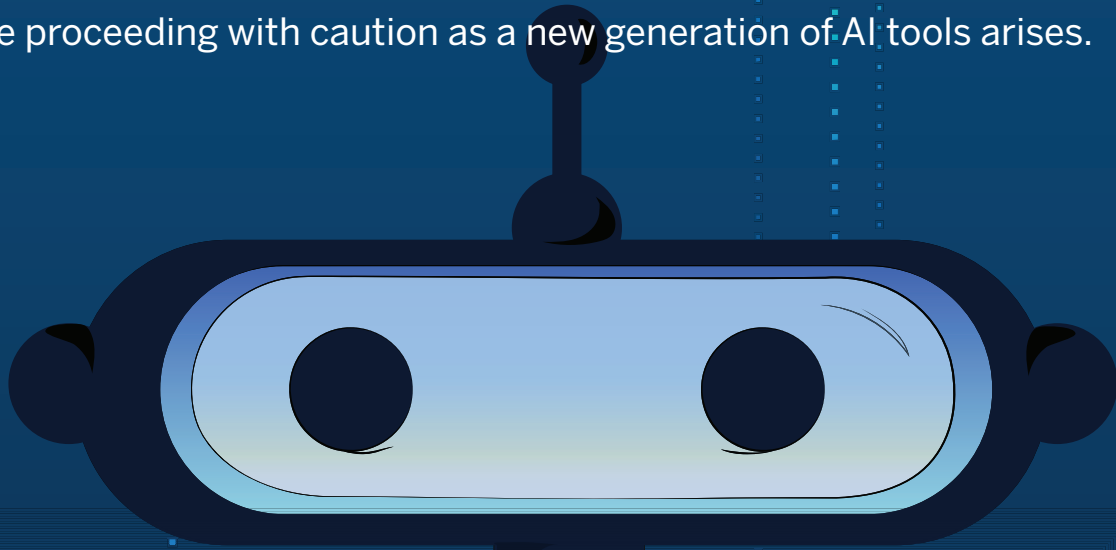# Preparing for banking's AI revolution

Banks are proceeding with caution as a new generation of AI tools arises.

By Matt Kuhrt

Ask Chatbot...

Original research from

AMERICAN BANKER

Preparing for banking's AI revolution

Introduction

ChatGPT's appearance on the scene in November 2022 brought generative AI into the forefront of the public's imagination. The ability to ask a computer a real question and have it answer in plain language or provide a new image seemed to open a new horizon of possibilities for improving digital experiences. It also introduced the public to a range of concerns over what it means to have a computer authoring new content, making authoritative statements that may be inaccurate or even committing plagiarism and violating copyright laws.

It's important to note that AI has actually been in use for several years (in some cases decades) quietly generating more efficient experiences by automating menial tasks and drawing obscure patterns out of large datasets. The users of these tools, however, haven't always been aware of what was going on behind the scenes. And what people do and don't know can affect how willing they are to use and deploy AI, particularly as its capabilities become much more sophisticated.

Caution is certainly warranted as AI advances, particularly for highly regulated industries like banking, requiring careful risk management as users look to maximize benefits. But banks also need to ensure that caution doesn't lead to paralysis.

As the benefits of a new generation of AI tools spread, consumers will expect AI-enhanced experiences across the board. And banks won't be immune to that expectation. As organizations across industries pursue AI applications, standards and best practices will evolve quickly, and talent capable of integrating AI innovations will likely be in short supply. How banks and their vendors position themselves now will be important factors in their ability to take full advantage of the opportunities ahead.

Why read this report? This study explores the unfolding role of AI in all of its current iterations in the banking industry. The data highlights the intersection of personal and professional AI use, sentiment regarding risks and benefits, current organizational use and the predicted impact on the workforce.

Sponsored by

Key findings

- The majority of bankers are still educating themselves about the breadth of AI solutions and how they can best be applied to support their businesses. The top actions they are taking include exploring AI technology providers in their industry, attending industry conferences or events on AI and educating key stakeholders to get buy-in across the organization.

- Outside of ChatGPT, which three-quarters of banking professionals have used, familiarity with specific generative AI tools is low and opportunities for education on what those products do are plentiful.

- Almost all banks expect that AI will help their business in one way or another, suggesting the full extent of AI's impact on the industry could wind up being impressively large.

- Bankers are also concerned with potential business risks associated with AI, particularly loss of personal touch with customers and the introduction of ethical issues.

- Banks appear to be in the early stages of preparing for widespread implementation of innovative generative AI tools, with almost two-thirds either still collecting information ahead of an implementation or starting small pilot projects.

- In general, a majority of banking professionals expect that AI adoption will ultimately impact the nature of jobs in the industry without greatly reducing the number of jobs.

Research methodology

This research was conducted online by Arizent, parent company of American Banker, in January 2023 among 127 banking professionals across job titles and functions. The average bank size measured by assets is $81.4 billion. The respondents break out as follows:

- 28% executives (C-level and SVP)
- 34% upper management
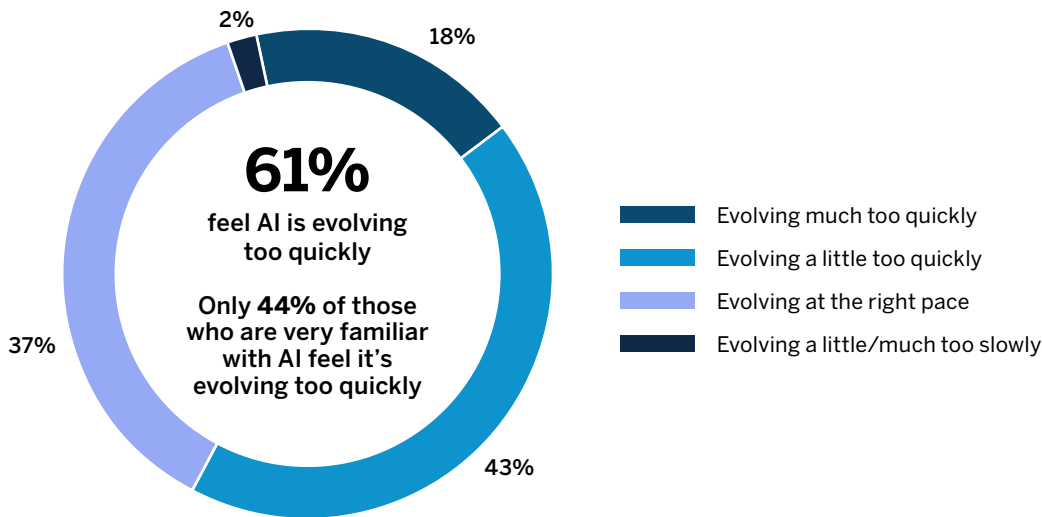- 21% managers
- 17% non-management staff

A bright spotlight on the evolution of AI

AI is far from a new technology. However, the introduction of the generative AI chatbot, ChatGPT, put powerful AI technology in the hands of the general public like never before. Suddenly, anybody with an internet connection could make incredibly powerful computers generate new content simply by asking a question. As with any potentially disruptive breakthrough, the possibilities seemed endless. Worries about potential unintended consequences weren't far behind either.

The banking industry has historically been slow to embrace new technology. The strict regulations that apply to banks' activities can complicate the adoption of new technological tools — especially when those tools may need access to sensitive data to do their job. It's not surprising, then, that most banking professionals are inclined to pump the brakes on the newest generation of AI-enabled tools (see Figure 1).

Figure 1: A majority of banking professionals feel AI is evolving too rapidly

Question: What do you personally think about the rate at which AI is evolving?

61% feel AI is evolving too quickly

Only **44%** of those who are very familiar with AI feel it's evolving too quickly

2%
18%
37%
43%

- Evolving much too quickly
- Evolving a little too quickly
- Evolving at the right pace
- Evolving a little/much too slowly

Source: Arizent, 2024 Unleashing the Power of AI
Base: Total respondents n=127; Those very familiar with AI n=36

A relative lack of familiarity with the detailed workings of AI technology may also be driving some of the hesitance to adopt new tools. ChatGPT is the only next-generation AI tool most banking professionals know by name. Nearly three-quarters (73%) say they have used it in a personal or professional setting. By contrast, only 35% have used the second most familiar tool, Bing AI.

As clear as it seems to be that ChatGPT changed public perception of AI, exactly how it did so may be less clear. Generative AI, which creates new content based upon existing patterns in a dataset, is the technology that distinguishes this generation of AI tools from the vast majority that have come before them (see the sidebar "What's different about generative AI?"). However, fewer than a quarter of banking professionals (23%) fully understand the difference between generative AI and AI in general.

What's different about generative AI?

In simple terms, AI refers to computing technologies that recognize patterns in datasets. A particular AI model's ability to recognize patterns accurately depends on both the quality and quantity of data used to train it. The continual increase in computing power has allowed AI models to consume ever larger amounts of training data — and therefore recognize patterns with ever greater accuracy.
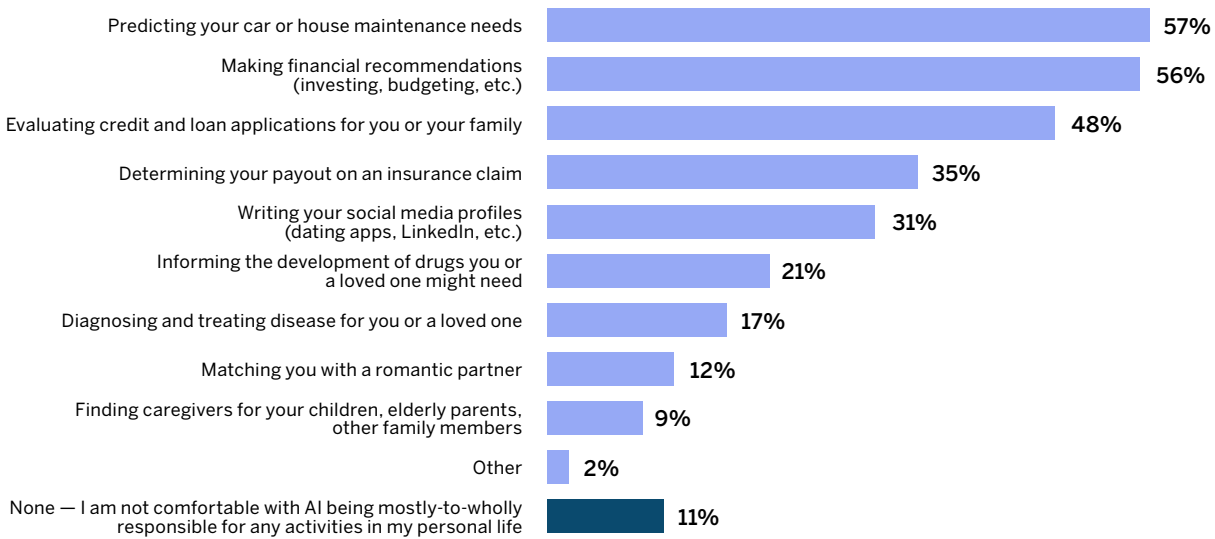
What distinguishes generative AI is what it does with the patterns once the model recognizes them. Generative AI uses those patterns to create something new — creating new content. That represents a major shift from use cases that depend on machines to recognize patterns that humans can then test for causal connections (e.g., whether a pattern of transactions suggest attempted fraud) or evaluate the patterns according to existing rule sets (e.g., whether a person's financial situation suits a bank's lending restrictions). It's also what worries bankers and their regulators — the new content autonomously created.

Attitudes toward AI vary widely by use case

Ultimately, attitudes toward AI's evolution have to do with the ways the technology actually gets put to use. AI has a well-established track record in areas such as automating rules-based processes and making recommendations based on established patterns of human activity. More than half of bank professionals say they would trust AI to predict car or home maintenance needs (57%) or make financial recommendations (56%). Nearly half (48%) would trust AI to evaluate credit and loan applications for themselves or their families (see Figure 2).

Figure 2: Trust in AI for personal use is limited when it comes to health and caregiving

Question: In your personal life, for which of these would you trust AI to be mostly-to-wholly responsible?

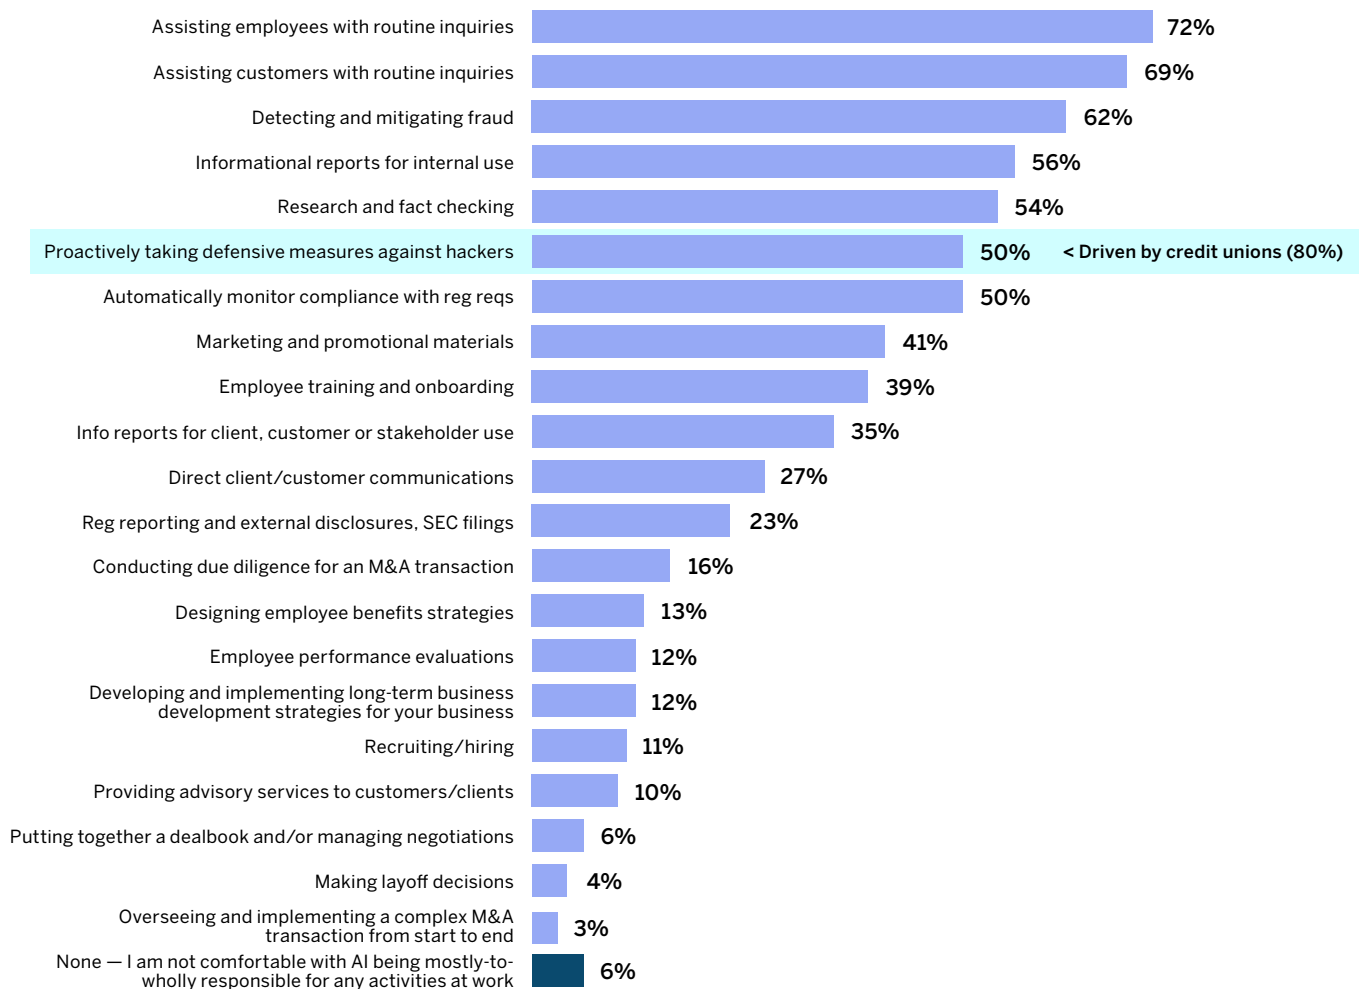| Category	Percentage
Predicting your car or house maintenance needs	57%
Making financial recommendations (investing, budgeting, etc.)	56%
Evaluating credit and loan applications for you or your family	48%
Determining your payout on an insurance claim	35%
Writing your social media profiles (dating apps, LinkedIn, etc.)	31%
Informing the development of drugs you or a loved one might need	21%
Diagnosing and treating disease for you or a loved one	17%
Matching you with a romantic partner	12%
Finding caregivers for your children, elderly parents, other family members	9%
Other	2%
None — I am not comfortable with AI being mostly-to-wholly responsible for any activities in my personal life	11%

Source: Arizent, 2024 Unleashing the Power of AI
Base: Total respondents n=127

When it comes to health and caregiving decisions, however, trust in AI diminishes greatly. Only 21% would trust the technology to inform drug development. Use cases like diagnosing and treating disease (17%) and finding caregivers for family members (9%) fare worse. The higher the potential human and emotional consequences, it seems, the more likely respondents are to prefer a human driving the decision-making process.

Attitudes toward professional use of AI demonstrate more trust for the technology overall (see Figure 3). Once again, however, trust levels fall as human consequences rise, particularly around ongoing human resources functions and high-touch customer interactions.

Figure 3: Attitudes toward professional use of AI vary widely by use case

Question: For purposes of this question, assume regulation on AI does not apply. Would you trust AI to be mostly-to-wholly responsible for any of the following for your business?

Use case	Percent
Assisting employees with routine inquiries	72%
Assisting customers with routine inquiries	69%
Detecting and mitigating fraud	62%
Informational reports for internal use	56%
Research and fact checking	54%
Proactively taking defensive measures against hackers	50% < Driven by credit unions (80%)
Automatically monitor compliance with reg reqs	50%
Marketing and promotional materials	41%
Employee training and onboarding	39%
Info reports for client, customer or stakeholder use	35%
Direct client/customer communications	27%
Reg reporting and external disclosures, SEC filings	23%
Conducting due diligence for an M&A transaction	16%
Designing employee benefits strategies	13%
Employee performance evaluations	12%
Developing and implementing long-term business development strategies for your business	12%
Recruiting/hiring	11%
Providing advisory services to customers/clients	10%
Putting together a dealbook and/or managing negotiations	6%
Making layoff decisions	4%
Overseeing and implementing a complex M&A transaction from start to end	3%
None — I am not comfortable with AI being mostly-to-wholly responsible for any activities at work	6%

Source: Arizent, 2024 Unleashing the Power of AI
Base: Total respondents n=127; Credit union respondents (all asset sizes) n=15

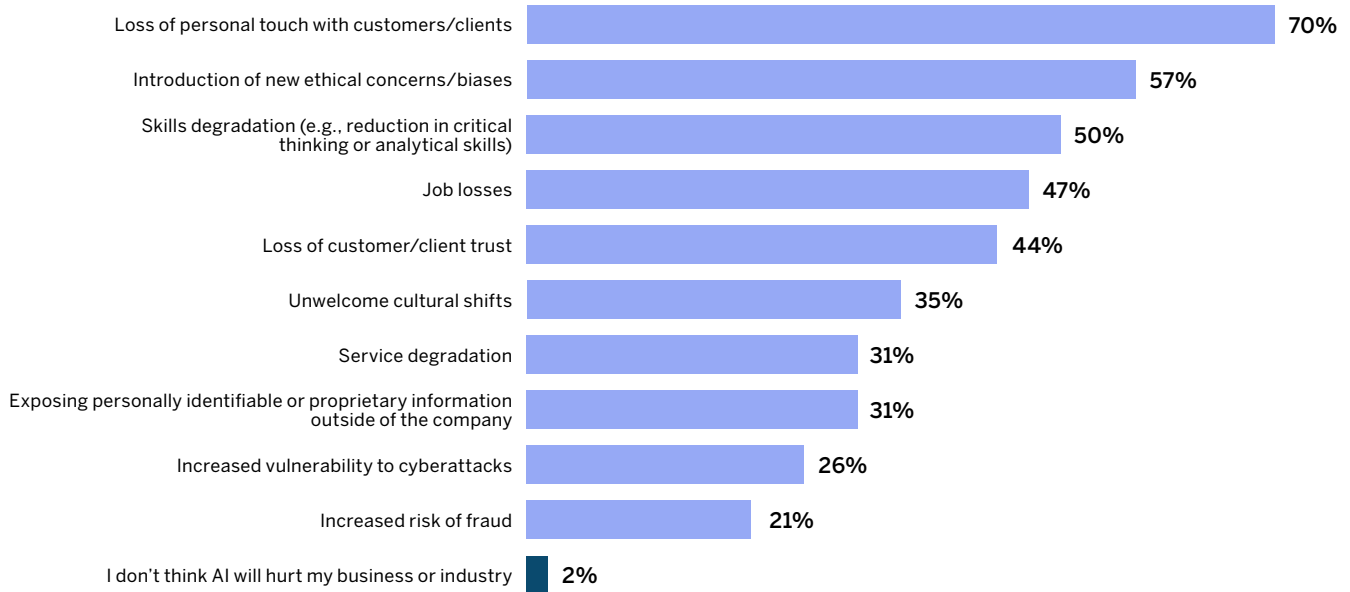Using AI technology to assist staff with information about routine inquiries or by focusing attention on important areas such as potential fraud, cyberattacks or regulatory issues appears relatively well established at this point. As use cases imply less human intervention, however, trust in AI tools falls precipitously. Complex tasks that involve synthesizing information, such as conducting due diligence around mergers and acquisitions, evaluating employee performance or providing advisory services to customers, inspire significantly less trust.

Bank professionals' concerns about the way AI might hurt their business run along similar fault lines. Although AI has become more sophisticated, bank professionals are hesitant about allowing it to take more consequential actions with less human oversight (see Figure 4).

Figure 4: Concerns about AI include loss of personal touch and ethical responsibility

Question: How do you expect AI, in all of its current iterations, might hurt your business?

Concern	Percentage
Loss of personal touch with customers/clients	70%
Introduction of new ethical concerns/biases	57%
Skills degradation (e.g., reduction in critical thinking or analytical skills)	50%
Job losses	47%
Loss of customer/client trust	44%
Unwelcome cultural shifts	35%
Service degradation	31%
Exposing personally identifiable or proprietary information outside of the company	31%
Increased vulnerability to cyberattacks	26%
Increased risk of fraud	21%
I don't think AI will hurt my business or industry	2%

Source: Arizent, 2024 Unleashing the Power of AI
Base: Total respondents n=127

This fault line appears to track the tectonic shift toward generative AI applications, which do more than identify patterns and make recommendations. When tools are making complex decisions or generating mission-critical content such as customer-facing communications or business strategies, the consequences of an undetected mistake become much higher.

Banks are proceeding with caution on generative AI

Although a majority of employees have been exposed to generative AI either in their personal or professional lives, organizational strategies to implement the technology are proceeding slowly (see Figure 5). In all, only 37% of banks are actually in the process of implementing the technology. That movement is strongest among global and national banks with more than $100 billion in assets, of which more than half of organizations (55%) have initiated at least small-scale implementations.

Figure 5: Most banks are only investigating or pilot testing generative AI tools

Question: How would you describe the expected pace of generative AI adoption at your organization over the next 12 to 18 months?
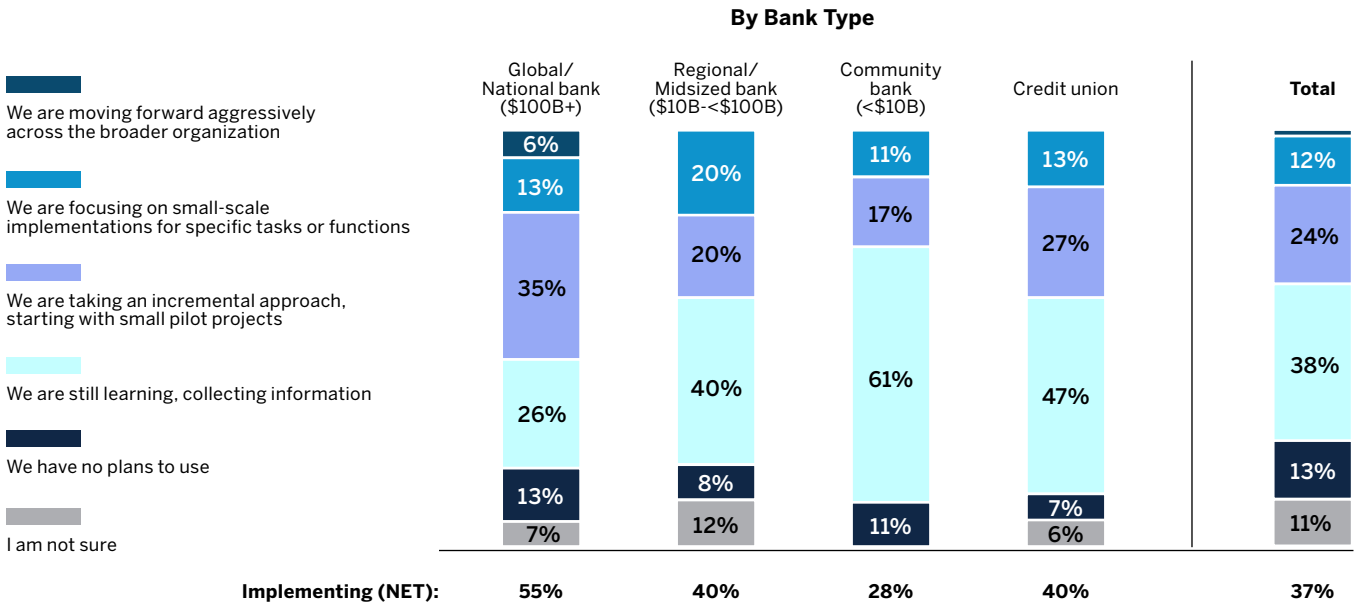
Source: Arizent, 2024 Unleashing the Power of AI
Labels for values less than 4% are not shown
Base: Total respondents n=127; Global/National bank ($100B+ in assets) n=31, Regional/Midsized bank ($10B to less than $100B in assets) n=25, Community bank (less than $10B in assets) n=18, Credit union (all asset sizes) n=15

This level of activity puts banking roughly in the middle of the pack among similar industries. Fintechs and insurers indicate they are more likely to be pilot testing or using generative AI than banks, while wealth managers and accountants are moving even more slowly on this front. This trend aligns with the general historical pace of technology adoption across these industries, with insurers and fintechs navigating significant competitive and cost pressures that offer more incentive for them to investigate new technologies that can improve margins or offer a differentiated customer experience.

There appears to be little consensus among bank leaders about how to position their organizations ahead of generative AI. Less than one-fifth are embracing or encouraging robust use of the technology. Meanwhile, about a quarter are approaching the technology slowly and cautiously and a few say they have no plans to use generative AI at all (at least in the next 12 to 18 months).
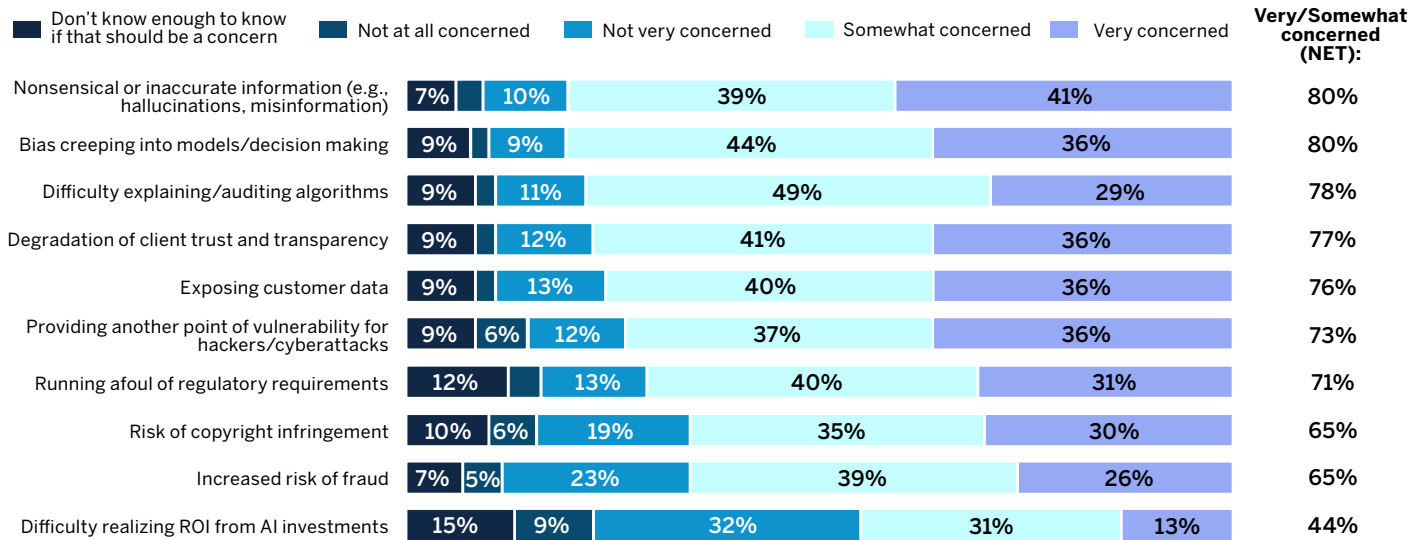
For roughly a third of organizations (35%), this caution has extended to banning or restricting the use of ChatGPT by employees. Global and national banks with assets of more than $100 billion are more likely to have these types of restrictions in place for now, which may reflect the fact that they are more likely to have an ongoing strategy in place to implement such tools. By most accounts, these restrictions will likely be temporary. Of those respondents who work at banks that have banned or restricted the use of generative AI, 62% expect restrictions to be loosened in the year ahead.

Generative AI introduces potential business risks

If ChatGPT demonstrated the potential of generative AI on the public stage, it also brought its flaws into sharp relief. The tool's potential to underline{misinform} commanded headlines and raised difficult questions.[1] Concerns about the provenance of the materials generative AI produces are prominent among bank professionals, with inaccuracies and ethical considerations topping the list (see Figure 6).

Figure 6: Banks have real and widespread concerns about business risks posed by generative AI

Question: How concerned are you about the following as potential risks to your business from using generative AI?

	Very/Somewhat concerned (NET):
Don't know enough to know if that should be a concern · Not at all concerned · Not very concerned · Somewhat concerned · Very concerned	
Nonsensical or inaccurate information (e.g., hallucinations, misinformation): 7% / 10% / 39% / 41%	80%
Bias creeping into models/decision making: 9% / 9% / 44% / 36%	80%
Difficulty explaining/auditing algorithms: 9% / 11% / 49% / 29%	78%
Degradation of client trust and transparency: 9% / 12% / 41% / 36%	77%
Exposing customer data: 9% / 13% / 40% / 36%	76%
Providing another point of vulnerability for hackers/cyberattacks: 9% / 6% / 12% / 37% / 36%	73%
Running afoul of regulatory requirements: 12% / 13% / 40% / 31%	71%
Risk of copyright infringement: 10% / 6% / 19% / 35% / 30%	65%
Increased risk of fraud: 7% / 5% / 23% / 39% / 26%	65%
Difficulty realizing ROI from AI investments: 15% / 9% / 32% / 31% / 13%	44%

Source: Arizent, 2024 Unleashing the Power of AI
Labels for values less than 4% are not shown
Base: Total respondents n=127

Addressing these areas will be critical for advancing generative AI implementations within banks. Notably, some of the greatest concerns revolve around the models themselves. Vendors or in-house teams that develop the AI models on which these tools rely will need to train and test the systems carefully to ensure accuracy. The ability to demonstrate which data informed a decision or output could play an important role in identifying and mitigating model bias. This type of information will also likely play a role in identifying and avoiding copyright infringement, a risk inherent to tools capable of producing new content.

[1] May 1, 2023, *New York Times*, "When A.I. Chatbots Hallucinate." https://www.nytimes.com/2023/05/01/business/ai-chatbots-hallucination.html

Risks such as exposing customer data or providing a point of vulnerability for cyberattacks will require banks to consider how they retrieve information from these systems as well as how they expose them to customers and staff. Data security will be a paramount concern, as will alignment with regulatory requirements.
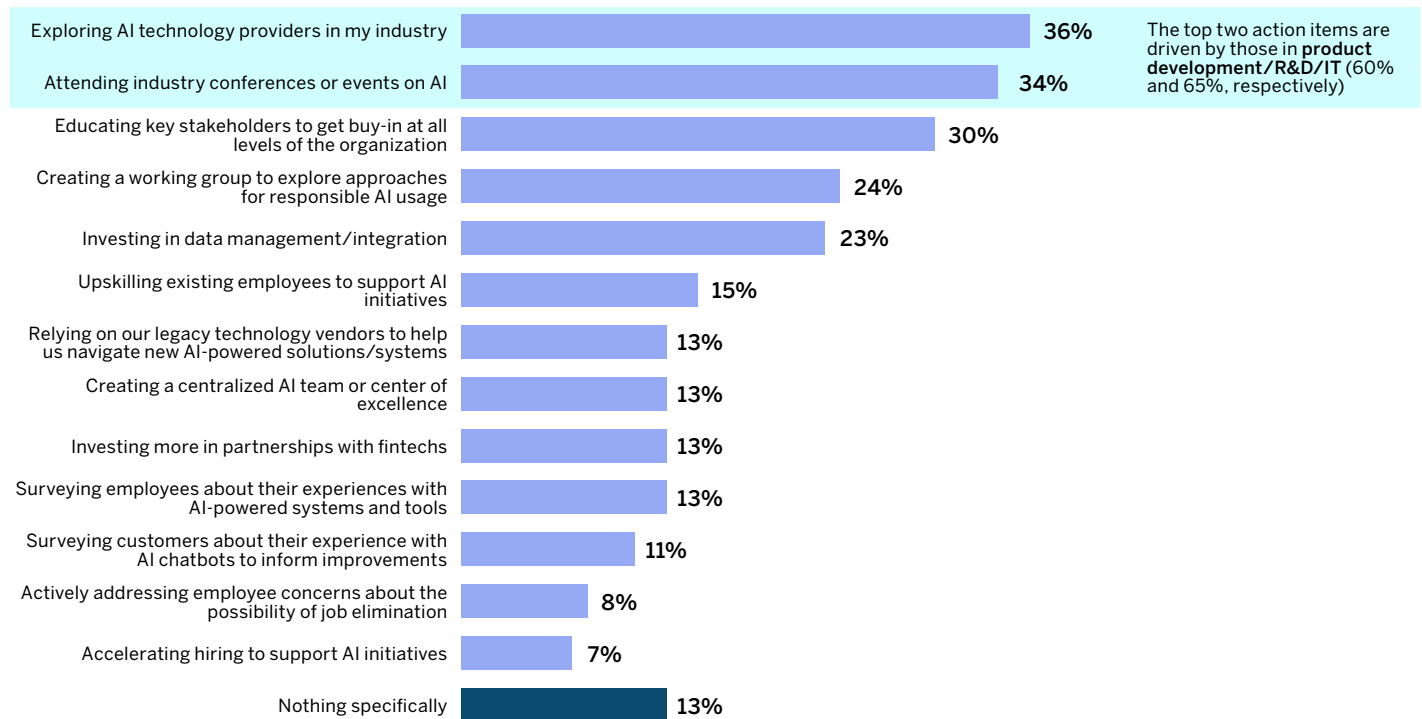
Banks appear to be looking for some degree of external help addressing many of these issues before they implement solutions. Three-quarters of respondents say they think international standards and stronger guardrails are needed to govern the use of AI in the industry. Organizations that are keen on AI adoption are especially likely to be focused on this need, suggesting that the establishment of widespread best practices for responsible use of generative AI could be a gating factor for accelerated adoption.

Setting the stage for the next generation of AI tools

The generally slow pace of adoption isn't an excuse for banks to remain complacent about generative AI-powered tools. Only 13% of respondents indicate that their organization has no plans to use generative AI. Those taking steps to advance their AI readiness are most commonly focusing on learning more about the technology so that they can strategize appropriately (see Figure 7).

Figure 7: Banks are still gathering information about generative AI

Question: What, if anything, is your organization doing to be more AI ready?

Action	%
Exploring AI technology providers in my industry	36%
Attending industry conferences or events on AI	34%
Educating key stakeholders to get buy-in at all levels of the organization	30%
Creating a working group to explore approaches for responsible AI usage	24%
Investing in data management/integration	23%
Upskilling existing employees to support AI initiatives	15%
Relying on our legacy technology vendors to help us navigate new AI-powered solutions/systems	13%
Creating a centralized AI team or center of excellence	13%
Investing more in partnerships with fintechs	13%
Surveying employees about their experiences with AI-powered systems and tools	13%
Surveying customers about their experience with AI chatbots to inform improvements	11%
Actively addressing employee concerns about the possibility of job elimination	8%
Accelerating hiring to support AI initiatives	7%
Nothing specifically	13%

The top two action items are driven by those in **product development/R&D/IT** (60% and 65%, respectively)

Source: Arizent, 2024 Unleashing the Power of AI
Base: Total respondents n=127; Those in a functional area of product development, R&D or IT n=20

Most organizations indicate they are gathering information, educating stakeholders and beginning to plot a course for the use of generative AI. More concrete steps such as training employees, engaging with technology vendors or creating a center of excellence are less common but are nevertheless underway in some institutions.

Very few banks indicate they are accelerating hiring to support AI initiatives at this time. Only 18% of respondents say their organization is willing to pay more for employees with AI fluency or skills at the moment. Those attitudes may be changing, however. Global and national banks are markedly more willing to pay for AI talent than smaller banks, possibly reflecting their head start on implementations. Furthermore, 30% of banks say that they are "not yet" willing to pay more for employees with AI fluency or skills, suggesting that they are open to the possibility in the future. C-suite executives are also more than twice as likely as their colleagues in other senior- or middle-management roles to leave the door open to future spending (56% vs. 25%), suggesting individuals with a longer strategic view are at least thinking about improving AI fluency within their organization.
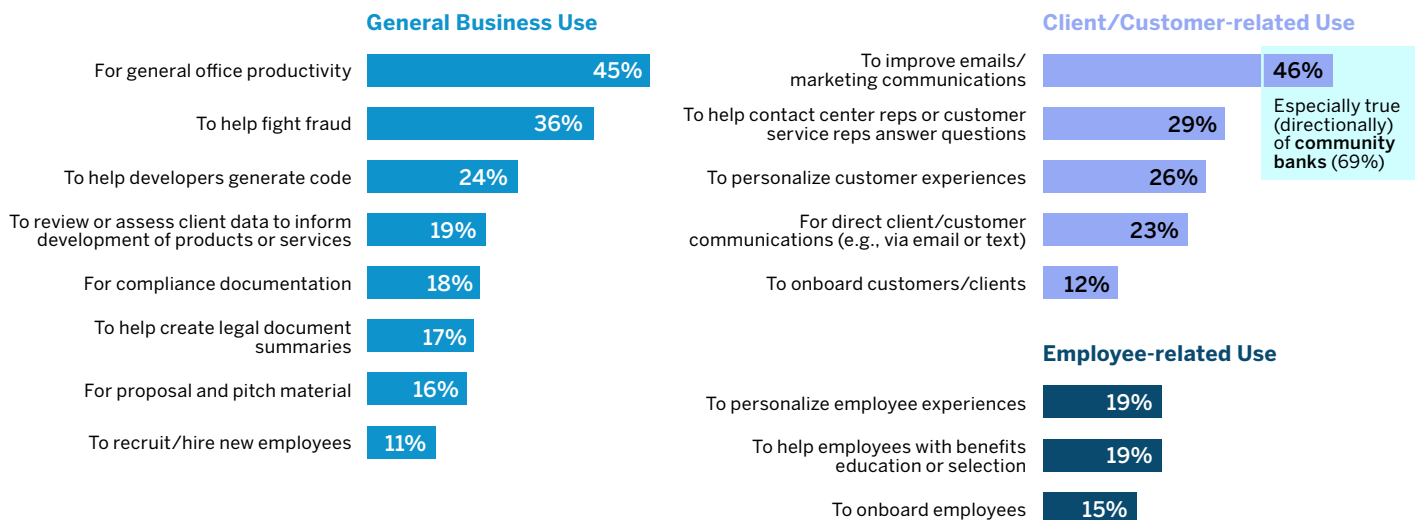
Waiting too long could become an issue, however. Insurers and fintechs are already signaling a greater willingness to pay more for AI-fluent employees. To the extent that their head start puts pressure on the talent pool, banks may find themselves competing for scarce talent by the time they ramp up their efforts. This situation could open the door to best-in-class vendors capable of providing solutions that address concerns about business risk and integrate smoothly with existing systems.

Looking ahead to the future of generative AI in banking

Over the near term, banks seem likely to focus their efforts on augmenting areas of the bank where AI tools already exist and where they can help employees do their jobs better (see the sidebar "The role of AI in the financial services ecosystem"). Early efforts have focused on applications that improve business processes, help improve the customer experience and help improve employee productivity (see Figure 8).

Figure 8: Early use cases are expected to focus on marketing, productivity and fraud prevention

Question: In what areas is your organization using or likely to use generative AI to support your business over the next 12 to 18 months?

General Business Use

For general office productivity	45%
To help fight fraud	36%
To help developers generate code	24%
To review or assess client data to inform development of products or services	19%
For compliance documentation	18%
To help create legal document summaries	17%
For proposal and pitch material	16%
To recruit/hire new employees	11%

Client/Customer-related Use

To improve emails/ marketing communications	46%
To help contact center reps or customer service reps answer questions	29%
To personalize customer experiences	26%
For direct client/customer communications (e.g., via email or text)	23%
To onboard customers/clients	12%

Especially true (directionally) of **community banks** (69%)

Employee-related Use

To personalize employee experiences	19%
To help employees with benefits education or selection	19%
To onboard employees	15%

Source: Arizent, 2024 Unleashing the Power of AI
Base: Total respondents n=127; Community bank respondents (less than $10B in assets) n=16

These potential applications offer banks a way to take advantage of generative AI's benefits while mitigating concerns. The use of AI tools to automate back-office tasks and fraud prevention is already well established. In those areas, models are likely to be trustworthy and already tuned to comply with regulatory requirements.

Expected client and employee use cases include a human in the loop, ensuring that generative AI tools augment a sense of personal touch in the client or employee experience rather than detract from it. The hesitancy in reducing or eliminating the human element in interactions with customers or clients extends to chatbots, widely considered an area of low-hanging fruit for generative AI-powered customer service. Where other industries have incorporated generative AI chatbots to enhance the customer experience, banks have been slow adopters of the technology to date. Only 24% are developing or implementing such tools, mainly global or national banks with more than $100 billion assets (45%).

Despite the slow uptake, banking leaders expect AI technology to continue to take over employee workloads over the medium term. On average, 9% of bank professionals believe AI can already do more than a quarter of the work they perform today. Further, 39% believe AI will be able to accomplish at least a quarter of their work within three to five years. In other words, banking professionals expect the amount of work AI is capable of doing within their organizations to grow considerably in the coming years.

Banking professionals are largely optimistic about those changes. More than three-quarters (76%) believe AI will redefine or change the nature of jobs in the industry without eliminating jobs on a large scale, while 36% believe AI will in fact produce new jobs or expanded opportunities. Optimism about job creation is most prevalent at global and national banks (52%), as well as more broadly at banks that are already implementing generative AI (51%). In other words, organizations that are the furthest along in their journey toward generative AI tend to see a rosier future than those with less experience.

The role of AI in the financial services industry

AI tools already support a variety of functional areas across the financial services industry, particularly when it comes to driving efficiency and automation, and fighting fraud. Here are some examples:

- AI tools capable of gathering information and comparing it to existing pattern-based rules are being used for automated decision making in areas such as small business lending and personal credit.

- AI tools trained on updated regulatory requirements are helping banks and credit unions manage business risks related to third-party vendors.

- Cybersecurity tools have long incorporated AI's ability to identify patterns to predict and identify fraudulent behavior, making it easier to mitigate cyberattacks.

- Pattern recognition is also helping banks detect, document and fight against potential instances of money laundering.

Conclusions

- Although ChatGPT has brought generative AI into the limelight, banks are moving cautiously as they navigate concerns about accuracy, ethics and the potential loss of the personal touch.

- A majority of banks are in the early stages of educating themselves about generative AI, forming strategies and piloting small-scale implementations. Generally, large national banks are leading the way on this front.

- The evolution of standards governing the responsible use of AI will likely play a role in increasing organizations' comfort level with the technology. Meanwhile, implementation efforts are likely to concentrate on areas where AI tools are already established or where a human can remain involved in the process.

- Banking leaders are beginning to position themselves for a future enhanced by generative AI tools. However, they will face competition for AI-fluent tech talent from adjacent industries such as insurance and fintech.

- Banking professionals generally agree that generative AI tools are likely to transform the workplace over the next three to five years in ways that are unlikely to result in large-scale loss of industry jobs.

About Arizent Research

Arizent delivers actionable insights through full-service research solutions that tap into its first-party data, industry SMEs and highly engaged communities across banking, payments, mortgage, insurance, municipal finance, accounting, HR/employee benefits and wealth management. Arizent has leading brands in financial services, including American Banker, The Bond Buyer, Financial Planning and National Mortgage News, and in professional services, such as Accounting Today, Employee Benefits News and Digital Insurance. For more information, please visit www.arizent.com.

Interested in learning more about how to put Arizent's full-service research capabilities to work for your company? Please contact: Janet King, Vice President Research, janet.king@arizent.com, 207-807-4806.