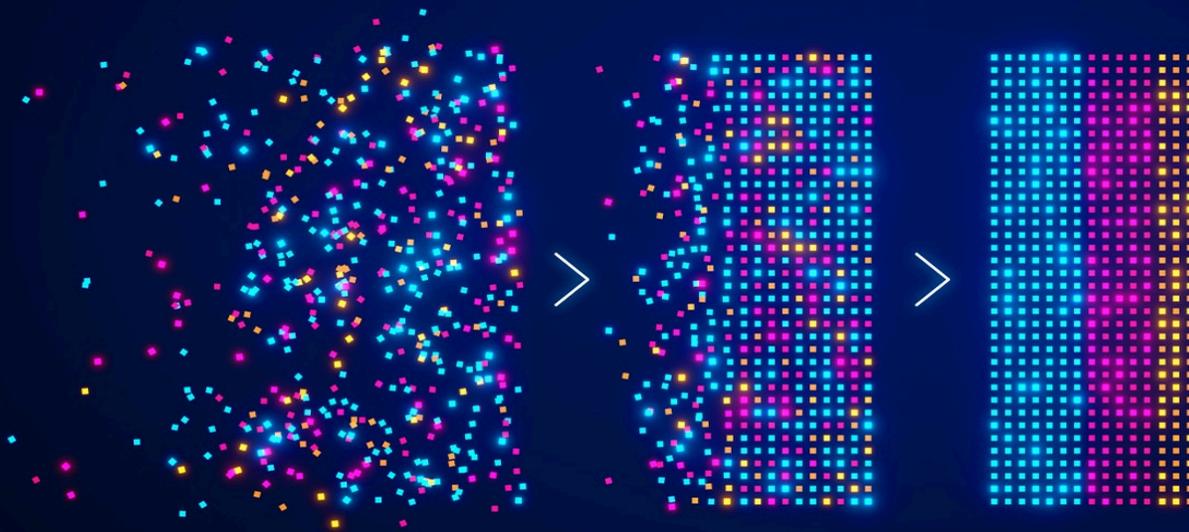


Risk & Resilience Practice

How generative AI can help banks manage risk and compliance

In the next five years, generative AI could fundamentally change financial institutions' risk management by automating, accelerating, and enhancing everything from compliance to climate risk control.

by Rahul Agarwal, Andreas Kremer, Ida Kristensen, and Angela Luget



Generative AI (gen AI) is poised to become a catalyst for the [next wave of productivity gains](#) across industries, with financial services very much among them. From modeling analytics to automating manual tasks to synthesizing unstructured content, the technology is already changing how banking functions operate, including how financial institutions manage risks and stay compliant with regulations.

It's imperative for risk and compliance functions to put guardrails around gen AI's use in an organization. However, the tech can help the functions themselves improve efficiency and effectiveness. In this article, we discuss how banks can build a flexible, powerful approach to using gen AI in risk and compliance management and identify some crucial topics that function leaders should consider.

Seizing the promise of gen AI

Gen AI has the potential to revolutionize the way that banks manage risks over the next three to five years. It could allow functions to move away from task-oriented activities toward partnering with business lines on strategic risk prevention and having controls at the outset in new customer journeys, often referred to as a ["shift left" approach](#). That, in turn, would free up risk professionals to advise businesses on new product development and strategic business decisions, explore emerging risk trends and scenarios, strengthen resilience, and improve risk and control processes proactively.

These advances could lead to the creation of AI- and gen-AI-powered risk intelligence centers that serve all [lines of defense](#) (LODs): business and operations, the compliance and risk functions, and audits. Such a center would provide automated reporting, improved risk transparency, higher

efficiency in risk-related decision making, and partial automation in drafting and updating policies and procedures to reflect changing regulatory requirements. It would act as a reliable and efficient source of information, enabling risk managers to make informed decisions swiftly and accurately.

For instance, McKinsey has developed a [gen AI virtual expert](#) that can provide tailored answers based on the firm's proprietary information and assets. Banks' risk functions and their stakeholders can develop similar tools that scan transactions with other banks, potential red flags, market news, asset prices, and more to influence risk decisions. These virtual experts can also collect data and evaluate climate risk assessments to answer counterparty questions.

Finally, gen AI could facilitate better coordination between the first and second LODs in the organization while maintaining the governance structure across all three. The improved coordination would enable enhanced monitoring and control mechanisms, thereby strengthening the organization's risk management framework.

Emerging applications of gen AI in risk and compliance

Of the many promising applications of gen AI for financial institutions, there's a set of candidates that banks are exploring for a first wave of adoption: regulatory compliance, financial crime, credit risk, modeling and data analytics, cyber risk, and climate risk. Overall, we see applications of gen AI across risk and compliance functions through three use case archetypes.

Through a *virtual expert*, a user can ask a question and receive a generated summary answer that's built from long-form documents and unstructured

data. With *manual process automation*, gen AI performs time-consuming tasks. With *code acceleration*, gen AI updates or translates old code or writes entirely new code. All these archetypes can have roles in the key responsibilities of risk and compliance:

- *Regulatory compliance.* Enterprises are using gen AI as a virtual regulatory and policy expert by training it to answer questions about regulations, company policies, and guidelines. The tech can also compare policies, regulations, and operating procedures. As a code accelerator, it can check code for compliance misalignment and gaps. It can automate checking of regulatory compliance and provide alerts for potential breaches.
- *Financial crime.* Gen AI can generate suspicious-activity reports based on customer and transaction information. It can also automate the creation and update of customers' risk ratings based on changes in know-your-customer attributes. By generating and improving code to detect suspicious activity and analyze transactions, the tech can improve transaction monitoring.
- *Credit risk.* By summarizing customer information (for example, transactions with other banks) to inform credit decisions, gen AI can help accelerate banks' end-to-end credit process. Following a credit decision, it can draft the credit memo and contract. Financial institutions are using the tech to generate credit risk reports and extract customer insights from credit memos. Gen AI can generate code to source and analyze credit data to gain a view into customers' risk profiles and generate default and loss probability estimates through models.
- *Modeling and data analytics.* Gen AI can accelerate the migration of legacy programming languages, such as the switch from SAS and COBOL to Python. It can also automate the monitoring of model performance and generate alerts if metrics fall outside tolerance levels. Companies are also using gen AI to draft model documentation and validation reports.
- *Cyber risk.* By checking cybersecurity vulnerabilities, gen AI can use natural language to generate code for detection rules and accelerate secure code development. It can be useful in "red teaming" (simulating adversarial strategies and testing attack scenarios). The tech can also serve as a virtual expert for investigating security data. It can make risk detection smarter by speeding and aggregating security insights and trends from security events and behavior anomalies.
- *Climate risk.* As a code accelerator, gen AI can suggest code snippets, facilitate unit testing, and assist physical-risk visualization with high-resolution maps. It can automate data collection for counterparty transition risk assessments and generate early-warning signals based on trigger events. As a virtual expert, gen AI can automatically generate reports on environmental, social, and governance (ESG) topics and sustainability sections of annual reports (see sidebar, "How generative AI can speed financial institutions' climate risk assessments").

Once companies have embedded gen AI in these roles and functions, they have seen a second wave of emerging use cases across other aspects of risk management. Gen AI can streamline enterprise risk by synthesizing enterprise-risk-management summaries from existing data and reports. It can

help accelerate the internal capital adequacy assessment process and model capital adequacy by sourcing relevant data. Banks can also use it to summarize risk positions and draft risk reports and executive briefings for senior management.

Another area in which gen AI can play an important role is operational risk. Banks can use it for operational automation of controls, monitoring, and incident detection. It can also automatically draft risk and control self-assessments or evaluate existing ones for quality.

Key considerations in gen AI adoption

While several compelling use cases exist in which gen AI can propel productivity, prioritizing them is critical to realizing value while adopting the tech responsibly and sustainably. We see three critical dimensions that risk leaders can assess to determine prioritization of use cases and maximize impact (exhibit).

Chief risk officers can base their decisions on assessments across qualitative and quantitative dimensions of impact, risk, and feasibility. This process includes aligning with their banks' overall visions for gen AI and associated guardrails, understanding relevant regulations (such as the EU AI Act), and assessing data sensitivity. All leaders need to be aware of the novel risks associated with this new tech. These risks can be broadly divided into eight categories:

- impaired fairness, when the output of a gen AI model may be inherently biased against a particular group of users
- intellectual property infringement, such as copyright violations and plagiarism incidents, as

foundation models typically leverage internet-based data

- privacy concerns, such as unauthorized public disclosure of personal or sensitive information
- malicious use, such as dissemination of false content and use of gen AI by criminals to create false identities, orchestrate phishing attacks, or scam customers
- security threats, when vulnerabilities within gen AI systems can be breached or exploited
- performance and “explainability” risks, such as models providing factually incorrect answers and outdated information
- strategic risks through noncompliance with ESG standards or regulations, creating societal or reputational risks
- third-party risks, such as leakage of proprietary data to the public realm through the use of third-party tools

Winning strategies for planning a gen AI journey

Organizations that can extract value from gen AI should use a focused, top-down approach to start the journey. Given the scarcity of talent to scale gen AI capabilities, organizations should start with three to five high-priority risk and compliance use cases that align with their strategic priorities. They can execute these use cases in three to six months, followed by an estimation of business impact. Scaling the applications will require the development of a gen AI ecosystem that focuses on seven areas:

Exhibit

Risk leaders can prioritize risk, impact, and feasibility considerations when planning gen AI implementation in a risk function.

Initial assessment to prioritize generative AI (gen AI) use cases based on impact, feasibility, and risk scoring¹

Risk

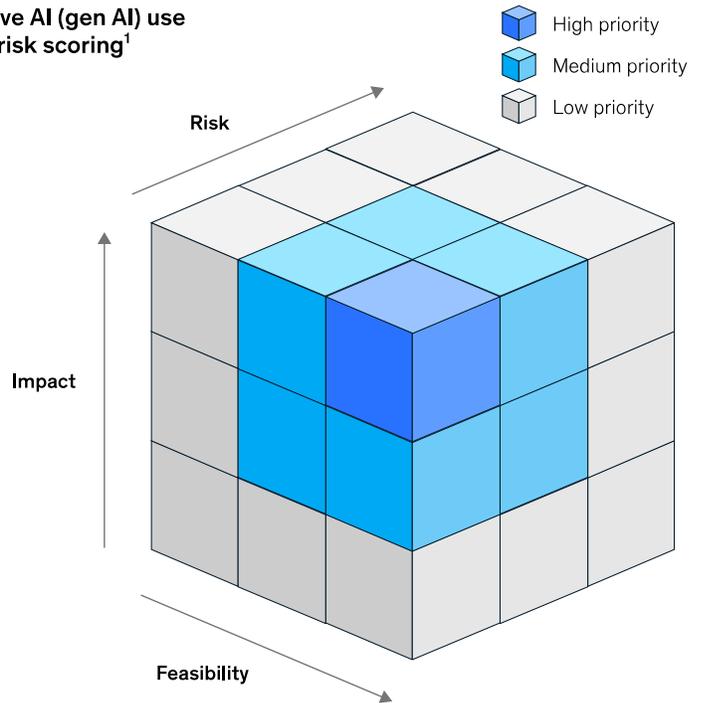
Source and use of data
Security threats
Performance and transparency
Strategic risk
Third-party risk

Impact

Revenue generation
Operational cost savings
Strategic priority of the organization
Scalability of gen AI ecosystem
Solves pain points not addressed by traditional AI

Feasibility

Data quality and architecture
Readiness of tech stack
Upskilling and hiring needs
Change management needs



¹Taking quantitative and qualitative dimensions into account.

McKinsey & Company

- a catalog of production-ready, reusable gen AI services and solutions (use cases) that can be easily plugged into a range of business scenarios and applications across the banking value chain
- a secure, gen-AI-ready tech stack that supports hybrid-cloud deployments to enable support for unstructured data, vector embedding, machine learning training, execution, and pre- and postlaunch processing
- integration with enterprise-grade foundation models and tools to enable fit-for-purpose selection and orchestration across open and proprietary models
- automation of supporting tools, including MLOps (machine learning operations), data, and processing pipelines, to accelerate the development, release, and maintenance of gen AI solutions

- governance and talent models that readily deploy cross-functional expertise empowered to collaborate and exchange knowledge (such as language, natural-language processing, and reinforcement learning from human feedback, prompt engineers, cloud experts, AI product leaders, and legal and regulatory experts)
- process alignment for building gen AI to support the rapid and safe end-to-end experimentation, validation, and deployment of solutions
- a road map detailing the timeline for when various capabilities and solutions will be launched and scaled that aligns with the organization's broader business strategy

At a time when companies in all sectors are experimenting with gen AI, organizations that fail to harness the tech's potential are risking falling behind in efficiency, creativity, and customer engagement. At the outset, banks should keep in mind that the move from pilot to production takes significantly longer for gen AI than for classical AI and machine learning. In selecting use cases, risk and compliance functions may be tempted to use a siloed approach. Instead, they should align with an entire organization's gen AI strategy and goals.

For gen AI adoption by risk and compliance groups to be effective and responsible, it is critical that these groups understand the need for new risk management and controls, the importance of data and tech demands, and the new talent and operating-model requirements.

Risk management and controls

With gen AI, a new level of risk management and control is necessary. Winning responsibly requires both defensive and offensive strategies. All

organizations face inbound risks from gen AI, in addition to the risks from developing gen AI use cases and embedding gen AI into standard workplace tools. So banks will need to evolve their risk mitigation capabilities accordingly.

The first wave heavily focuses on human-in-the-loop reviews to ensure the accuracy of model responses. Using gen AI to check itself, such as through source citations and risk scores, can make human reviews more efficient. By moving gen AI guardrails to real time and doing away with human-in-the-loop reviews, some companies are already putting gen AI directly in front of their customers. To make this move, risk and compliance professionals can work with development team members to set the guardrails and create controls from the start.

Risk functions need to be vigilant to manage gen AI risks at the enterprise level. They can fulfill that obligation by taking the following steps:

1. Ensure that everyone across the organization is aware of the risks inherent in gen AI, publishing dos and don'ts and setting risk guardrails.
2. Update model identification criteria and model risk policy (in line with regulations such as the EU AI Act) to enable the identification and classification of gen AI models, and have an appropriate risk assessment and control framework in place.

Find more content like this on the
McKinsey Insights App



Scan • Download • Personalize



3. Develop gen AI risk and compliance experts who can work directly with frontline development teams on new products and customer journeys.
4. Revisit existing know-your-customer, anti-money laundering, fraud, and cyber controls to ensure that they are still effective in a gen-AI-enabled world.

Data and tech demands

Banks shouldn't underestimate the data and tech demands related to a gen AI system, which requires enormous amounts of both. Why? For one, the process of context embedding is crucial to ensure the accuracy and relevance of results. That process requires the input of appropriate data and addressing data quality issues. Moreover, the data on hand may be insufficient. Organizations may need to build or invest in labeled data sets to quantify, measure, and track the performance of gen AI applications based on task and use.

Data will serve as a competitive advantage in extracting value from gen AI. An organization looking to automate customer engagement using gen AI must have up-to-date, accurate data.

Organizations with advanced data platforms will be the most effective at harnessing gen AI capabilities.

Talent and operating-model requirements

Since gen AI is a transformational technology requiring an organizational shift, organizations will need to understand the related talent requirements. Banks can embed operating-model changes into their culture and business-as-usual processes. They can train new users not only on how to use gen AI but also on its limitations and strengths. Assembling a team of "gen AI champions" can help shape, build, and scale adoption of this new tech.

We expect gen AI to empower banks' entire risk and compliance functions in the future. This implies a profound culture change that will require all risk professionals to be conversant with the new tech, its capabilities, its limitations, and how to mitigate those limitations. Using gen AI will be a significant shift for all organizations, but those that navigate the delicate balance of harnessing the technology's powers while managing the risks it poses can achieve significant productivity gains.

Rahul Agarwal is an associate partner in McKinsey's New Jersey office, **Andreas Kremer** is a partner in the Berlin office, **Ida Kristensen** is a senior partner in the New York office, and **Angela Luget** is a partner in the London office.

The authors wish to thank Adrija Banerjee, Stephan Beitz, Adrian Foerster, Yilin Li, Anke Raufuss, Ibtesam Siddiqui, and Claudia Satrústegui for their contributions to this article.

Copyright © 2024 McKinsey & Company. All rights reserved.