of more than €500 million per year [...] We have also identified 100 use cases for generative artificial intelligence experiments with large language models used, for example, by ChatGPT or Bard." Source

DBS Bank "DBS today runs more than 300 AI and machine learning projects, which it says yielded a revenue uplift of SG$150 million ($112.53 million) last year and saved SG$30 million ($22.51 million) in risk avoidance, for example, from improved credit monitoring." Source

JPMorgan Chase "We're ahead of our plan to deliver on our commitment to deliver $1 billion in business value through AI... The firm has increased its artificial intelligence and machine learning use cases by more than 34% year over year, with more than 300 use cases in production." Source

Société Générale "As of 2022, the Group's portfolio has around 340 Data and AI Use Cases (UCs) in production, of which 170 are AI-based, all working to best apply our strategy with an expected value creation of €500 million." Source

HSBC "We currently have nearly 1,000 AI use cases across our global operations, in use or in testing, covering a range of areas from detecting financial crime and fraud, helping our customers budget better, and more. We're also testing and learning across a range of Generative AI (GAI) use cases across HSBC, and are in the process of scaling up a small number in secure environments." Source

Other banks – such as Intesa Sanpaolo, Royal Bank of Canada, Deutsche Bank and Capital One – have consistently published statements about specific use cases, including the many different applications of AI/ML across the banks as well as some of the operational or business outcomes. They have also detailed some of the value or outcomes that they have generated from those applications. They have not yet made specific statements about the total number of use cases for AI/ML in the bank, or the associated type of return.

AWARENESS AND COMMON LANGUAGE MATTER

Comparing even the limited statements made by these banks should be done with caution. As banks begin to report on the number of AI use cases that they have in production the definitional inconsistencies present a challenge. These comparability challenges include:

→ **Level of detail matters:** the taxonomy of use cases is complex, especially as customer-facing use cases may involve multiple underlying models, many of which may be use cases in their own right. Meanwhile the style of deployment may vary – in one bank a chatbot might cover the entire customer relationship (one use case) whereas in another chatbots might be deployed per product (current account versus mortgage say) or per process (account queries versus complaints for example).

FIG 03. ILLUSTRATIVE VIEW OF HOW BANKS' DEFINITIONS OF AI USE CASES IMPACTS REPORTED VOLUME

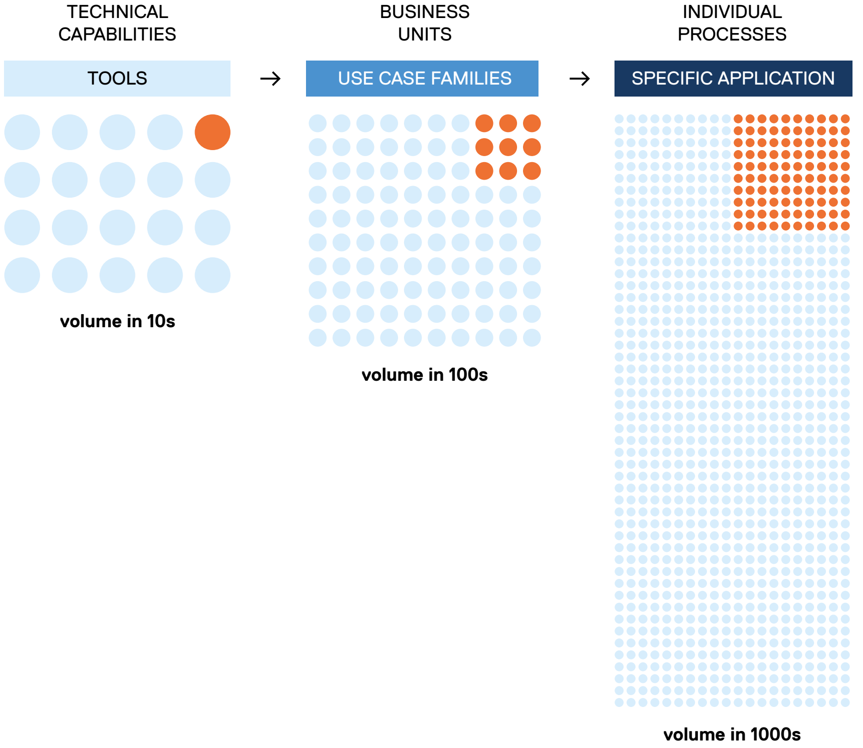

Fig. 03 (above) shows the different ways in which a given bank might represent its use cases, and how it may differ from another in terms of the total number of 'use cases' it would then report. One bank might take a single technical capability, e.g. a semantic search model, and deploy it in 100 different functions across the banks, leading to 100 use cases. Another bank might build another technical capability, e.g. lifetime value prediction, and deploy it as a feature of their centralised data platform, leading to 1 use case with many different areas of impact.

This difference in approach and calculation means that the reported numbers of use cases across banks are not strictly comparable. Where one bank might build an underlying technical capability such as value prediction, including a model-based tool and consider this to be an AI/ML use case (the highlighted single block), another might build a value prediction model that supports life-time value prediction for commercial banking customers, high net-worth individuals and others in specific business units (several highlighted blocks) and, finally, another bank might apply lifetime value prediction to a whole range of specific applications across its online and mobile banking offerings, and for different banking products such as current accounts, trading or others (many highlighted blocks).

→ **Stage of maturity matters:** use cases can be at various stages of production, and a use case might refer to a product that is being planned, is fully deployed or in an intermediate stage of development. The gap between a use case undergoing proof of concept, and one that has been tested, validated and scaled up is a wide one.

EXAMPLE OF VARIABLE LEVELS OF MATURITY WITHIN A USE CASE: TRADE ALLOCATION, WITH EMAIL CLASSIFICATION AND EXTRACTION.

A bank in the early stages of AI implementation, without a working system to address trade allocation tasks, may have to manually classify relevant data: individually opening emails and attachments to search for details on the trades in question.

A mature bank may have a system in place to make all extracted counterparty affirmations machine readable, once the content has been manually reviewed, interpreted and entered into a structured file that can be uploaded to a risk system to allocate a customer's trades.

The more advanced banks, in terms of AI implementation, may have deployed a system to read each incoming email message, classify those containing trade allocation information, extract and analyse that information and save it into a structured file without intensive manual intervention. This information can then be sent on to the trade booking system for processing.

The most advanced, 'AI-first' organisations are not just improving a singular process, but changing the way the bank is run over time; by gathering data from an operation like trade allocation, AI can eventually contribute to organisation-wide knowledge used across multiple functions such as painting a consolidated list of counterparty risk or predicting future trade allocations to counterparties based on past behaviour.

The resulting value delivered comes in a variety of areas: from reduced time to allocation on trades, increased efficiency in terms of the resource cost of processing trades and conducting risk reporting. The collection of operational data on trade allocation may also be used to drive actions based on historical activities; such as identifying patterns in trade data. Subject Matter Experts (SMEs) can also divert more time and attention to cases in which incomplete allocation instructions have been received, or an error seems to have occurred.

→ **Complexity matters:** use cases can also vary in complexity. Some involve many more layers of data ingestion, learning, testing or implementation. In addition, internally facing use cases that do not involve heavy regulated functions, sensitive data or some other factor which might be simpler to deploy. This will impact on the risk profile, further complicating the picture from a management perspective.

→ **Source of use case:** banks deal with thousands of suppliers, many of whom are now rolling out AI use cases. Some will be critical to a bank's operations. Typically buying in these services will be less onerous on internal resources (at least initially) than building internally. Whether these should be counted for a bank's stated number of use cases deployed is obviously not set in stone but will impact on comparability of, for example, a bank's AI operating model.

→ **Calculation of return:** as we will discuss in Chapter 2, there are a number of different methodologies for measuring return on investment (ROI) into AI/ML use cases. Some banks choose to categorise returns purely in strict financial terms, or as a blend of financial and non-financial factors. This significant variability in calculating ROI means that public statements may not be fully comparable.

Evident, in conversation with a number of leading banks, has explored what use cases mean, how they are thought about and prioritised, and which factors go into driving success. The language surrounding "use cases" is inconsistent in much the same way that banks' public communications about "AI" can involve different definitions.

> "It is crucial that the definition of 'AI' remains broad and encompassing, due to the constant evolution of the technologies, and the rapid advancements in their application to banking."
> *Ash Booth, Global Lead Applied AI/ML in Markets at JPMorgan Chase*

The opportunity for using AI in banking is huge. The number of options can be bewildering, and the opportunity cost of pursuing the wrong initiatives has the potential to be very high. The following KPIs are a guide to how banks can assess their level of maturity when it comes to evaluating their AI portfolio, and the extent to which they are seizing the AI opportunity.

KEY PERFORMANCE INDICATORS FOR MAPPING USE CASES

	KPI	DETAIL
Process KPIs	Does the bank have a central repository of the organisation's AI/ML use case portfolio?	*A centralised dashboard that shows specifications on the use case and dependencies.*
	Does the bank have a common internal language for defining a 'use case'?	*A consistent set of terms for defining if an initiative should be measured and monitored as a use case, and if so what type (which may trigger a wider set of protocols, for example for risk management or governance oversight.)*
	Does the bank have a common internal language for defining 'artificial intelligence'?	*A consistent set of terms for defining what 'AI' means within the organisation, and how it relates to 'data' and other analytical techniques.*
Quantitative KPIs	Total number of use cases	*By type, by technology, by launch date, by source (e.g. internal and suppliers)*
	Distribution of use cases across business lines or functions	*By business unit, by function, ownership and usage*

MEASURES OF USE CASE SUCCESS

Every use case should produce measurable outcomes within the bank. Each use case will be different, and may produce multiple potential outcomes. Broadly, AI use cases can drive value across five different buckets:

→ **Income Uplift:** AI can be used to create and augment new revenue streams. For example, use cases can contribute to portfolio returns through strengthening research and analysis, or helping to boost product cross-selling and upselling. They can drive fee-based income in wealth management, investment advice and insurance. This is the edge that banks are chasing. As AI becomes more advanced and more players enter the mix, this edge is getting ever narrower.

"All of the different entities within the bank may have different priorities, so we ask, at a group level, to always generate value. What that value is depends on the different strategies involved: a risk department will be more focused on risk management and efficiency, while a business department will be more focused on developing banking income, on top of efficiency."
Etienne Guibout, Head of Group Innovation Data & AI at Société Générale

→ **Efficiency Gains / Net Cost Reduction**: one of the main opportunities from AI use cases is the potential to improve operational efficiency; requiring fewer resources to perform a given process, for a greater overall yield or to reach an outcome more quickly. Banks have been chasing margin gains for decades by making incremental improvements and AI has opened up a vast new frontier of applications for efficiency-boosting augmentation.

BNY MELLON: PAYMENT INSTRUCTION AUTOMATION

At BNY Mellon, teams identified an opportunity to deploy AI in payment processing. Part of the business deals with tens of thousands of payment instructions from clients coming to the bank in the form of email attachments. And no two formats are the same, so key inputs like the payment information, beneficiary account numbers are all in different formats, in different parts of the document.

The complexity of the problem meant that standard, off-the shelf document ingestion solutions, struggled with the variety of inputs. BNY Mellon developed an AI enabled solution that determines whether an incoming file includes a payment instruction, then extracts all key information (amount, account number etc.) and makes it available to the processor for validation. This use case is in production, and has reduced payment processing time by 80%.

→ **Risk Reduction / Avoidance:** mitigating exposure to risk is essential to operating competitively. AI techniques can provide a toolkit for understanding, controlling and mitigating a number of risk factors by enabling solutions like anomaly detection, secure onboarding, anti-money laundering checks and cross-analysis of suspicious activity reports. Its capacity to analyse huge amounts of data can also be utilised to reduce the risk profile of a bank by identifying previously difficult to see regulatory or contractual instances which leave a bank open to risk.

"AI/ML solutions enable capital market leaders to apply their controls against every instance and consider all information, whereas today these controls are limited by what can be manually processed and issued, and only evaluated with the knowledge of the specific team member."
Ryan McQueen, Head of Product at DeepSee

BNY MELLON: AI ASSISTED LIQUIDITY MANAGEMENT TOOL REDUCES RISK AND IMPROVES EFFICIENCY

BNY Mellon has a responsibility to set aside reserves every single day. This end-of-day cash position is impacted by a myriad of client activities throughout the day, with varying levels of predictability. The BNY Mellon team developed an AI solution which helps predict what the end of day cash position would be earlier in the day.

Importantly, BNY Mellon kept humans in the loop through the process. It was up to the corporate treasury team whether to take the tool's recommendation: AI was augmenting their decision making.

With the successful deployment of the solution, BNY Mellon's corporate treasury team now benefits from managing excess liquidity every day – crucially without introducing any additional risk.

→ **Customer Satisfaction:** a customer's experience can be a key differentiator, especially when they may be increasingly spoilt for choice. Using AI to help deliver a frictionless experience at multiple touchpoints along the customer journey helps build trust and reduce churn: reducing operational downtime, accelerating onboarding processes, and providing round the clock service can all be better (or more cheaply) served by AI.

HSBC: ATM CASH OPTIMIZATION

For example, at HSBC the organisation is using AI to optimise cash in its ATMs globally. HSBC developed a model to accurately predict cash flow at ATMs, ensuring that they had the right amount of cash in line with changing customer needs, and optimising this on an ongoing basis.

Customers experienced less moments when there was no cash available, contributing to an increase in Net Promoter Scores. HSBC also found a correlation between NPS and Revenue Growth, showing how this application of AI has impacted HSBC's bottom line, while growing its relationship with customers.

→ **Staff Satisfaction:** banks are in constant competition to attract and retain the best talent. AI use cases can deliver talent empowerment outcomes through improving the workplace experience. This could be through giving employees the tools to improve their existing workflow, or adapting the profile of the work needed to attract particular talent categories.

> “We intentionally ask at the onset of a use case what would a service or process look like after an AI solution is implemented, and outcomes or improvements we will see that we can measure. Some of our impact considerations may not necessarily be financial, but rather if a solution is going to really enable our employees, and create a positive employee experience.”
> *Michael Demissie, Managing Director, AI Hub at BNY Mellon*

FIG 04. EXAMPLES OF OUTCOME MEASURES

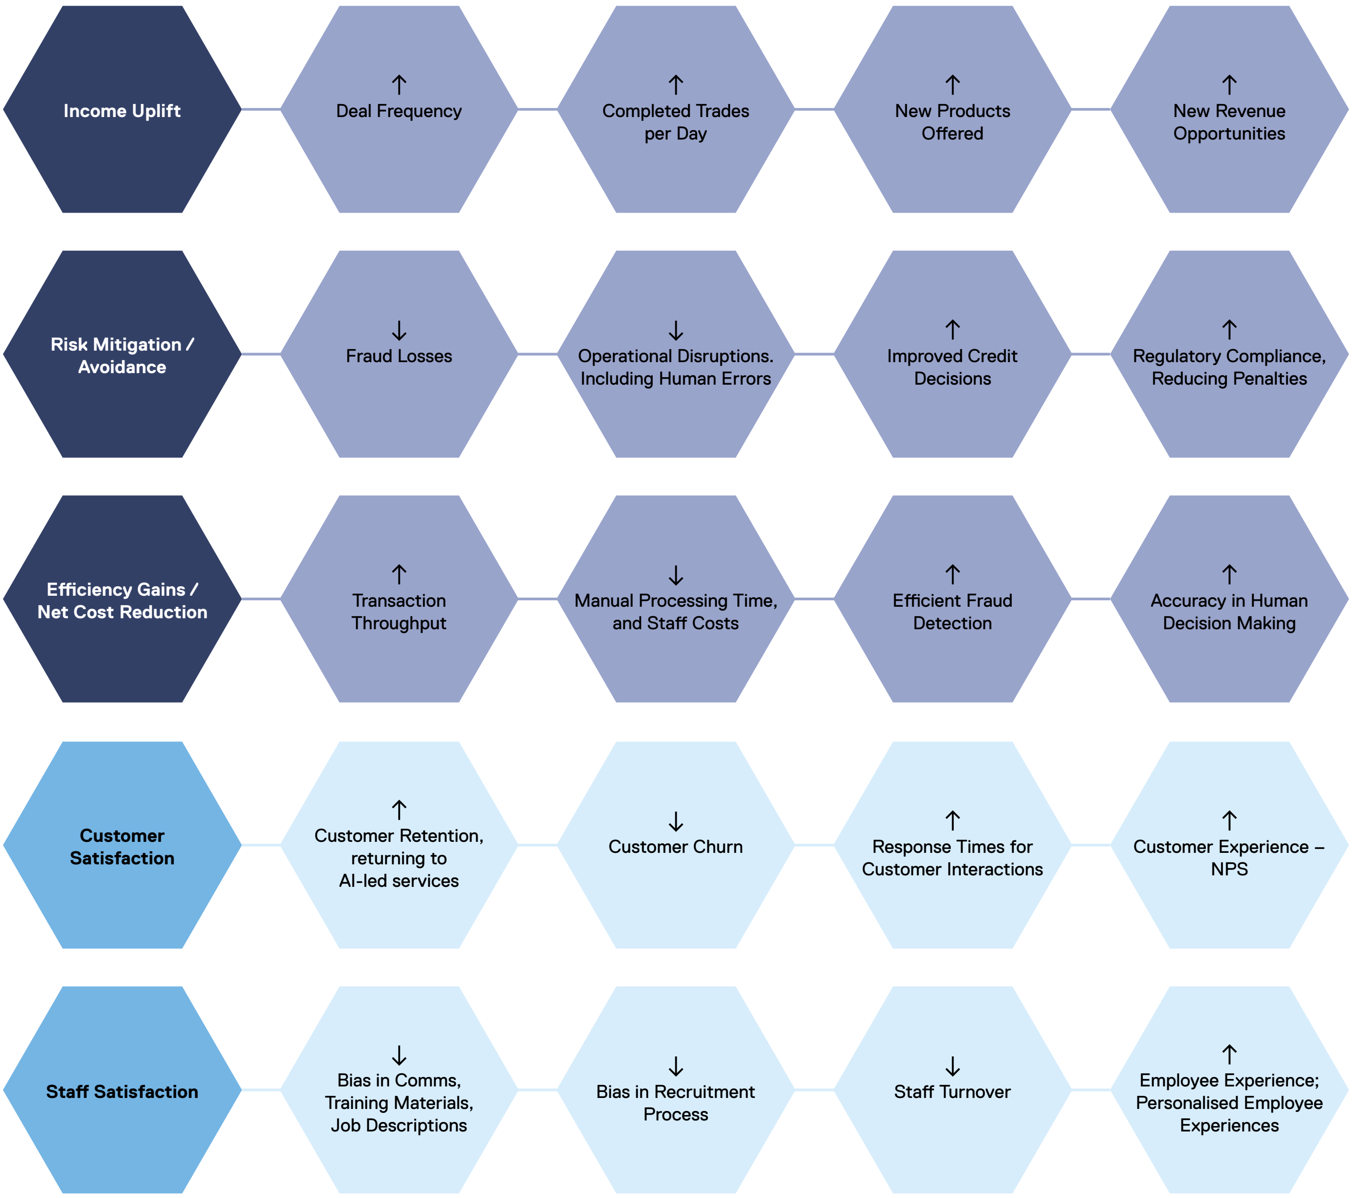

CHALLENGES IN MEASURING OUTCOMES

Accurately measuring outcomes can be difficult. The inherent challenges include:

→ **Assessing return using assumptions:** throughout the process of assessing the projected outcome of a use case, decision makers may need to rely on estimates rather than concrete values. For many use cases there may not be adequate available information, and a hypothesised approach may be required.

"For LLMs or other models, the approach is always the same. The idea is to generate some concrete assumptions which can be refined through time, based on initial experiments.

For example, consider the time saved processing documents with AI – we know that LLMs can augment the performance of employees through multiple applications: summarising & extracting insights from unstructured data sources, interpreting texts & transcripts, generating content and source code, and improving customer engagements thanks to facilitated interactions. As we explore the full potential of LLM technologies, we continue to test our initial assumptions with reality."

Etienne Guibout, Head of Group Innovation Data & AI at Société Générale

→ **Isolating the impact of AI:** in many use cases the deployment of AI may only be a part of the overall solution. For example, if AI lies at the heart of a new business proposition that allows a bank to cross-sell solutions to the existing client base how would one account for the impact of AI? Is it the entirety of the new revenue produced or should some be apportioned to other input areas: salesforce deployed or the wider product infrastructure? Apportioning economic value add inside a bank is not a new challenge but equally what we know about the accounting complexities (and the politics) of these exercises should suggest caution on claimed outcomes.

→ **Evaluating potentially unexpected benefits**: AI use cases can have multiple impacts outside of the primary purpose. For example, an AI tool which democratises knowledge for employees may drive value in terms of talent retention. When use cases are novel and banks are experimenting with new approaches, these outcomes may be emergent and unexpected before they are measured.

→ **Accounting for errors & deterioration:** part of this assessment needs to include an assessment of the cost of a model, or use case's, error. Machine errors, bias, misaligned objectives and otherwise faulty implementations of artificial intelligence could result in errors, and models are subject to deterioration over time. Assessing deterioration will require an ongoing assessment of the model's performance, rather than calculating at a single point in time.

→ **Understanding both *defensive* and *positive* measures**: measuring outcomes from "positive" AI use cases which generate new business or uncover new opportunities can be relatively straightforward. For example, identifying revenue uplift from refining a trading algorithm can be measured with some degree of accuracy. However, when considering "defensive" measures, serving to protect existing revenues or gains, measurement can be more challenging. Sufficient A/B testing, or performance measurement before and after deployment of the AI use case can help address this challenge. For example, banks could assess the change in risk through benchmarking the number of incidents or errors which naturally occur through a process, against the AI-enabled solution, to determine its incremental benefit.

"Testing is in the DNA of tech companies, banks have a way to go to get there."
Angelique Augereau, former Chief Analytics Officer at Capital One

ASSESSING RETURN AND FEASIBILITY OF GENERATIVE AI?

Determining the value add from generative AI can be difficult as the technology is emerging. Banks are experimenting. They are searching for where generative AI use cases can deliver the most value. However, as banks get to grips with the technology this needs to be balanced with clear assessment of risks, ensuring that they are well managed. Establishing what the expectation of return is can help direct efforts, rather than blindly experimenting with multiple use cases.

Goldman Sachs' approach is methodical and thoughtful: "We're all anxious to see results right away ... [however it will require] feeling comfortable about the accuracy... in which we feel comfortable that the information is correct and the risks are actually well managed."
– *Marco Argenti, CIO at Goldman Sachs* (VentureBeat 2023)

Metrics are crucial. Without an agreed approach to measure outcomes from AI use cases, a bank's internal effort to accelerate adoption cannot be optimised. The following KPIs are a guide to how banks can assess their level of maturity when it comes to measuring outcomes.

KEY PERFORMANCE INDICATORS FOR MEASURING USE CASES

	KPI	DETAIL
Process KPIs	Does the bank use a common measurement framework to assess the ROI of AI use cases?	*A centralised, universal framework for assessing performance of AI use cases used across the organisation. Clarity on how accounting and other issues are handled and dealt with on a consistent basis.*
	Does the bank assess all AI use cases against this ROI framework, at deployment and over time?	*A clear internal process where the performance of AI use cases is quantified when in operation, and then tracked over time.*
Quantitative KPIs	The total revenue uplift from AI use cases	
	The total cost reduction (or efficiency gains) as a result of AI use cases	
	The total risk reduction (or avoidance) as a result of AI use cases	
	The total customer satisfaction improvement from AI use cases	
	The total staff satisfaction improvement from AI use cases	

SOURCING IDEAS FOR SUCCESS

Use cases start with a problem that needs to be solved or a new technical development to be exploited; in some ways measures of demand and supply. Demand comes from a problem or opportunity identified by the business that requires a solution. Supply from new capabilities or innovations that become technologically available, perhaps through a bank's research or the advent of a technological break-through such as Generative AI. To satisfy this demand, or make use of supply, the ideation process begins, clearly, with an idea.

However, not all ideas are created equal. Generative AI has accelerated the ideation process: the ubiquity of ChatGPT has opened up people's eyes to the potential of AI. LLMs are but one technology, and for some business problems there are other, more suitable AI solutions. Diagnosing which solutions are best requires a forensic understanding of the problems banks are facing.

"Around the time ChatGPT was released, we saw a sharp increase in interest from areas that had not typically leveraged AI – many were curious to understand how AI could help address their business problems. Generally we have been finding those "grassroots" level use cases – ideas that don't come from the traditional data science world – are focused on the democratisation of knowledge, where you can suddenly have incredibly meaningful conversations by harnessing the capability of AI."
Dan Jermyn, Chief Decision Scientist at Commonwealth Bank of Australia

Sometimes the hardest part is really knowing what the problem is. Bringing together the AI and business strategy is critical for driving successful ideas for use cases. Ideas generated without some engagement with the business users may get to the proof of concept (POC) stage, but ultimately not be operationalised as they do not tangibly impact a bank's operations or pain points. They may also suffer from tissue rejection when proposed to local management teams.

"I like to start with impact, first understanding the business value of a use case and what is needed to solve it. Doing so across use cases informs the foundational capabilities needed."
Angelique Augereau, former Chief Analytics Officer at Capital One

Individuals at a bank will be able to identify that there is a problem they face, however they might not be able to clearly identify what the cause of the problem is. For example, traders may identify there are bottlenecks within a trade processing and settlement lifecycle which stop them from settling on time, but forensically parsing them out can be a challenge. If banks do not correctly identify and solve for the right problem, banks face wasted effort and will miss out on the full potential of AI.

USING AI TO IDEATE

To solve this, some banks are using AI as a means to diagnose the problems they face. Banks have vast quantities of relevant and specialised data, specific to their operations which have captured much of the historic processes, successes, and failures at the bank. AI's capability for processing information at scale lends itself to helping decision makers to diagnose and assess the anatomy of a problem, empowering them to make more informed decisions on where a use case might be deployed.

> "One of the ways we use AI to help uncover areas of opportunity is by analysing the bank's internal data. There is a huge amount of data about how the bank operates: where it has operated really well, and where there are potential issues that can be improved. We have been using AI to identify issues such as where the underlying data used to train the LLMs needs more specificity, or if it could be interpreted in different ways depending on context. As well as creating the system, we are also making the banks' underlying processes better."
> *Dan Jermyn, Chief Decision Scientist at Commonwealth Bank of Australia*

HOW CAN BANKS USE AI TO DIAGNOSE EMAIL COMMUNICATIONS

Senior leaders in capital markets manage a portfolio of complex processes across three major categories: Onboarding, Trading, and Settling. There are no comprehensive industry platforms that can handle the exceptions and all of the myriad of problems and data formats. Therefore, email serves as the default medium to coordinate and resolve all exceptions due to its innate flexibility. Email does not care if data is written out in paragraphs, pasted in screenshots, nor attached as PDF documents. This flexibility is why senior leaders manage large shared inboxes located around the globe that ingest thousands (or more) emails each day. A person cannot read all of these emails to create a complete and transparent view of the patterns occurring both today or historically. AI can.

AI can be put to work to determine where AI should be put to work. This sounds counterintuitive, but new models can categorise this flow of communication and operation data, informing managers in many ways including the following:

1. Where would my process benefit the most from automation dollars?
2. Am I focusing on the most important counterparties or tasks?
3. Where is my operational risk today?
4. Am I meeting agreed to client SLAs?
5. Where is my team encountering chokepoints in resolving exceptions?

"The hardest part is deciding where to start. Using AI to diagnose the problem in more detail can help organisations ensure the problem they are solving is the right one."
– Ryan McQueen, Head of Product at DeepSee

However, most banks are not at AI ground zero. New use case ideas build on top of years of AI experimentation and operationalisation across many banks. While this often-decentralised experimentation environment has driven some success, many banks we speak to do not know where or what those successes are, and feel that they may be solving the same problem multiple times, in different business units. If banks baseline their existing strengths and capabilities, there is an opportunity to understand where these can be reapplied and drive success in other business lines.

ORGANISING TO DRIVE IDEATION

Banks need to organise to bring together different teams to ideate use cases. Creating a common language and knowledge can streamline this process, and help to build trust among teams.

> "Unquestionably, developing and using AI will require collaboration alongside some independence by business unit."
> *Jeff McMillan, Chief Analytics and Data Officer at Morgan Stanley*

To support and encourage aligning teams, banks are taking at least three organisational approaches:

1. INCREASE AI LITERACY ACROSS THE BANK

In order for individuals across the bank to come up with ideas for AI use cases, there needs to be a level of understanding of the types or problems AI can (or cannot) help to solve. This can be fostered by general upskilling and training initiatives for the banks. At HSBC, the bank has recognized the need to upskill the business on AI, launching the "AI Literacy Pathways" program – suitable for both technologists and non-technologists. Employees can sign up to different certifications of AI literacy dependent on how deeply they want to engage with the program.

> "HSBC launched its AI Literacy platform to help train, upskill, and prepare HSBC teammates to confidently use and adapt to AI. This platform provides multiple learning pathways to help non-technical teammates who are "AI Curious" as well as to upskill seasoned technologists with new AI specialisms and capabilities. There is a cross-industry "war for AI talent" and this places great importance on developing and investing in AI talent."
> *Ronnie Chung, AI Lab Lead at HSBC*

The best banks embed this as part of the remit of a centralised AI unit, such as a Centre of Excellence or Innovation team. This can be done through multiple avenues, including, but not limited to; hosting "Incubation sessions" and webinars with business units to foster ideas; conducting "outside-in" benchmarking to identify interesting use cases, and identify the gap between the current state and the potential opportunities and sharing use cases from across the business.

"We make it less about the AI jargon, but focus on the problems that AI can solve. We have an active discussion with the business leaders saying, 'let's really understand your most impactful problem statements' and we examine if they map to AI capabilities we can deliver. To do this, our team needs a robust understanding of what AI can do, and an ability to engage, understand, and structure problem statements as well as determine the key success factors, how you're going to measure outcomes."
Michael Demissie, Managing Director, AI Hub at BNY Mellon

2. INCREASE THE BUSINESS KNOWLEDGE OF AI EXPERTS

Information can flow in more than one direction, AI teams need to understand front line business problems, in detail, to fully conceptualise the problem their AI capability is needed to solve. Banks foster the business knowledge of AI experts through different avenues, including appointing specific roles who bridge the technical and business side who have a robust understanding of both AI capabilities and SME problems, or embedding Data Science capabilities in business lines. Enabling this expertise can ensure use cases are tied to tangible business problems.

At Société Générale, Data Leaders are embedded within each of the Business and Support units of the Group, tasked with ensuring the unit is using AI & data to serve the unit's strategy, and identifying the value that AI & data could enable.

"We have a point of contact in each of the Business and Support units of the group – the Data Leader – who plays an instrumental role and is in charge of ensuring the unit is committed to using data & AI to serve the strategy of the unit. Leaders identify the pockets of value which could be enabled by data and AI."
Etienne Guibout, Head of Group Innovation Data & AI at Société Générale

Meanwhile, Commonwealth Bank of Australia has recently enabled more agile ways of working across its business. Previously, there was a structure of semi-formal embedding of data scientists within the lines of business: they would go to where the work was to deliver on a project. The bank has now formalised a new model where data scientists are embedded within various business units – to identify and deliver AI use cases – and are supported by a central home in the data science practice, which fosters capability uplift and career growth, including individual contributor tracks. This new model enables data scientists to develop decentralised expertise, while benefiting from centralised capabilities. This type of evolution is often seen in organisations increasing their AI maturity.

"Since we've implemented this formalised model, we are seeing a lot of value already, so the cascading impact of what we're doing has really accelerated quite dramatically over the last 12 months."
Dan Jermyn, Chief Decision Scientist at Commonwealth Bank of Australia

At Capital One data scientists are embedded throughout the business, and work very closely to understand the business line objectives, and help to imagine where the business could be going. The aim is for data scientists to be technical experts as well as domain experts.

"Our data scientists at Capital One are embedded with the business. They are aligned to them. They deeply understand the needs and objectives of the business."
Zach Hanif, VP, Model, Machine Learning, and Software Development Platforms at Capital One

3. BUILD AN AI COMMUNITY

Convening and building an AI community is a powerful tool. Leading banks are creating learning forums, bringing in external expertise from the likes of Microsoft or other major tech players to upskill their AI community. Internal knowledge sharing, such as central use case libraries, can serve as a source of inspiration for future use cases, and help practitioners identify where capabilities may already exist across the business.

Société Générale has an AI/Data and Digital steering committee who meet quarterly to agree where value could be generated from AI, data and digital. The committee also shares the flagship initiatives that can inspire other divisions across the bank, and brings in external speakers on occasion to share the latest news and discuss the latest developments in AI. Additionally, the team created a tool four years ago which has become a directory of all the data and AI use cases at the bank, enabling discovery of use cases and helping employees contact stakeholders involved in developing the use case.

"The community of Digital & Data leaders share their achievements and learnings within the AI/Data and Digital Steering committee who meet quarterly. The committee helps to steer our key results from value generated through AI, data and digital, and is an opportunity for leaders to share their flagship initiatives. Sometimes we invite external speakers to speak to the committee, for instance at the beginning of this year we invited Microsoft to speak with us on the potential for LLMs and their partnership with OpenAI."
Etienne Guibout, Head of Group Innovation Data & AI at Société Générale

At HSBC employees share best practices in specialised messaging channels, bringing together practitioners across the bank to learn from each other. Part of this is encouraging teams to use existing pre-approved models to build on and utilise for use cases.

"In late 2022, HSBC launched a global AI Centre of Excellence, a pro-active community where AI teams collaborate, share ideas and lessons, and interact with AI and Gen AI subject matter experts as products and capabilities are researched, developed and tested."
Suzy White, Global COO, Global Banking & Markets (GBM) at HSBC

THE POWER OF CULTURE AT COMMONWEALTH BANK OF AUSTRALIA

Commonwealth Bank of Australia has a federated structure for AI capabilities. Data scientists are embedded in business units, developing both domain and technical knowledge while working towards a common goal – "building a brighter future for all".

"It's really, really challenging to build a motivated, engaged team of data scientists within a large organisation unless you get the culture right"
– Dan Jermyn, Chief Decision Scientist at Commonwealth Bank of Australia

This structure facilitated the development of Commonwealth Bank of Australia's AI-enabled solution for transaction abuse. The customer advocacy team identified instances where customers were subject to abuse through small transactions in the payment description. Abuse ranged from outright aggressive messages, to more insidious abuse.

The AI team developed a machine learning model to flag harassment transactions for the customer teams to handle sensitively. The model was designed to tackle the nuance of emotional abuse by not only capturing typical abusive language, but more insidious patterns of abuse by analysing sentiment, toxicity, and the relationship of the payer to the payee. Longer term, the model has become a reusable asset that can be used across other use cases, including customer feedback, wider transaction interpretation, and financial crime.

"Having a customer-first mindset that the whole organisation is behind is critical to deploying AI safely and responsibly across the organisation. It's made a whole lot easier because no matter where you sit in the organisation, we all strive to deliver better outcomes for our customers and communities."
– Dan Jermyn, Chief Decision Scientist at Commonwealth Bank of Australia

There remains a simple bottom line: there is no shortage of opportunity. However, there is always a capacity constraint; financial, technical, personnel and time-based. Difficult choices need to be made. That is the next challenge.

KEY PERFORMANCE INDICATORS FOR IDEATING USE CASES

	KPI	DETAIL
Process KPIs	Is there a formal process at the bank to capture ideas for AI use cases?	*A process which enables the sharing of AI use case ideas across teams, supporting ideation of AI use cases from both technical and non-technical employees.*
	Is there a clear approach to provide support to staff members generating ideas for AI use cases?	*Resources allocated to helping generate higher quality ideas, possibly via coaching or internal information resources (such as a use case library).*
Quantitative KPIs	Number of use case ideas generated within a given period	*Regardless of source, assessment of how many ideas are generated across the business. Indicative of the ease for ideas to be sourced from across the business. How does this break down by business unit and geography.*
	Proportion of use case ideas generated by technical and non-technical employees	*Having ideas flowing from non-technical employees is indicative of bringing all employees on the AI journey. Technical employees will also want to feel empowered.*

THE NEED FOR STANDARDISED USE CASE PRIORITISATION

As we have seen, there are plenty of potential AI use cases. As awareness of AI increases, ideas for use cases will continue to proliferate as employees understand its potential to solve problems. Banks need to build a shared and comprehensive way to prioritise these use cases—both at a business unit and at group level—if they are not to drown in the sea of possibilities.

Banks should aim to apply a standardised (and potentially centralised) way of evaluating use cases. This is not only to prioritise where existing use cases can be scaled up but also to prioritise which new ideas that are generated across the business are the right ones to pursue. Using a common framework can help banks realise outcomes faster, and help teams communicate and justify use case prioritisation to senior management and wider stakeholders.

> “The value framework is a good tool to discuss with senior management to explain why we should prioritise use cases, how they are evolving through time, and gather their feedback.”
> *Etienne Guibout, Head of Group Innovation Data & AI at Société Générale*

THREE LENSES BANKS USE TO PRIORITISE USE CASES

While the specific frameworks, decision-making strategies, governance bodies and terminology will vary from bank to bank, broadly there are three lenses through which use cases are evaluated:

1. **Return versus Feasibility:** is it worth it?
2. **Capability Building:** will this make the organisation stronger?
3. **Ethics and Governance:** does this meet rigorous ethical and compliance standards?

FIG 05. THE IDEAL PRIORITISATION FRAMEWORK PROCESS

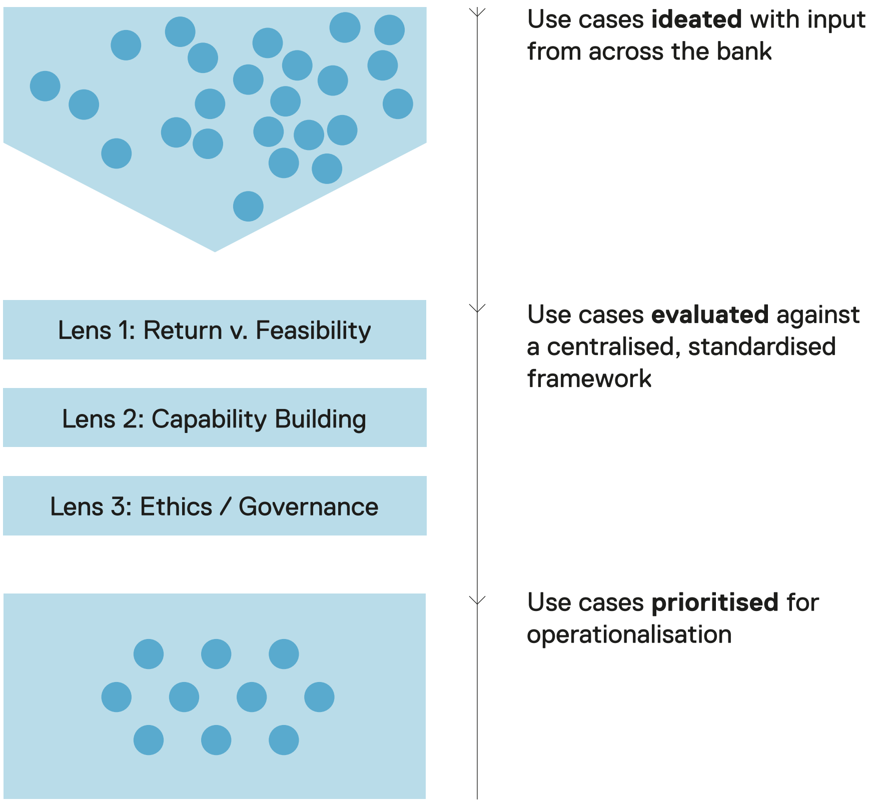

LENS 1. RETURN VS. FEASIBILITY

This lens takes into account two main factors:

→ **Return:** the potential net value of the use case (ideally defined in dollars)
→ **Feasibility:** the challenge of developing the use case (and the risks involved)

> "You have to ask 'how do-able is the project?', and 'how much value can it deliver for customers?'. The things that produce a lot of value are hard to do; the things that produce not that much value are typically easier. There is a sweet spot in the middle. The art of getting this right is to be quite bold and ambitious and then be really focused on ensuring that what we do will deliver better outcomes for our customers."
> *Dan Jermyn, Chief Decision Scientist at Commonwealth Bank of Australia*

Broadly, banks can understand the potential return from a use case as the relationship between the cost of developing it, and the value it could deliver to the business. Chapter 2 includes an in-depth study into calculating the potential value of a use case. We need here to consider approaches to calculating the feasibility (measured in costs and risks) of delivering that value.

FIG 06. HOW TO CALCULATE RETURN FOR AI USE CASES

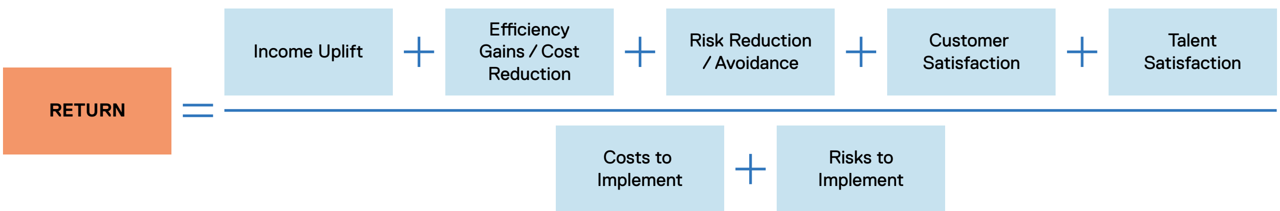

Understanding the feasibility of a use case requires an assessment of the likelihood a use case will fail to be delivered, alongside the effort needed to deliver the use case.

Broadly, when assessing these feasibility factors, banks need to consider:

→ **Effort required:** how much work is needed will be key, as will an understanding of who needs to do it across all functions.

→ **The level of complexity:** complex use cases i.e. use cases which impact multiple business lines and processes, or require integration of multiple AI models or techniques may be less feasible. Regulatory challenges can also be more salient with complex use cases, requiring greater time and effort to ensure standards are met. Determining a use case's complexity can help accurately ascertain how feasible the predicted time requirement may be.

→ **Expertise levels:** banks need to assess what staffing and skills are currently available when assessing the feasibility of AI use cases. Development expertise can be thought of in three buckets, Process SMEs, Data Scientists, and Developers. Data scientists often are thought of first in delivering AI solutions as they train models to augment the solution. However, they often do not have the expertise on the desired output to evaluate the initial training data, provide human in the loop feedback to tune the model, nor evaluate the output to ensure that the answer is accurate and appropriate for a regulated use case. These require managerial and SME input.

Expertise and bandwidth has multiple angles which should all be considered as factors affecting feasibility. It is not only the engineers that matter; but also availability of support function resources around compliance, ethics, and regulatory teams.

→ **The model in use:** powering every AI use case are models. Which models are used for developing the AI use case can influence the feasibility of its development and implementation. Developing new models, with all the associated compute costs for training the new models, is resource and time intensive. If use cases can build off existing models, banks could also benefit twice over from not having to train new models, and benefitting from existing governance or approvals, as novel models require intense regulatory scrutiny before being deployed. Banks may also find using third party models could be a better solution, however this also comes with associated costs, and potential explainability, auditability, or data privacy hurdles to manage.

→ **The infrastructure needed for the use case:** banks need to ensure the building blocks are in place to be able to develop and continue to run internal solutions. An assessment of the existing networking infrastructure, compute power, and storage capacity needed to run the AI use case is required. If current infrastructure capabilities are not sufficient to manage the AI capability, external solutions will need to be explored.

→ **The data availability:** typically, training AI models requires access to significant volumes of data. This is especially true if banks choose to train models from scratch rather than fine-tune existing models. The amount of data required will depend on multiple factors, such as the complexity of the model and what labelling or annotation needs are required. Data scientists need to be able to reliably access the data at scale, and this can be a particular challenge in organisations like banks where data can be siloed or inaccessible.

→ **The data quality / viability:** data quality is a foundational consideration with any AI use case. In many organisations, the level of data quality required to move a use case from development into production cannot just be assumed. It needs to be validated. Additionally, this data needs to be labelled and reviewed, which can be compounded by expertise constraints highlighted previously. Understanding how readily available the data is, in a usable format for data scientists, is essential when assessing a project's feasibility.

→ **The time to production and time to produce outcomes:** throughout feasibility analyses banks need to consider the length of time it will take for a use case to get into production. Other considerations of feasibility will impact this assessment—data and expertise availability or regulatory hurdles to overcome for example.

→ **State of the technology:** it is not unheard of that senior management can aspire to use cases where the technology is simply not yet ready. Managing expectations around this can be a key role for AI Leadership.

LENS 2: CAPABILITY BUILDING

When prioritising use cases, banks should also assess how use cases could accelerate a bank's journey towards AI maturity. Use cases which increase scale of impact through applications across the bank, or pave the way for future use cases, should be prioritised. A key component for pursuing an "AI-first" strategy is the development of continual, self-reinforcing use cases that expand the level of knowledge and capabilities in the bank, and have a cascading impact accelerating use case development and adoption.

> "A lot of the requirements for solving the next problem are related to gathering knowledge, specifically when your data itself becomes knowledge, and then it becomes about solving not just the next problem, but the interconnected inefficiencies in settling trades or trading ideas. If you capture this knowledge, it opens up a whole lot more capabilities to solve that next level problem."
> *Ryan McQueen, Head of Product at DeepSee*

As banks progress along the learning curve, the time taken and resources needed to deliver use cases will decrease as banks accumulate knowledge and capabilities. As banks scale, this ability will continue to accelerate. Once baseline capabilities are established for the initial use cases, banks can iterate through use cases with greater efficiency and speed, compounding on their learning as they progress.

FIG 07. LENS 2 JOURNEY: ACCELERATION OF USE CASE IMPLEMENTATION:

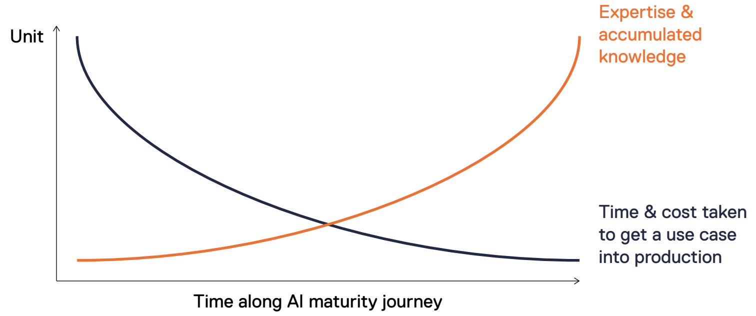

The net impact of this will be to change the calculations about returns and allow a more sophisticated assessment of the relative cost of projects. If a project will help structurally shift the likelihood of future success for AI projects then it may make sense to prioritise it instead of a project with a more immediate return on investment. Part of the challenge of getting this right is making sure that the management teams making prioritisation choices are suitably incentivised on the appropriate time horizon.

> "When making trade-off decisions, what you incentivise is what you get. If the bonus targets are all short term then don't be surprised if necessary, but unsexy, long term requirements are never met."
> *Tim Gordon, Partner at Best Practice AI*

LENS 3: ETHICS/GOVERNANCE REVIEW

Finally when prioritising use cases, banks need to consider the ethical or governance ramifications of the use case. Banks are held to stringent regulatory standards, and have a duty to deliver trusted services to their customers. The transformative power of AI brings this issue even more to the forefront when prioritising use cases. The technology has broad-reaching capabilities, and its potential risks have been debated, often from before their practical manifestation. While a proposed use case may be able to deliver significant value for the firm, if this comes at a cost on issues such as customer privacy, or opens up the bank to regulatory penalties, then the use case needs to be deprioritised.

Take a use case which would be used for assessing credit scores: while on the face of it, the use case could be used to increase efficiencies, the risk of creating models which could create biased outcomes against certain groups, could mean the use case gets reweighted or blocked through the ethics review. In contrast, a use case which would drive greater financial inclusion, yet on the face of it may not drive significant profit margins or a similar return measure, could be boosted on the prioritisation framework as it can be presented as the "right thing to do" with potential brand and reputational impact.

These three lenses have clear overlaps. In practice banks will have many different ways to cover these issues – and varied terminology to describe them. However a clear process is key to ensure that banks are driving in the right direction. Once a decision has been made on which use cases to go after, banks face the next challenge: getting a use case out into production.

The following KPIs are a guide to how banks can assess their level of maturity when it comes to the effectiveness of their prioritisation process.

KEY PERFORMANCE INDICATORS FOR PRIORITISING USE CASES

	KPI	DETAIL
Process KPIs	Does the bank have a common (centralised) evaluation framework to prioritise AI use cases?	*Common, centralised framework should include a consideration of the AI use cases' impact v. feasibility, whether it builds capabilities, and an ethics / governance evaluation, and be implemented consistently across the organisation.*
Quantitative KPIs	Proportion of use case ideas that are approved for POC development	*This will measure quality of ideation output, delivery capacity and quality of management process around then prioritisation process. A comparison with the next KPI may suggest a weak or indecisive decision-making process.*
	Time taken for a use case to get approved for POC development	*Assessment of the speed and effectiveness of corporate take up of new ideas. Speed to market will imply a more mature and competitive AI delivery capability, as well as potentially higher capacity.*

STEPS TO OPERATIONALISATION

Once banks have selected which use cases to put into production, the next stage is operationalisation. Only as use cases come into production can banks unlock the value of AI, and incrementally build on success throughout the organisation.

> “It’s not real unless it’s live.”
> *Ash Booth, Global Lead Applied AI/ML in Markets at JPMorgan Chase*

The specific steps involved in operationalising an AI use case will depend on multiple factors including existing capabilities, architecture, and the nature of the use case itself. If banks can streamline the operationalisation process, they can expedite the output of use cases into production, driving scalability and impact along the AI maturity journey.

Broadly speaking, the process of operationalising an AI use case involves:

→ Data collection, labelling, and cleaning
→ Training and Development of Models
→ Deployment in a testing environment (e.g. Proof of Concept in the lab setting)
→ Monitoring outcomes, and adjusting models where appropriate
→ Model Validation and Governance
→ Full deployment to teams, including upskilling on the use case’s capabilities and changes to the surrounding operational architecture

Reducing the time taken to iterate through these steps, without compromising on quality or integrity, will help banks harness the competitive advantage AI can bring. However, it is a challenge, and requires a mindset shift from organisations which historically can be less than agile in the face of change.

OPERATIONAL ENABLERS

In our recent interviews, senior AI leaders across the banks highlighted four "top-of-mind" enablers in order to operationalise AI at scale.

4 KEY ENABLERS TO BUILD AN AI-CENTRIC OPERATING MODEL:

1. Invest in AI tools and platforms
2. Create a fit-for-purpose model validation framework / governance
3. Build an enabling data and infrastructure strategy
4. Form appropriate partnerships to fill capacity gaps and accelerate delivery

1. INVEST IN AI TOOLS AND PLATFORMS

In order to accelerate the process of moving a use case from ideation to production, and scale the operationalisation of AI use cases across the banks, many of the banks we've spoken to are investing in common platforms and tooling that centralise foundational AI capabilities, and make these accessible to teams across the bank.

CAPITAL ONE'S FOUNDATIONAL PLATFORMS BEST PRACTICE:

When it comes to democratising AI and ML across the enterprise, Capital One prioritises foundational, common operating platforms. This enhances accessibility, which aims to make data and applications intelligible and available to users and developers across the business, as well as to more nimbly and effectively prototype ideas, develop, refine, and validate through the model development life cycle. Part of this democratisation means learnings from different teams can be shared, so the same problem is not solved in two different areas of the organisation.

"These types of platforms allow for developers to move through the entirety of the model development life cycle in a well-managed way. If you have an idea, you start prototyping it here, you develop it, you refine it, you validate it, and you verify that it is working the way you want it to. They ultimately provide consistent environments for designing, deploying and managing ML models, repeatedly and in larger volumes."
– Zach Hanif, VP, Model, Machine Learning, and Software Development Platforms at Capital One

Crucially, however, these kinds of platforms ensure the time from idea to testing is shortened – freeing up more of associates' time to spend on experimenting and testing. For example, Capital One's fraud efforts have benefitted from this speed – in a complex space which processes a high volume of transactions, even the smallest margins of improvement can drive high value for the fraud teams. Having the ability to accelerate iterations through model development enabled fraud teams to rapidly adapt models to the changing environment of fraudsters, delivering value for customers and the business.

> "We're not only accelerating the data science effort from idea to production, but also building up capabilities, platforms and infrastructure, everything you need to deliver AI solutions at scale in a responsible way."
> *Michael Demissie, Managing Director, AI Hub at BNY Mellon*

Centralised model repositories enable AI practitioners to access existing models that could then be deployed across different areas of the bank. For example, access to a term extraction model might serve ESG analysis functions in a trading division, as well as payments settling or automated savings guidance functions for customers.

The leading banks are going one step further and providing a single, common operating environment which reduces the time for data scientists to set up the environment, or workbench, needed to experiment with and train AI models.

> "Developing foundational AI and machine learning platforms enables teams to achieve several critical things to promote access to this kind of technology. It gets associates on the same tech stack; facilitates collaboration, reusable components, and standardisation; and it helps bring down silos. It also enables the company to have an idea, prototype, experiment and validate — going from experiment to efficacy through insights and data."
>
> *Zach Hanif, VP, Model, Machine Learning, and Software Development Platforms at Capital One*

2. CREATE A FIT-FOR-PURPOSE MODEL VALIDATION FRAMEWORK / GOVERNANCE STRUCTURE

The banking sector is subject to many regulatory demands, and for a disruptive technology like AI, maintaining customer protection and regulatory standards is of the utmost importance. Tensions between wanting to "go fast" and "go well" mean use cases can be stalled upon reaching the point of model validation – where a use case is assessed to determine whether it conforms to governance requirements and ethical frameworks.

> "It's hard to get from POC to production because model validation is hugely important in financial services."
>
> *David Rice, Global COO, Commercial Banking (CMB) at HSBC*

Our conversations suggest that banks experience these governance issues very differently. Commonwealth Bank of Australia, for example, emphasised their strength in enabling innovation within existing model governance frameworks. Executives at less mature banks, especially those with geographic complexity built in, were more keen to suggest that a new, more dynamic approach to governance might be needed if the banks were to become truly AI-first. The propagation of differing regulatory and legislative approaches to AI across the world can only complicate these challenges.

With the proliferation of Generative AI, the need for dynamic model governance frameworks and regulation has become even more apparent. Having common data frameworks and centralised knowledge repositories can help with navigating the governance landscape. Building a well informed and engaged support team ecosystem can also expedite the governance process, helping to identify any governance sticking points early on in the process.

> "With Generative AI, we debated whether it needs distinct governance. Ultimately, we found that while general principles apply, there are added considerations and risks. Therefore, we are putting in place additional assessments and strong guardrails before a GenAI model can be deployed."
>
> *Monique Shinavandan, Chief Data & Analytics Officer at HSBC*

Key decisions early on are also likely to prove wise investments. Training staff across the organisation, especially on Responsible AI approaches and embedding these in design approaches from the start will ensure that the AI Governance teams can focus on support and encouragement rather than simply policing.

> "The earlier banks engage with Responsible AI the cheaper it will be deliver it."
> *Tim Gordon, Partner at Best Practice AI*

3. BUILD AN ENABLING DATA & INFRASTRUCTURE STRATEGY

Data is the fuel that drives AI, and most banks are years into decades-long processes to improve data infrastructure and shift data onto the cloud. However, the rise of Generative AI has placed increased demands on data availability to meet the pace of AI development. Models need to be trained on high quality data, which meets the rigorous regulatory requirements of the financial sector – otherwise, as the adage goes, "garbage in, garbage out".

> "One of our biggest challenges is creating a data asset that is fit for purpose to enable the algorithmic and analytical capability that is available to us."
> *– David Rice, Global COO, Commercial Banking (CMB) at HSBC*

While the need for investment in data and infrastructure isn't new, the reality is that banks are still on the journey towards data maturity, and it remains top of mind for banks looking to increase their pace and scale of AI deployment.

> "If you start at the very top with good data governance aligned with good data management capabilities, then you'll have good data quality as an outcome that is going to help with how AI takes in that information, learns from it, and puts it together in meaningful outputs. This is because the data is based on a set of reasonable standards that enable consistency in terms of what the recommendations are that AI generates, resulting in usable and useful metadata and data content. You have the foundation set."
> *Thomas Dunlap, Founder & Managing Partner at DIACSUS*

An AI-enabled data strategy needs to prioritise the following:

→ **Quality data sets:** when Data Scientists have to grapple with the arduous task of wrangling data, cleaning and tagging datasets on a large scale, banks are diverting a valuable resource which could be better allocated towards model optimisation. In addition, banks will increasingly access external data sets, potentially for research purposes or simply to help build better quality models or to deliver new use cases. The rules for governance, privacy, quality control and security will be just as important as for internal data – and building an eco-system to deliver them is an added challenge.

→ **Scalable access to data:** banks' legacy infrastructure can impede data strategy. Where data is not in a shared data lake or equivalent, AI models may struggle to access the volume of data needed to train and operate AI models. Where banks have not fully digitised business processes it can be a challenge to create the necessary data sets for AI models to be trained on. Democratising access to quality data across business lines will drive innovative use cases, and enable their operationalisation.

→ **Embedded privacy and governance frameworks:** protecting data privacy and security when building AI models is critical for maintaining trust in the technology, and ensuring outputs are compliant with privacy regulation. Every data strategy needs to have this at its heart. If this is not implemented, banks face the longer term risk of eroding customer trust in financial institutions, as well as opening themselves up for regulatory and compliance penalties.

So how can banks get there? A strong data strategy requires three key components to drive success:

→ **Good collaboration models:** everyone, from enterprise leadership to practitioners, needs to be involved and to support the data strategy;

→ **Strong policies:** having well positioned, robust data governance policies and standards;

→ **Application of the capabilities:** the most challenging part of a data strategy is ensuring that the policies and guidelines are consistently adopted across the business.

Getting data architecture, pipelines, and APIs in order so that you can deploy AI at scale, and bringing this within the same governance framework, with aligned security measures, will ensure that banks can fully capitalise on the AI opportunity.

4. BUILD APPROPRIATE PARTNERSHIPS

The fourth enabler that repeatedly crops up in conversations with AI leaders is establishing partnerships to scale up AI deployment, fast.

Buying or partnering can help provide access to missing resources or tools and speed up delivery. External firms may have built up specialist skill sets, overwhelming data advantage or simply have become the industry's de facto standard. Relationships between vendors and banks can vary from simple "off the shelf" purchasing to deeper partnerships with multiple levels of exchange and cross-fertilisation of data, expertise and even new revenue streams from commercial sales to 3rd parties.

Partnerships can prove to be particularly valuable when a particular use case or technology is earlier in the maturity cycle. At the beginning of the cycle, banks are less likely to have the skills in-house to develop the technology, so they may be more likely to turn to providers whilst developing internal capacity.

> "In the beginning of the hype curve, the third parties, your partners, are critically important as you get started in your journey."
> *Steve Van Wyk, Global CIO at HSBC*

Over time, banks may have had the time and resources to build internal teams to grapple with the use case and build capabilities. This is not a set rule, as use cases progress along different journeys, and banks also need to consider other questions, including how much control they would want to maintain over the technology.

Generative AI has brought the trade-offs of partnerships into sharp relief: while organisations can lean on third party providers like OpenAI, questions have been asked about whether these capabilities might be brought in house. Questions of data risks when using the technology, as well as its explainability, need to be considered when deciding on the appropriate approach. With regulatory guidance like the Federal Reserve's final guidance on third-party risk management placing the majority of risk management responsibility on banks, banks may be more comfortable building internal capabilities to be closer to model development.

However, erring too much on either building or buying AI capabilities can decrease resilience in the long term.

If too many solutions are built, technologies are limited to in-house resources and expertise. Banks need to be strategic in prioritising internal resources for building use cases that will create a clear competitive advantage versus commodity applications which are less differentiated, for example, marketing spend optimization.

If too much is bought, developing a competitive capability will be a lot more challenging, and banks will be subject to the risk of rising vendor costs in the longer term.

There is no set answer for what a potential procurement strategy should be. We are seeing banks both significantly investing in internal capabilities, as seen with JPMorgan Chase's extensive AI research team, as well as buying AI solutions.

> "If someone else could build it better than we could, we can leverage that."
> *Jeff McMillan, Chief Analytics and Data Officer at Morgan Stanley*

If buying, choosing the right vendor is essential for this journey: vendors are adopting a consultative position as experts within the space, engaging with banks to help diagnose internal issues as well as find solutions. Vendors are also well placed to provide (hopefully independent) guidance for banks on where their current strategy and use cases approach is in comparison to peers. Accessing the knowledge and network of the right strategic vendor can provide valuable additional resources and insights into the development of AI use cases.

> "They benchmark you both in terms of where you're at and what you're thinking, and your ambition."
> *David Rice, Global COO, Commercial Banking (CMB) at HSBC*

Banks need to carefully evaluate their overall mix of built, bought, and partnered solutions. The profile of this mix will change over time, as banks expand their ecosystem and develop internal capabilities, however the more resilient banks will embrace variety.

The following KPIs are a guide to how banks can assess their level of maturity when it comes to the effectiveness of their operationalisation process.

KEY PERFORMANCE INDICATORS FOR OPERATIONALISING USE CASES

	KPI	DETAIL
Process KPIs	Does the bank have a centralised platform for developing AI use cases which employees across the bank can access?	*A standardised platform which gives delivery teams democratised access to capabilities which help to build, test and refine AI models easily.*
Quantitative KPIs	Proportion of approved use case ideas that end up in POC	*Likelihood that the operational enablers (tooling, governance framework, data strategy) are aligning to drive AI use cases through to POC.*
	Proportion of approved use case ideas that end up in deployment	*Likelihood that the operational enablers (tooling, governance framework, data strategy, partnerships) are aligning to drive AI use cases through to deployment.*
	Time taken for an AI use case to move from approval to POC	*Time taken to go through the stages involved in creating a POC, including identifying data sources, labelling relevant data, and training the AI model.*
	Time taken for an AI use case to move from POC to deployment	*Length of time needed for AI use cases to reconcile governance requests and scaling for deployment in a live environment.*
	Proportion of use cases that meet initial cost expectations	*How good was the prioritisation calculation process.*
	Proportion of use cases that meet initial time expectations	*How good was the prioritisation calculation process.*
	Proportion of uptake among target users of the AI use case	*How well received is the model in the target bank employee population. This measure will vary by use case.*

ENDNOTE

This report, based on interviews and conversations with senior AI leaders, explored the different approaches banks are taking towards mapping, measuring, ideating, prioritising and operationalising AI within banks. While best practice is still emerging, we aim to provide a common framework to enable banks to evaluate and compare the progress they are making towards delivering AI outcomes vs peers.

Please reach out to find out more about how to get involved in the outcomes benchmark, our Membership offering, and how we're creating the definitive independent benchmark for tracking industry-wide AI adoption and readiness. The next Evident AI Index will be released in November 2023, expanding the ranking to include 50 banks.

Alexandra Mousavizadeh
CEO & Co-founder
Alexandra.mousavizadeh@evidentinsights.com

Annabel Ayles
COO & Co-founder
Annabel.ayles@evidentinsights.com