



Embracing generative AI in credit risk

July 1, 2024 | Article

Credit risk organizations are already adopting gen AI technologies. How can they deploy them safely and at scale?

DOWNLOADS

[↓ Article \(6 pages\)](#)

Some technologies are so compelling that they quickly take on a life of their own. Generative AI (gen AI) made the leap from the laboratory to the mainstream in late 2022, when Open AI launched a public beta of its ChatGPT service. Within two months, it had more than 100 million users,^[1] making it the fastest-growing product in human history.

By the first quarter of 2023, big technology companies were integrating gen AI capabilities into their own products and offering programmatic access to generative models for business customers. A year on, gen AI is making its mark in multiple industries, including those that have

traditionally taken a relatively conservative approach to the adoption of emerging technologies—credit risk, for example.


McKinsey recently surveyed senior credit risk executives from 24 financial institutions, including nine of the top ten US banks. We asked these executives about their organizations’ adoption of gen AI, its current use cases, their future plans for it, and the challenges they expected.

Twenty percent of the respondents have already implemented at least one gen AI use case in their organizations, and a further 60 percent expect to do so within a year (Exhibit 1). Even the most cautious of these executives believe that gen AI will be part of their companies’ credit risk processes within two years.

Exhibit 1

In our survey, 80% of credit risk organizations expect to implement gen AI technologies within a year.

Expectations for implementing generative AI (gen AI),¹ % of respondents



¹Question: When do you expect to have implemented your first generative AI use case? (n = 24).
Source: October 2023 roundtable of chief credit risk officers

Use cases in credit risk

As these financial institutions gear up to use gen AI, they are considering potential applications across the full credit life cycle. In

general, such applications use large language models (LLMs) to combine, summarize, and analyze unstructured data and natural language. They can also output complex forms of natural language (such as reports, emails, and summary documents) and generate structured data or instructions for other software tools. Our survey revealed several potential use cases for gen AI in credit risk.

In *client engagement*, gen AI might be used to offer customers hyperpersonalized product mixes based on their profiles and activity histories. Gen AI systems could support relationship managers by drafting individualized outreach communications, summarizing meetings, and suggesting next steps. Gen AI–powered virtual experts could help customers identify and determine suitable products.

During *credit decision and underwriting* processes, gen AI tools could review documents and flag policy violations or missing data. They could draft outreach communications seeking clarifications or missing information from customers. And they could help compile information about customers, conduct credit analyses, and draft several sections of credit memos before credit officers review them. Agent-based gen AI systems can autonomously follow task sequences to extract information from sources, calculate relevant ratios, compare outcomes with typical thresholds, and summarize results in credit memos. These capabilities can all be developed in natural language, using plain English, with limited need for programming and advanced modeling skills.

Once credit is approved, gen AI can streamline and accelerate contracting processes. Gen AI systems can draft legal contracts, for example, or create outreach communications to inform customers about credit decisions and next steps that may be required.

In *portfolio monitoring*, gen AI tools can support portfolio managers in multiple ways, such as automating the creation of routine performance


and risk reports or drafting summaries (based on portfolio managers' analyses) of portfolio optimization options. Gen AI systems could even produce subsegment-specific optimization strategies in line with an organization's risk appetite and optimize an existing early-warning system (EWS) by consuming real-time unstructured information (such as news or market reports) to identify borrowers with elevated risk or borrower segments that may require attention.

Finally, gen AI tools can support *customer assistance* processes—for example, by drafting personalized outreach communications to customers in the event of issues. Gen AI systems could also identify suitable restructuring options and then guide customers through the restructuring process. In addition, some institutions are using gen AI to coach the interactions of their agents with customers, both in real time and with post-call analyses.

Respondents to our survey say they are exploring gen AI applications in all these areas. Portfolio monitoring is currently the leading area of activity among the respondents: nearly 60 percent are pursuing these use cases. Credit application processes are the next-largest area of reported activity, along with controls and reporting: just over 40 percent of our respondents report ongoing or planned projects in both areas (Exhibit 2). Across business lines, respondents see slightly more potential for gen AI in wholesale than in retail credit.

Respondents are planning gen AI projects across the credit risk life cycle.

Areas of generative AI (gen AI) application in the credit risk life cycle,¹ % of respondents



¹Question: Which areas are you applying (or planning to apply) generative AI in? (n = 24). Respondents can choose more than one option.
Source: October 2023 roundtable of chief credit risk officers

McKinsey & Company

The current state of gen AI in credit risk

Gen AI has arrived in the credit risk world but has yet to transform it. Executives surveyed were candid about the current state of their gen AI use cases, which are mostly narrow, noncustomer-facing solutions addressing specific operational pain points.

One bank, for example, has developed a proof-of-concept gen AI tool that can prepopulate climate risk questionnaires for commercial clients. The bank's relationship managers are required to periodically complete such questionnaires as part of their climate risk monitoring. The gen AI system, based on an LLM, extracts relevant information from the client's annual reports and other disclosures. These source documents are preprocessed to identify relevant sections, which are presented to the model along with carefully designed prompts asking it to find and summarize key information. The model provides a synthesized response, including relevant citations to the source material. Finally, human subject matter experts review and validate the results.

Another use case, which several banks have explored, is the use of gen AI in drafting credit memos. In commercial banking, the first line must often invest significant amounts of time in collecting information, performing analyses, and writing memos for credit decision and underwriting purposes. Gen AI tools can perform tasks such as extracting, collecting, and sourcing information; analyzing financial information; visualizing data; and drafting sections of memos by following preset instructions. Portfolio managers can then review the drafted memo, together with an estimated confidence level offered by the gen AI tool, before finalizing it. In addition to freeing up capacity for other activities, this tool can improve the consistency and accuracy of the memos generated and, potentially, speed up the credit decision process.

The programming of such systems can be executed in natural language by using agent-based systems, without a need for programming or advanced modeling skills. Meta-agents can coordinate carefully crafted agents, specializing in specific tasks, to achieve results from multistep tasks. In the case of credit memos, these can include extracting information, calculating ratios, and summarizing information. An additional agent layer can boost risk control mechanisms and help reduce common gen AI pitfalls, such as hallucination.

Since the introduction of such gen AI systems, the banks using them have reduced the time required to answer climate risk questions by approximately 90 percent, from more than two hours to less than 15 minutes. The system's answers are fully correct 90 percent of the time.

Challenges

Executives acknowledge that scaling up the application of gen AI in credit risk will be challenging. The most significant barriers, highlighted by 75 percent of our respondents, concern risk and governance. Major risk categories associated with the use of gen AI include the following:

- impaired fairness in algorithms that confuse or mislead users
- IP infringements, such as copyright violations or plagiarism
- privacy violations resulting from the use of personal or sensitive information to train models
- the generation of malicious content
- security threads and related vulnerabilities
- performance and explainability issues
- the risk of using proprietary data from third parties
- environmental, social, and governance (ESG) effects, such as increased carbon emissions or workforce disruptions

The key risks can lead to regulatory, legal, reputational, and business consequences when not effectively managed.

Sixty-seven percent of the participants highlighted potential shortages of gen AI capabilities inside the organization. Further challenges, cited by around 50 percent of participants, include difficulties defining uses cases and value at stake.

We asked participants to elaborate on the need for frameworks or guardrails to manage risks in gen AI applications. The top concern, cited by 79 percent of respondents, was data quality, followed by model risk issues (mentioned by 58 percent), such as transparency, audibility, fairness, and explainability.

A lack of formal and coordinated organizational support for gen AI in credit risk organizations exacerbates some of these challenges. Only a third of the respondents' institutions have established a center of excellence (CoE) to manage gen AI use cases. Less than 10 percent of the respondents report that their organizations now define gen AI use cases centrally. Most of them are initiated in a decentralized manner, so common practices and lessons are not leveraged.

Building a gen AI ecosystem

To capture the full potential of gen AI in credit risk, financial institutions must move beyond today's ad hoc approach and develop a common set of practices to prioritize, develop, deploy, maintain, and reuse gen AI applications. Eight such practices are essential:

- *An AI road map.* This should align with the organization's broader business strategy, explain the required capabilities and solutions, and provide a timeline for development, launch, and deployment at scale.
- *Aligned processes* for building gen AI tools. These should support rapid—but safe—end-to-end experimentation, comprehensive validation, and the deployment of solutions.
- A secure, gen AI-ready *technology stack* that supports hybrid-cloud environments, so companies can access the computing power they need to train models and then run them at scale. Such a stack should be able to manage unstructured data, train and execute models, and pre- and postprocess data.
- Integration with enterprise-grade *foundation models and tools*. These are large deep-learning neural networks, such as large language models that underpin advanced gen AI systems, along with software

toolkits that support their customization and deployment. Gen AI applications use such models directly or build on them to develop proprietary solutions.

- Robust *automated supporting tools*. These include machine-learning ops (MLOps, systems to manage training and the development of models) and appropriate data infrastructure and processing pipelines to support the development, release, and maintenance of use cases.
- A *governance and talent model* that can deploy cross-functional expertise to support gen AI development. The people providing this expertise might include software developers, natural-language-processing (NLP) specialists, teams to run reinforcement learning based on human feedback (RLHF), cloud computing specialists, AI product leaders, and legal and regulatory experts.
- A *modular solution architecture*. This allows parallel development and customizable connections across different layers, such as the UX layer and the business logic layer.
- The product of these practices should be a *library of production-ready, reusable gen AI services and solutions*. The items in the library can be plugged into a range of business scenarios and applications across the credit value chain.

Developing and deploying these eight practices will take most institutions time, but even deploying some of them can improve efficiency and effectiveness significantly. For example, institutions that have implemented two practices report that successful gen AI deployments have been accelerated by 30 to 50 percent. First, such institutions follow a modular solution architecture that includes three layers: a user experience layer, a business logic layer, and an infrastructure layer. These are enabled by the organization's operating model. Second, the institutions reuse existing components and adopt

open-source libraries; for example, developers can pick and choose from a multitude of ready-made modules (such as data retrieval pipelines, prompt libraries, and guardrails) from open-source tools to build an end-to-end gen AI solution rapidly—often in just one or two weeks.

Major credit risk players are embracing generative AI at speed. The technology has transformative potential, promising to improve efficiency, accuracy, and personalized services across the credit life cycle. Although early adopters have begun to reap these benefits, widespread implementation requires financial institutions to overcome significant challenges related to risk governance, talent acquisition, and the creation of a comprehensive ecosystem to support gen AI applications. By addressing these hurdles and fostering an environment that encourages innovation and collaboration, financial institutions can unlock the full potential of gen AI, setting a new standard for excellence in credit risk management.

How relevant and useful is this article for you?

☆☆☆☆☆

1. Krystal Hu, “ChatGPT sets record for fastest-growing user base—analyst note,” Reuters, February 2, 2023.
-

ABOUT THE AUTHOR(S)

Andreas Kremer is a partner in McKinsey’s Berlin office; **Arvind Govindarajan** is a partner in the Boston office; **Himanshu Singh** is a partner in the New York office, where **Ida Kristensen** is a senior partner; and **Elaine Li** is a consultant in the Silicon Valley office.

EXPLORE A CAREER WITH US

[Search openings](#)