



Overseeing AI: Governing artificial intelligence in banking

Ensuring ethical, fair and well-documented AI-based decisions will gain urgency in the post-pandemic era. A review of global regulatory guidance given so far reveals the key risks and recommendations.

- AI will separate the winning banks from the losers, 77% of executives in the industry agree
- Covid-19 may intensify the use of AI, making effective governance all the more urgent
- A review of regulatory guidance reveals significant concerns including data bias, “black box” risk and a lack of human oversight
- Guidance has so far been “light touch” but firmer rules may be required as the use of AI intensifies

The ability to extract value from artificial intelligence (AI) will sort the winners from the losers in banking, according to 77% of bank executives surveyed by The Economist Intelligence Unit in February and March 2020. AI platforms were the second highest priority area of technology investment, the survey found, behind only cybersecurity.

At the time, the covid-19 pandemic was in full effect in Asia and the rest of the world was beginning to understand its gravity. Since then,

the depth and extent of covid-19's impact on consumer behaviour and the global economy have come into clearer focus.

Covid-19 has already triggered an uptick in digital banking—in the US Citibank is reported to have seen a tenfold surge in activity on Apple Pay during lockdown, for example.¹ But the disruption to businesses and households has only just begun, and banks will need to adapt to rapidly changing customer needs.

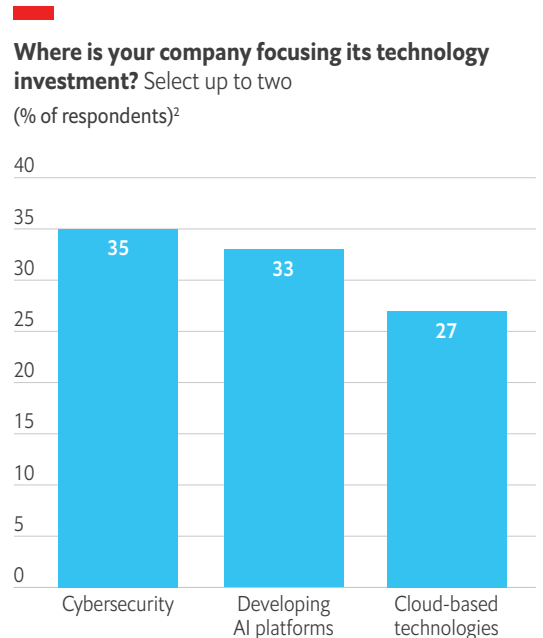
As such, the criticality of AI adoption is only likely to increase in the post-pandemic era: its safe and ethical deployment is now more urgent than ever.

In common with most matters of governance and safety, banks will look to regulatory authorities for guidance on how this can be achieved. Until a few years ago regulators adopted a wait-and-see approach but, more recently, many have issued a number of studies and discussion papers.

In order to assess the key risks and governance approaches that banking executives must understand, The Economist Intelligence Unit undertook a structured review of 25 reports, discussion papers and articles on the topic of managing AI risks in banking. The main findings are summarised in the table on page 3 and examined throughout this article.

Risks known and unknown

The nature of the risks involved in banks' use of AI does not differ materially from those faced in other industries. It is the outcomes that differ should risks materialise: financial damage could be caused to consumers, financial institutions



Source: The Economist Intelligence Unit.

themselves or even to the stability of the global financial system.

Our review reveals that prominent risks include bias in the data that is fed into AI systems. This could result in decisions that unfairly disadvantage individuals or groups of people (for example through discriminatory

“Some AI models have a complexity that many organisations, including banks, have never seen before.”

Prag Sharma, senior vice-president and emerging technology lead, Citi Innovation Labs

¹ Antoine Gara, “The World’s Best Banks: The Future Of Banking Is Digital After Coronavirus”, *Forbes*, June 8th 2020. <https://www.forbes.com/sites/antoinegara/2020/06/08/the-worlds-best-banks-the-future-of-banking-is-digital-after-coronavirus/>

² Top three responses shown. For full results “Forging new frontiers: advanced technologies will revolutionise banking”, The Economist Intelligence Unit, 2020. <https://eiperspectives.economist.com/financial-services/forging-new-frontiers-advanced-technologies-will-revolutionise-banking>

Bank governance of AI: main issues raised by regulators

Governance area	Challenges	Regulator guidance ³
Ethics	<ul style="list-style-type: none"> AI systems should be developed, deployed and used in ways that respect human autonomy, prevent harm and ensure fairness and explicability 	<ul style="list-style-type: none"> Apply an “ethical by design” approach to AI model development
Fairness	<ul style="list-style-type: none"> Bank decisions based on AI models should not disadvantage any individual or groups of individuals without justification To guard against discrimination—eg in loan approvals—AI models should be free from bias 	<ul style="list-style-type: none"> Establish an ethics committee to validate AI use cases and monitor their adherence to ethical standards Institute the periodic review of AI decisions from a fairness perspective by ethics or other committees or by individual experts Data sources should be evaluated and tested regularly to ensure data is representative
Explainability and traceability	<ul style="list-style-type: none"> Banks must be able to fully explain any AI-driven decision that affects customers or other individuals who provide data The steps leading to a decision should be able to be tracked from initial data gathering through to the actual decision 	<ul style="list-style-type: none"> All steps taken in the development of an AI model should be thoroughly documented Employ “feature ranking” by which the relative importance of each input variable to different types of decisions is calculated
Human in the loop	<ul style="list-style-type: none"> Human monitoring and review are needed to ensure AI models are performing correctly, particularly in situations where fairness of decisions affecting bank customers may come into question 	<ul style="list-style-type: none"> Mandate regulator monitoring and the review of model functioning by domain experts Set boundaries to model outcomes, beyond which further decisions are delegated to a domain expert
Data quality	<ul style="list-style-type: none"> Avoiding bias and ensuring fairness of decisions depends on the accuracy and integrity of data used in AI models 	<ul style="list-style-type: none"> Sound bank-wide data governance practices should be sufficient for data used in AI models Such practices should also be applied to data used in model training, testing and validation, and to third-party data used in models
Model security	<ul style="list-style-type: none"> AI models are susceptible to unique threats, such as stealing (a model’s illicit replication) and poisoning (in which anomalies are introduced to model data) 	<ul style="list-style-type: none"> Robust bank-wide cybersecurity defences should be adequate to protect AI models
Skills	<ul style="list-style-type: none"> Deficits of AI-specific talent and expertise pose a challenge not just in terms of building and maintaining AI models, but also for the models’ oversight, control and audit 	<ul style="list-style-type: none"> Banks must ensure that senior management, risk management and compliance functions have an adequate level of AI expertise Expertise may be recruited from outside, but internal training in the workings of AI models should be provided to existing risk management and audit staff Banks should also ensure their boards of directors are familiarised with the functioning of AI models
Compliance	<ul style="list-style-type: none"> To train their AI models banks need to feed them large amounts of customer data; the management of this data must fully conform to privacy regulations (such as the EU’s GDPR) 	<ul style="list-style-type: none"> Apply a “compliance by design” approach to AI model development
Governance structure and management	<ul style="list-style-type: none"> The use of AI models must be subject to oversight and governance practices that are no less stringent than those for any other area of bank operations 	<ul style="list-style-type: none"> The use of AI should be integrated into banks’ existing risk management frameworks Clear roles and responsibilities for the adoption and use of AI should be assigned to individuals or bodies within the existing governance structure
Accountability	<ul style="list-style-type: none"> Responsibility for AI decisions must be delegated to humans and not machines Such responsibility must ultimately be held by banks’ boards of directors 	<ul style="list-style-type: none"> Banks should apply their AI standards and requirements consistently to internally developed and externally sourced AI applications Boards’ accountability should extend to externally developed and/or sourced AI applications

³ The Economist Intelligence Unit reviewed 25 reports, discussion papers and articles published in the past three years by banking and financial sector supervisory authorities, central banks and supranational institutions, as well a handful of reports from universities and consultancies, on the theme of managing the risks associated with AI. The points of guidance listed below are distilled from studies and discussion papers published in the past two years by financial sector supervisory authorities, regulatory agencies and supranational institutions. These include the European Banking Authority, Monetary Authority of Singapore, DNB (Bank of the Netherlands), Hong Kong Monetary Authority, CSSF (Financial Sector Supervisory Authority of Luxembourg), BaFin (Federal Financial Supervisory Authority of Germany), ACPR (Prudential Supervision and Resolution Authority of France), and the European Commission.

lending). “Black box” risk arises when the steps algorithms take cannot be traced and the decisions they reach cannot be explained. Excluding humans from processes involving AI weakens their monitoring and could threaten the integrity of models (see table for a comprehensive list of AI risks in banking).

At the root of these and other risks is AI’s increasingly inherent complexity, says Prag Sharma, senior vice-president and emerging technology lead at Citi Innovation Labs. “Some AI models can look at millions or sometimes billions of parameters to reach a decision,” he says. “Such models have a complexity that many organisations, including banks, have never seen before.” Andreas Papaetis, a policy expert with the European Banking Authority (EBA), believes this complexity—and especially the obstacles it poses to explainability—are among the chief constraints on European banks’ use of AI to date.

Governance challenges

The guidance that regulators have offered so far can be described as “light touch”, taking the form of information and recommendations rather than rules or standards. One possible reason for this is to avoid stifling innovation. “The guidance from governing bodies where we operate continues to encourage innovation and growth in this sector,” says Mr Sharma. Another reason is uncertainty about how AI

will evolve. “AI is still at an early stage in banking and is likely to grow,” says Mr Papaetis. “There isn’t anyone who can answer everything about it now.”

The documents that banking regulators have published on AI range from the succinct (an 11-page statement of principles by MAS—the Monetary Authority of Singapore) to the voluminous (a 195-page report by Germany’s BaFin—its Federal Financial Supervisory Authority), but the guidance they offer converges in several areas.

At the highest level, banks should establish ethical standards for their use of AI and systematically ensure that their models comply. The EBA suggests using an “ethical by design” approach to embed these principles in AI projects from the start. It also recommends establishing an ethics committee to validate AI use cases and monitor their adherence to ethical standards.⁴

For regulators, paramount among the ethical standards must be fairness—ensuring that decisions in lending and other areas do not unjustly discriminate against individuals or specific groups of people. De Nederlandsche Bank (or DNB, the central bank of the Netherlands) emphasises the need for regular reviews of AI model decisions by domain experts (the “human in the loop”) to help guard against unintentional bias.⁵ The Hong Kong Monetary Authority (HKMA) advises that

⁴ European Banking Authority, *EBA Report on Big Data and Advanced Analytics*, January 2020. https://eba.europa.eu/file/609786/download?token=Mwkt_BzI

⁵ De Nederlandsche Bank, *General principles for the use of Artificial Intelligence in the financial sector*, July 2019. https://www.dnb.nl/en/binaries/General%20principles%20for%20the%20use%20of%20Artificial%20Intelligence%20in%20the%20financial%20sector2_tcm47-385055.pdf

⁶ Hong Kong Monetary Authority, *Reshaping Banking with Artificial Intelligence*, December 2019. https://www.hkma.gov.hk/media/eng/doc/key-functions/financial-infrastructure/Whitepaper_on_AI.pdf

Monitoring the modellers at Citi

Bias can creep into AI models in any industry but banks are better positioned than most types of organisations to combat it, believes Prag Sharma, senior vice-president and emerging technology lead at Citi Innovation Labs. “Banks have very robust processes in place, learned over time, that meet strict external [regulatory] and internal compliance requirements,” he says.

At Citi, a model risk management committee reports directly into the bank’s chief risk officer and operates separately from the modellers and data science teams. Consisting of risk experts as well as data scientists, the committee’s task, says Mr Sharma, is to scrutinise the models

that his team and others are developing exclusively from a risk perspective. The committee’s existence long predates the emergence of AI, he says, but the latter has added a challenging new dimension to the committee’s work.

Maximising algorithms’ explainability helps to reduce bias. Once a team is ready to deploy a model into production, the model risk management committee studies it closely with explainability one of its key areas of focus. “[The risk managers] instruct us to explain all the model’s workings to them in a way that they will completely understand,” says Mr Sharma. “Nothing hits our production systems without a green light from this committee.”

model data be tested and evaluated regularly, including with the use of bias-detection software.⁶ “If you have a good understanding of your underlying data,” says Mr Sharma, “then a lot of the algorithmic difficulties in terms of ethical behaviour or explainability can be addressed more easily”.

Ensuring the right level of explainability, as Mr Papaetis suggests, is arguably banks’ toughest AI challenge. Most of the regulator guidance stresses the need for thorough documentation of all the steps taken in model design. The

Ensuring the right level of explainability is arguably banks’ toughest AI challenge.

EBA, says Mr Papaetis, recommends taking a “risk-based approach” in which different levels of explainability are required depending on the impact of each AI application—more, for example, for activities that directly impact customers and less for low-risk internal activities.

Are more prescriptive approaches needed?

Regulators generally consider banks’ existing governance regimes to be adequate to address the issues raised by AI. Rather than creating new AI-specific regimes, most regulators agree that current efforts should instead be focused on updating their governance practices and structures to reflect the challenges posed by AI. Ensuring that the individuals responsible for oversight have adequate AI expertise is integral, according to DNB.

Fit for purpose? Applying existing governance principles to AI

In Europe, bank adoption of AI-based systems may be described as broad but shallow. In a recently published study the European Banking Authority (EBA) found that about two-thirds of the 60 largest EU banks had begun deploying AI but in a limited fashion and “with a focus on predictive analytics that rely on simple models”.⁷ This is one reason why Andreas Papaetis, a policy expert with the EBA, believes it is too early to consider developing new rules of governance for the EU’s banks that focus on AI.

Mr Papaetis points out that the EBA’s existing guidelines on internal governance and ICT (information and communications technology) risk management are adequate for AI’s current level of development in banking. The existing framework is sufficiently flexible

and not overly prescriptive, he says. “That enables us to adapt and capture new activities or services as they develop. So at the moment we don’t think that AI or machine learning require anything additional when it comes to governance.” In any event, the EBA’s approach, he says, will be guided by European Commission policy positions on the role of AI in the economy and society.⁸

Mr Papaetis does not exclude the possibility of more detailed regulatory guidance on AI in areas such as data management and ethics. Should bias and a lack of explainability prove to be persistent problems, for example, regulators may need to consider drafting more specific rules. But at present, he says, “no one can predict what challenges AI will throw up in the future”.



No one can predict what challenges AI will throw up in the future.

*Andreas Papaetis, policy expert,
European Banking Authority*

⁷ *ibid.*

⁸ The commission launched a public debate on AI policy options in early 2020 and is expected to announce its policy decisions later this year.

More regulatory guidance will almost certainly be needed in the future, says Mr Sharma, and some of it may require the drafting of rules. He offers as an example the uncertainty among experts around the retraining of existing AI models. “Does retraining a model lead to the same or a new model from a risk perspective?” he asks. “Does a model need to go through the same risk review process each time it is retrained, even if that happens weekly, or is a lighter-touch approach possible and appropriate?”

Should a major failure be attributed to AI—such as significant financial losses suffered

by a group of customers due to algorithm bias, evidence of systematic discrimination in credit decisions or algorithm-induced errors that threaten a bank’s stability—authorities would no doubt revisit their previous guidance and possibly put regulatory teeth into their non-binding recommendations.

The need to monitor and mitigate such risks effectively makes it incumbent on regulators to build AI expertise that meets or betters that of commercial banks. As AI evolves, strong governance will inevitably be demanded at all levels of the banking ecosystem.

While every effort has been taken to verify the accuracy of this information, The Economist Intelligence Unit Ltd. cannot accept any responsibility or liability for reliance by any person on this report or any of the information, opinions or conclusions set out in this report. The findings and views expressed in the report do not necessarily reflect the views of the sponsor.