Deep Science Publishing
https://doi.org/10.70593/978-93-49910-62-1 _7

DeepScience
Open Access Books

Chapter 7: Balancing innovation with compliance: Governance, risk management, and security in artificial intelligence-augmented banking infrastructure

7.1. Introduction

As FinTechs have been leveraging technology to provide more efficient ways to deliver banking services to their customers, traditional banks have been delegating their risk infrastructure responsibilities to others or building technological capabilities to innovate their services. While innovative technology aims to lower the cost to develop a base product, a novelty is expected in personalized service delivery or decision-making, availability, and multi-functionality in a single applet. The use of Artificial Intelligence tools to facilitate those techno-commercial characteristics is becoming more and more widespread. However, as a critical piece of the technology stack made available and accessible as a Service, AI in its various forms is also contributing to a shadow banking ecosystem of banks and financial institutions creating challenges for the conduct of monetary policy and the regulation of economy-wide financial stability risks. This has led banks to explore the use of AI to be innovative in their service invention delivery (Brynjolfsson & McAfee, 2014; Iansiti & Lakhani, 2020; Marr, 2020).

Aging core banking infrastructure setups that have become monolithic and complex, and that cost resources to change, and a demand for novelty, personalization, and seamlessness, are driving this trend. However, financial services are a highly regulated space for good reasons, often using the conduct of monetary policy transmission as an argument. These reasons give rise to a pushback against the use of novel AI-powered tools in infrastructure operations. Regulatory authorities are grappling with managing these challenges to balance innovation with compliance in banking infrastructure powered by AI. Being a powerful tool, the shadow-side of AI also brings into place serious regulatory concerns. It can amplify human errors, biases, and prejudices, and move towards deskilling bank and operator personnel managing risks. The ability of large language models to fabricate responses and make errors with authority calls into question their explainability, accountability, and auditability.

However, use of open architecture and clouds create security weaknesses and require renewed vigilance on the part of banks and oversight agencies designed to protect customers against fraud and breaches of trust. In parallel, the collection and use of customer data to personalize and predict their needs intensify the obligations of banks under laws that mandate responsible management of personally identifiable information. These potential impediments to globalization and the exercise of customer-centricity have heightened the demand for the development of solution frameworks that reconcile innovation with regulatory compliance (Ross, 2016; Tapscott & Tapscott, 2016).

Fig 7.1: Balancing Innovation with Compliance

7.1.1. Background and Significance

The rapid metamorphosis of the financial sector was accelerated by the digital revolution during the pandemic, which highlighted inequities in financial services. In addition, increasing competitive pressure has compelled banks and financial services providers to innovate and differentiate themselves, but the demands of shareholders, regulators, and customers have placed checks on the unfettered exercise of innovation. These forces have come together to herald the advent of AI-augmented financial services. New services will be delivered by traditional banks in collaboration with fintechs that specialize in niche areas. Behind the scenes, new and innovative architectures that embrace an API-led paradigm are emerging and are designed to meet the unique and evolving needs of particular markets, customer segments, and products.

AI will also help to radically re-engineer back-office processes and infrastructures to create faster, more robust, and predictable systems. Cloud computing and cloud-native services will extend the range of innovation and options available to banks and financial services providers. By embedding AI capabilities in the devices and tools that companies use for day-to-day transactions, these processes will require less human intervention, removing silos related to internal geographies and resources, and will expose the sensor data needed to validate predictions and support decision-making at the point of interaction with customers. Open architecture and shared data will create intelligent interdependence that encourages transparent and predictable behavior among stakeholders.

7.2. Understanding AI in Banking

Banking as a sector is risk-averse and cautious regarding new trends, even trends with potential to disrupt banking business. AI as a concept started being studied from the middle of the last century, but it's coming to practice came at the beginning of the present century. Already in 2003 some researchers were suggesting AI applications which could achieve an improvement in the risk management of banking and finance. More specifically, they stated that in domains such as credit reporting/enforcement and financial fraud detection, the AI branch of anomaly detection is being successfully used in different forms, generating significant savings. A decade after these initial applications, the set of AI applications is significantly wider. More than just suggesting and predicting banking segments, or monitoring suspicious behavior, nowadays AI can support the development of financial regulation, or support regulators decisions, can indirectly control excess credit during bubble periods, and can even support clients in risk decisions. Not to mention the full range of broadened fraud detection and risk group identification, including the historical AI applications.

In addition to that, despite the AI revolution, there are new traditional banking procedures that require AI tools to analyze their advantages. For example, internal structure applications, such as talent management, are augmenting the use of traditional models to use AI tools to assess the use of diverse traditional algorithms, comparing their resolution to the natural intelligence model. This way, the potential improvements being reached by financial institutions are not only in the traditional banking activity of credit risk management, which have been advanced by AI/ML usage, but actually also in many other areas covered by traditional banking activities.

7.2.1. Historical Context of AI in Banking

The synthesis of the convergence of finance and technology was founded in the late 1950's. The computer acceptance corporation dialogs about the possibilities of electronic data processing. Computers would allow banks to manipulate large volumes of data. With the first banking radial, the Computer has broadcast the concept of universal banking. The banks of the next century would offer more than simple check-cashing services: They would know their customer's saving and investment needs better than they did. They would mainly use the customers financial data and the technology available to develop sophisticated investments. The revolution in banking practice which the Computer foresaw was not that of Automated Teller Machine but of the data communication revolution circa 1980. Nevertheless at that time, some were already following a different course. They would Automation, banks like Chemical Bank and Wells Fargo would replace teller operations with ATMs, concentrating serving services on what they could do well.

Using machine processing and plentiful data resources, they would apply credit marketing techniques, loan officers and printing operations they would combine management to direct many small loans and what were then considered foreign markets: Consumers lacking adequate banking facilities. The future for these small banks was not bright. Their banking was static; Market potential depressed by limitations: High interest rates stifling demand for loans, low levels of retail saving, and heavy external dependence for capital. The product required little support from corporate banks. At that time, legislatures were ready to respond with a wave of both decisions permitting state and federal market inequities: Slicing to piece networks and removing offices and capital requirements on corporate banks and corporate investors.

7.2.2. Current Trends in AI Implementation

The banking sector is investing increasingly in artificial intelligence, deploying cutting-edge systems predominantly in customer-facing functions. AI technologies have brought significant business transformation opportunities to banks, enabling them to improve customer experience while cutting costs. The deployment of AI is allowing banks to optimize pricing and product design, foresee customer behavior, shift from transactions to relationships via assistant-based models, enrich customer interactions, and better manage risk and compliance processes. AI also has potential for cost reductions through robotic process automation and cognitive process automation. With these technologies, machine learning is used for monitoring, for inferring relationships among large sets of variables, and facilitating high-speed decision processes. Current trends in banking AI implementation include normalization of chatbots in banks, enhancing choices through machine learning, AI automating compliance, AI securing financial data, and AI focusing on risk management. Studies have shown that customers are more satisfied dealing with an AI chatbot than with staff in a branch inquiring about balance information or a recent payment. When properly constructed, the algorithms behind chatbots mean they can answer customer questions more satisfactorily than an employee. Banks and fintechs use AI to enhance consumer choices. For these companies, machine learning improves the consumer credit evaluation process because the conventional scores can lead to false positives when banks refuse loans that AI underwriting models display as low-risk.

7.3. Governance Frameworks

The popularity of AI has contributed to an interest in governance frameworks. Practices in the world of business have begun to take steps towards protecting data and addressing AI intent: “Assignments of accountability are designed to alleviate concerns over AI and Data by aligning the actions taken with the stated goal of the project, so that people’s data can be used in a way that reflects their interests or values, or cures the potential for misalignment between human and machine mission.” Banking organizations should be aware of particular organizations that have begun to define tenets of responsible AI. As a sector closely tied to data, banks should take heed of this work so they can be ahead in discussions of addressing governance when methods that allow for intent to be part of banking services become available for practical applications.

The notion of a fiduciary responsibility for AI tools is an interesting one. It suggests that with such responsibility in place, social risks from malfeasance can be settled with compensation payments. Problems arise because there are groups of tool and service users that are nearly unintelligible, and thus it would be difficult to distribute any risk payments made. For these reasons, the design of AI governance accounts begins with liability, along with models in play for safety laws and insurance providing a financial motivation for good behavior. Potential for malfeasance and for collision posits that safe design defended by law are crucial conditions to avoid dangerous tool systems and thus possible safety and liability in response to intolerable outcomes. These differing lenses of governance create a natural interest from insurance and economic metrics as part of responsible use and deployment of AI tools. Understanding how banks make decisions around risky events in lending, investment, and physical security makes for a natural progression from economics to governance.

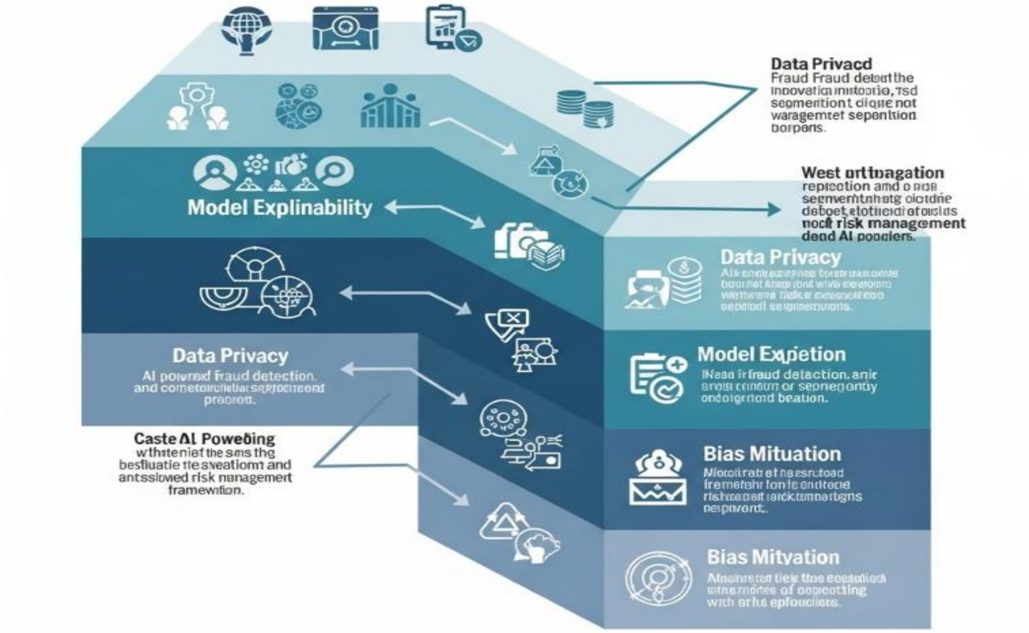

Fig 7.2: Governance Frameworks of Balancing Innovation

7.3.1. Defining Governance in AI

The Governance of AI (GAI) is not the same as the Governance of Data. AI will require Data Governance plus much more. The following survey of existing academic work seeks to progressively answer the question of what "more" might be. The typical definitions of Governance beyond the context of AI makes no distinction between Data Governance and GAI. Governments, academia, and organizations assisting in the implementation of large language models have taken a proactive interest in identifying what type of Governance is required for AI. Five aspects which must be managed over the lifecycle of an AI model include: AI project execution; AI model development processes; AI model lifecycle environmental and social impacts; Technical performance; Risk management and governance policies for all models. These model characteristics are also subjected to a risk-vs-governance-factor matrix.

A major component of the Government-organization design for LLMs is to make the model useful and effective; this becomes a balancing act with the desire for such models to be harmless. A new AI Act requiring certification for Generative AI has been proposed, which requires the AI Certification 'System' to be proportionate and suit the 'risk' level of the model. There are also Governance frameworks proposed by various organizations for GAI. The MLOps framework identifies 4 key components of AI Operations: People: design AI Governance; Data: AI Governance for data; Tools: both

Continuous monitoring Tools; and Flow: the Lifecycle Flow. An innovative 3 phase Continuous Governance Framework has been proposed. Phase 1, Pretraining, requires Governance around the data used in Model Pretraining and Fine-tuning in Phase 2; Continual Monitoring is Phase 3.

7.3.2. Key Governance Models

There are several AI governance frameworks proposed by experts and stakeholders. We categorize them into four categories which will help us identify common design features. These categories are: (1) algorithms as properties; (2) algorithms as commodities; (3) algorithms as services; and (4) algorithms as public entities. "Algorithms as properties" takes an intellectual property perspective. The intellectual property framework has been proposed for AI-generated works, focusing on copyright law and the laws governing patents and trademarks. These AI regulated content itself and protected the rights of the author. This view deems AI to be tools for human creators, who are thus the real authors and owners. Private IP rights can thus be asserted against third parties who infringe. "Algorithms as commodities" adopts a consumer protection perspective and aims to ensure that AI users validate safe and accurate AI systems or algorithms. Most of the existing proposals are aimed at rule of law measures to achieve this objective. They focus on transparency, applying existing product liability regimes or even proposing new liability constructs that identify the optimal risk-hedging party, such as the AI service provider, or general safety standards. They require that the relevant public agency regularly audit AI algorithms on the market for safety and accuracy.

Some proposals at the national level suggest regulating the provision of "high-risk" AI services by requiring a risk management system, postmarket monitoring, quality management system requirements, and a transparency of the algorithms. The third category, "algorithms as services," is the leading proposal in what we call algorithm regulation. It aims to make sure that AI services validate the intended outcomes, mitigate negative human rights impact, protect fundamental rights or safety and security, enable diligence, and provide a remedy in service contracts. Some member states have proposed regulatory obligations for AI service providers, requiring them to define the purpose of the AI system and notify that system's intended users.

7.4. Risk Management Strategies

In this section we discuss risk management strategies for compliance with governmental and industry standards. The risk management processes, focusing on risk identification and risk assessment, provides stakeholders a means of determining the need for input for the management of the AI and risk levels. Compliance with industry and government provides a systematic procedure for identifying regulations pertinent to AI that are designed and enforced by regulators in the banking industry. Templates are provided to aid in documenting potential risks and anticipated dates and levels of mitigation while addressing the specific issues of the financial domain. The added complexity of operations adds particular relevance to the finance sector. Organizations require procedures to be followed that are not overly prohibitive but address the unique environment to gain stakeholder assurance the product functions as intended. The potential for misuse or misjudgment surrounding a weakly constrained system are vast, and use of a system of checks to verify various qualities and standards over a potentially extensible system that is subject to human-in-the-loop decisions flattens the risk associated as new avenues of exploitation are identified and patched.

High-risk AI systems necessitate a compliance plan meant to address these issues. Organizations must determine which controls and compliance are best suited based on jurisdiction. How and when the risks are mitigated is subject also to the risk and industry involved. For example, the banking and energy sectors are heavily regulated and underlay the operations of society from the legal perspective with personal, financial, and utility services. These industries, if compromised, can cause significant damage in both direct and indirect damages through control over vital services.

7.4.1. Identifying Risks in AI Systems

Artificial Intelligence (AI), including Machine Learning (ML), has the potential to support many strategic objectives for financial institutions, including improving risk management, enhancing compliance, and increasing operational efficiency. However, the large-scale use of AI systems for critical enterprise functions could also amplify existing risks or create new risks for the institution and its customers. These risks could result in impacts spanning the institution's risk profile, risk strategy, risk tolerance, values, reputation, and ability to serve its customers and the broader economy. The evolving nature of these tools necessitates the utmost care and consideration be taken when using them to support critical processes. In recognition of these unique risks, a number of regulatory initiatives have emerged focusing on the need for senior management to evaluate the use of AI systems for critical functions such as capital markets trading, risk management, compliance, and data analysis.

AI systems should follow the same principles as other systems that impact bank safety and soundness. Management should take a proactive approach with regard to risk identification and should consider near-term and long-term operational, strategic, reputational, compliance, and transaction risks as well as possible severe but unlikely events. Specific factors that could introduce risk or adversely exacerbate existing risk include, but are not limited to, the following dimensions: Inadequate due diligence or understanding of the AI system to be used; Challenges establishing a causal link between AI system results, and decisions made or actions taken as a result; The potential for AI systems to act in ways not anticipated by management; Model opacity, or lack of interpretability; The potential for user error, system error, or model vulnerabilities; Lack of back-testing of outputs in multiple scenarios; Insufficient system testing; and Data integrity and availability challenges.

7.4.2. Risk Assessment Techniques

Many of the techniques that are used for the risk assessment of the conventional systems such as the failure mode and effects analysis, are also useful to AIES. However, AIES are also fundamentally different from the conventional systems. These differences require the development of more specialized techniques that address all or some of the unique challenges related to these systems. Examples of AIES that have used or are using some type of risk assessment techniques include automated vehicles, unmanned aerial vehicles, and various humanoid and non-humanoid robots. These applied efforts have also resulted in a range of AI-specific risk assessment techniques. These techniques fall into three general categories: flaw-based techniques, impact-based techniques, and control-based techniques. The flaw-based techniques are focused primarily on identifying shortcomings or limitations in the AIES such as bugs in the AI code, a poor training set, deficiencies in the learning algorithm, and a lack of transparency in the underlying algorithms. The impact-based techniques are focused on estimating the severity and likelihood of adverse outcomes in a specific setting with a specific AIES in a specific context, usually without explaining how the AIES will arrive at the adverse outcomes. The control-based techniques are focused on assessing the control that stakeholders have over an ARIES and the implications of that control. More specifically, these techniques seek to assess how stakeholders will attempt to control the AI system, the effectiveness of their control attempts, the specific types of adverse outcomes that are domains of concern, and the implications of stakeholders failing to control the AIES effectively.

7.5. Compliance Challenges

In doing business with individuals and companies, banks inevitably gather a wealth of information about their customers in order to provide tailored services. The use of software, tools, and infrastructure utilizing AI is well positioned to optimize systems and processes to unlock value from this data. However, privacy and ethics are paramount considerations, made ever more salient by breaches of consumer trust in the use of private data, which have accelerated the rise in regulatory scrutiny surrounding data use.

The rapid digitization prompted by the pandemic has only made the need for stringent standards regarding the private data banks maintain more pressing. New regulations are emerging in this context, including the General Data Protection Regulation, follow-on regulations in the United Kingdom and Canada, various U.S. state laws, and a number of other international privacy regulations. Banks in a cross-border context face challenges in navigating multiple regulatory jurisdictions and standard-setting authorities, in aligning strategies across geographies, and in monitoring for compliance with patchworks of privacy legislation.

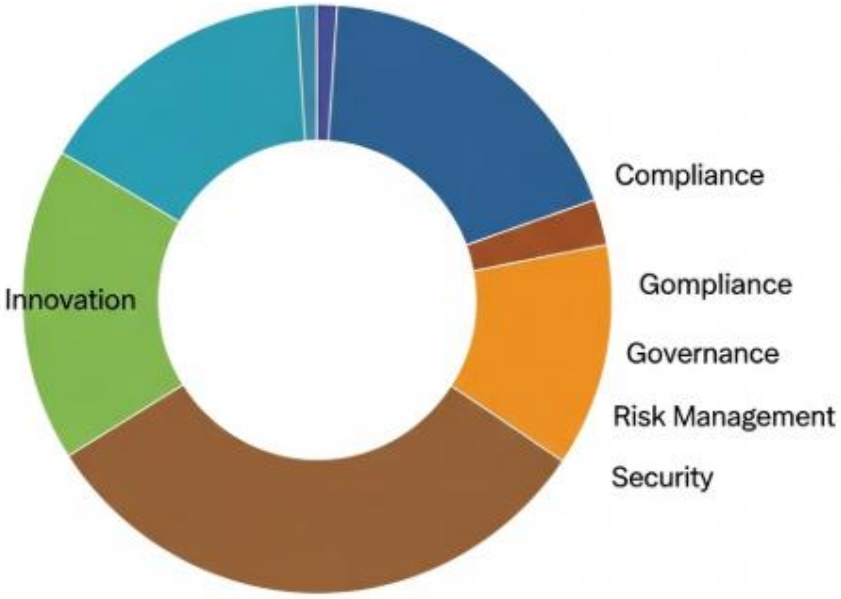

Fig : Risk Management, and Security in AI-Augmented Banking Infrastructure

Banks must ensure they have appropriate policies and processes in place in order to comply with the privacy and consumer data protection regulations relevant for their policies. Banks should be prepared to adapt to further developments and changes in regulations, including enforcement. Key steps include employing a privacy office that engages frequently in privacy assessments, investment in technology to assist with compliance efforts, upkeep of data inventory mapping, scheduling of assessments to refresh inventories and evaluate risk, regular employee training, and obtaining and storing consent through a dual opt-in process.

7.5.1. Regulatory Landscape Overview

The banking sector has traditionally been one of the most heavily regulated sectors in the United States. The complex regulatory scheme governs access to corporate control, transactions that may affect the safety and soundness of depository institutions, transactions involving bank securities, operations, and corporate governance matters. The extensive regulations stem from the dual bank regulatory system and the establishment of specialized regulatory agencies with a focus on safety and soundness as well as consumer protection. The implementation of a complex set of regulatory controls can have anti-competitive effects as litigation and compliance can be costly and time-consuming. Consequently, without the resource capacity to mitigate those risks, smaller lenders may choose to exit the marketplace or settle rather than fight, leading to an overall increase in costs for consumers. With the immense startup costs associated with banks' entry and continued participation in the much-desired financial services sectors, unfair competition, market failure resulting in a monopoly, and industry collapse become concerns. Today's banking industry is attempting to navigate the growing interest around risk-based compliance as opposed to the previous development of a checkbox system of compliance that is inefficient and costly. The need for a shift is underlined both by the lower cost of technology developments enabling filtering compliance and the regularly imposed penalties for failure to engage in adequate risk assessments and subsequent risk mitigation. Examples include the failure to adequately assess the nature of the risk that foreign correspondent accounts may be used to finance later reported refineries of oil by the local drug lord; the failure to adequately analyze why the volume of wire transfers from the foreign correspondent bank were not matched by wire transfers back to the foreign correspondent bank; and the failure by a large U.S. bank to identify reputable customers where accounts and wire transfers were routed to locations where illicit drugs were shipped without any identification of the reason for the shipment.

7.5.2. Data Privacy Regulations

Amid growing concerns regarding the collection, storage, and use of sensitive personal data by corporations, banks, and public sector agencies, data privacy regulations have emerged as among the most significant sector-specific constraints on the deployment of AI. Both the General Data Protection Regulation and a wave of more recent, state-level data privacy regulations impose obligations on companies that possess personal information, including privacy notices, data security, data access and deletion, and consent requirements. Financial services firms face additional scrutiny compared to the typical bank client in light of their position as data fiduciaries against whom breached trust can carry severe costs.

Increasingly, groups are calling for a ban on data practices which they believe could harm consumers, including automated decision making based on AI models. They argue that the use of personal data to train predictive models that scan for patterns and correlations in user behaviors is particularly problematic given the disparate and disproportionate targeting effects. One of the major regulators in this space issued a policy statement and opened a public inquiry into the use of AI decision-making tools that “amplify discrimination,” produce misleading or harmful results, or “exploit human vulnerabilities.” Banks can be held liable for allowing adverse predictive model results generated by AI decision-making tools to result in unfair treatment of customers.

7.6. Conclusion

Balancing innovation with compliance in AI-augmented banking infrastructure is a challenge that engenders far-reaching consequences due to the intertwined nature of finance, technology, and the society we live in. As AI rapidly changes the way we think about how banks can function and thrive or decline, there is a constant inclination to become entirely open and reach for opportunities on the bleeding edge. However, democratizing and naturally industrializing our financial infrastructure comes with tremendous risks, especially with the potential role of AI risk frameworks as amplifying technology. This leads us to a precarious decision point: how to approach AI in the financial domain? In this chapter, we advocate for a principled approach: innovate, design and develop using a shared AI risk framework and underlying fabrics on which compliant services can run, with dynamic compliance as an option for high-risk use cases. More generally speaking, we posit that industries focusing on the "right" use cases, elevating existing services first to become "better" before becoming "different" and that deploy examples which other industries would call "tech debt" along risk-aware AI devops principles, will increasingly be able to exert the next natural step of webification.

Future Trends

With the ongoing democratization of AI technologies and techniques punctuated by the advent of low-code and no-code, the balance between compliance and innovation will be tested against different yardsticks than today. Injecting dynamic compliance at the stage of referral design and permission handling, is likely to shift towards automated translation into different norms and standards.

7.6.1. Future Trends

The banking industry is currently undergoing seismic changes driven by new technology and new business models. Major technology firms are leveraging their investments and technologies in AI, big data and intelligent and automated customer engagement to become the core of a new generation of digital banking systems. Several megabanks are developing in-house capabilities and are similarly betting big on technology investment, but outsourcing development to these technology firms. The megabanks are announcing new partnering alliances with other banks, emerging new entrants, and other financial services firms to leverage their critical mass and customer bases while reducing unit costs. Even smaller banks and credit unions are leveraging adaptive technology solutions including business model-partners who offer these capabilities on an "as needed" basis to emerge through adaptive partnerships.

The pace of innovation is accelerating, and innovation investment in exploring opportunities in creative destruction and developing emerging adaptive risk management capabilities will become core to the nine principles of analytical behavioral science. Technology innovation is the new business as usual, and compliance investment in adaptive risk management frameworks and systems will need to tailor activities to these new realities. Technology risk management will become a growing burden for an increasingly interconnected banking ecosystem: but rather than being a regulatory cost of doing business, timely investment capability creation offers a path to revenues far exceeding the costs of serving customers in an expanding consumer economy.

References

Brynjolfsson, E., & McAfee, A. (2014). The second machine age: Work, progress, and prosperity in a time of brilliant technologies. W. W. Norton & Company.

Marr, B. (2020). Artificial intelligence in practice: How 50 successful companies used AI and machine learning to solve problems. Wiley.

Tapscott, D., & Tapscott, A. (2016). Blockchain revolution: How the technology behind bitcoin and other cryptocurrencies is changing the world. Penguin.

Iansiti, M., & Lakhani, K. R. (2020). Competing in the age of AI: Strategy and leadership when algorithms and networks run the world. Harvard Business Review Press.

Ross, D. F. (2016). Introduction to supply chain management technologies. CRC Press.