

AI in Financial Services

2026 | Produced by DALBAR.AI





Table of Contents

Introduction 3

Executive Summary 5

AI & Financial Services: A Brief History 6

DALBAR’s AI Research Findings..... 11

How the Industry Currently Uses AI13

Where Expectations and Practice Diverge17


Appendix.....19

Introduction

Spend ten seconds scrolling down your LinkedIn feed and “AI” will undoubtedly be mentioned in some capacity. Artificial intelligence has shifted from novelty to necessity in the daily lives of many Americans. It’s unavoidable at this point.

According to *TD’s AI Insights Report*, 78% of Americans use AI-powered tools on a daily basis. And why wouldn’t they? This technology can draft emails, plan trips, answer medical questions, and help people navigate virtually every corner of modern life. The cost to use it is low – even free in some circumstances – and it’s always accessible and ready to help.

Yet when the subject turns to personal finance, the comfort level seems to evaporate. DALBAR’s survey found that **only 38% of respondents have a positive outlook** when it comes to financial companies increasing their use of AI over the next few years. Furthermore, when asked “Are there any concerns that you have about financial companies using AI?”, a whopping 79% responded “Yes” (we’ll dig into what those concerns are later).





Clearly, consumers have already decided that AI is useful. However, they have not yet decided that it is *trustworthy* – at least not when it comes to finance. This trust gap is not a temporary hesitation that will resolve itself as consumers become more familiar with AI. It is instead gap that is rooted in concerns about **privacy, transparency, and the desire for human accountability**.

Financial decisions carry weight that other tasks do not. A wrong recommendation can mean a missed retirement goal, a denied loan, or a compromised nest egg. Likewise, irresponsible deployment can result in identity theft and privacy loss. These are very real concerns. Consumers instinctively understand this asymmetry. Their behavior, backed by data, reflects a level of caution that has not shown up in other industries adopting AI at a similar pace.

For financial professionals and institutions, this AI paradox is not a warning to slow down, but rather a call-to-action to **deploy AI in a way that respects consumer sentiment**. Firms and advisors who understand this distinction and build their AI strategies around closing this trust gap will have a real competitive advantage.



Executive Summary

This report's objective is to measure the distance between **what consumers expect** from AI in financial services, and **what the industry currently delivers**. That distance can be summarized in three consumer-side findings, three industry-side findings, and one sentence that captures the gap between them.

What Consumers Expect

In addition to analyzing external studies and findings, DALBAR conducted a consumer-facing survey about their opinion of AI in financial services. These findings suggest:

1. Consumers want disclosure far more than the industry currently provides.
2. Consumers are not rejecting AI outright but rather drawing sharp lines around specific functions.
3. Human review or override remains one of consumers' top comfort factors.

Where the Industry Stands

1. Institutional AI adoption is accelerating and concentrated internally.
2. Governance policy is outpacing governance enforcement.
3. Institutions are most mature where consumers already trust AI, and most cautious where trust is thin.

The Core Gap

Across every functional area examined in this report, one pattern holds: institutions have concentrated their AI investment in the internal background functions where consumer trust was never seriously in question. However, there seems to be hesitation in the disclosure, transparency, and consent infrastructure needed to earn trust in the AI applications customers can visibly see. **Trust is the gold standard when it comes to AI in financial services.**

The good news is that this trust gap is closeable. The recommendations that conclude this report translate these findings into specific, actionable steps for closing the trust gap.



AI & Financial Services: A Brief History

To understand why privacy, transparency and human involvement have become such flashpoints in today's conversation about AI in financial services, it helps to look at how the industry got here. This is not a general history of artificial intelligence, but a tightly scoped look at the specific developments that shaped how consumers came to feel the way they do about AI handling their money.


The Rules-Based Era

One of the financial industry's first widely visible encounter with consumer-facing AI came through robo-advisors. *Betterment* and *Wealthfront*, both founded in 2008, introduced automated investment platforms that built and rebalanced portfolios using algorithmic rules rather than human judgment; *Betterment's* product publicly launched in 2010, with *Wealthfront* following in 2011. *FINRA's* own 2016 *Report on Digital Investment Advice* characterized these platforms plainly: robo-advisors "currently largely use rules-based models to develop recommendations". These rules-based models made decision logic more constrained than today's generative systems. Given a set of inputs, they produced a predictable, explainable and consistent output.

Banks followed a similar path with customer-facing chatbots. *Bank of America* unveiled its virtual assistant *Erica* in 2016, which, after launching to customers in 2018, had surpassed six million users and 35 million customer requests by early 2019. *Capital One* piloted *Eno* in early 2017 before a full rollout that October, and *Citibank* followed with *Citibot SG* in 2018 for its Singapore consumer banking customers.


These tools were functional, but narrow. They could check a balance or flag a fraud alert, but they operated within scripted boundaries and did not reason beyond those boundaries. Compared with generative systems, early rules-based tools had narrower and more constrained failure modes. This predictability, however limited in scope, is part of why **fraud alerts remain the one AI interaction most Americans still feel positive about** today. When we asked consumers about their comfort level of financial companies using

AI, a minority (17%) of respondents were "Uncomfortable" with it being used to detect fraud and flag suspicious activity.



This same pre-COVID era saw AI expand quietly into the industry's back office. A 2018 *Chartis Research* study on AI in risk management found that roughly 70% of industry risk and technology professionals were already deploying AI in some professional capacity. FINRA's research into regulatory technology documented a similar trend toward AI-assisted compliance monitoring, anti-money laundering surveillance, and know-your-customer processes. None of this was visible to the average consumer. It operated behind the scenes, which is part of why, even today, roughly one in five Americans (21%) say they have never *knowingly* interacted with AI at a financial company at all.

WHEN YOU'VE INTERACTED WITH AI AT A FINANCIAL COMPANY, HOW DID IT USUALLY FEEL?



The Generative Shift

The rules-based model predictability shifted with the arrival of generative AI and large language models (LLMs).

Since 2018, the *Institute of International Finance* has conducted recurring surveys on machine learning and AI use in financial services, which have evolved into the *IIF-EY Annual Survey Report on AI/ML Use in Financial Services*. Their most recent survey finds that newer agentic AI and open-weight models are joining earlier generative AI and LLM capabilities, contributing to what the authors describe as “quicken adoption curves” for AI across the industry.

Unlike a rules-based chatbot, a generative model does not simply retrieve a scripted answer. It generates one, drawing on patterns learned from enormous datasets, which means its outputs are far more capable and conversational; but also far less predictable and considerably harder to explain after the fact.



This shift did not go unnoticed by regulators. The U.S. *Government Accountability Office's* May 2025 report to Congress on AI use and oversight in financial services found that AI adoption has expanded into automated trading, credit decisions, and customer service – delivering real efficiency gains while introducing new risks around bias, model transparency, and accountability. The report went so far as to recommend that the *National Credit Union Administration* update its model risk management guidance to keep pace with the broader range of models now in use.

Fallibility Consumers Can No Longer Ignore

The generative shift also came with a very public education in AI's capacity for confident error (commonly known as “hallucinating”). Recently, a Florida pastor sued *OpenAI* and CEO Sam Altman, claiming that *ChatGPT* repeatedly downplayed his worsening symptoms and discouraged him from seeking medical care. According to the complaint, *ChatGPT* recommended limiting movement. Doctors later attributed the nearly fatal pulmonary embolism to immobility.

And that isn't an isolated case. A separate 2026 lawsuit alleges a chatbot coached a 19-year-old into a fatal drug overdose, and a case documented in the *Annals of Internal Medicine* describes a patient who said he substituted bromide for chloride after following AI-generated dietary advice. This substitution allegedly led to bromide poisoning and psychosis.


Speaking of allegations, even the legal profession has not been immune. In one of the earliest and most cited examples, a New York attorney was sanctioned in 2023 after submitting a court filing built on entirely fabricated case citations that *ChatGPT* had generated with complete confidence. An ongoing public database now tracks over 1,500 similar court cases worldwide involving AI-fabricated material.

None of this happened in a corner of the internet hidden from consumers. These stories have run on network news, in national newspapers, and across social media. They have done more than embarrass the companies involved. They have taught the public something concrete: **AI can be wrong in a fluent, authoritative tone** that gives no outward sign of the error underneath. Consumers have absorbed that lesson from healthcare and

the courts, and it transfers naturally to money. If a chatbot can miss a life-threatening symptom or invent a legal precedent, it isn't a stretch for a consumer to wonder whether the same system might just as easily miscalculate a retirement projection or misjudge a loan decision. The caution consumers show toward AI in personal finance, then, is not simply inherited institutional distrust. It is a reasonable response to a technology the public has now watched fail in highly visible ways.


Why This History Matters

The common thread from rules-based robo-advisors to today's generative systems is a steady **trade of predictability for capability**. Early AI in financial services could do less, but both consumers and professionals could reasonably assume they understood how it worked. Today's AI can do dramatically more, inching its way closer to human-level reasoning and personalization, but it does so in ways that are frequently opaque even to the institutions deploying it, and the public has now seen firsthand what happens when that opacity goes ignored. That tradeoff is precisely what has pushed privacy, transparency, and human oversight to the center of consumer sentiment, and it is what connects this history directly to the findings presented in this report.



DALBAR's AI Research Findings






In the full version of this report, this section presents DALBAR's proprietary consumer research as the evidentiary core of our findings. The survey was fielded July 23 to August 2, 2026 among respondents ranging from 21 to 88 years old, with a median cluster in the 35 to 44 age bracket and a mean age of 46. It reflects a highly engaged financial consumer base: 98% hold a checking or savings account, 97% carry a credit card, 90% hold an investment or retirement account. Most respondents are also comfortable with digital finance management, with 73% reporting they prefer to manage their finances digitally rather than in person or by phone.



WANT MORE DATA?

The full report includes all of DALBAR's AI research findings

16 survey questions • 8 data charts • Complete section analysis



View our survey findings and gain deeper insight

- ✓ Consumer response analysis
- ✓ Supporting industry research
- ✓ Industry regulations surrounding AI usage
- ✓ Actionable steps to fill the AI trust gap



Get the full report at [DALBAR.AI](https://dalbar.ai) →



How the Industry Currently Uses AI

The previous sections captured what consumers believe, want, and fear from AI in their financial lives. Now let's explore what institutions are currently doing. Drawing on *FINRA's* review of AI applications in the securities industry, the *GAO's* May 2025 report to Congress on AI use and oversight in financial services, and the *IIF-EY Annual Survey Report* on AI Use in Financial Services, the picture that emerges is one of a heavily internal-facing, risk-managed rollout that has not yet caught up to consumer-facing ambition. Recent *IIF-EY* survey data indicate that 90% of financial institutions are now deploying AI in production, and among those, 84% have also implemented generative AI, reflecting rapid mainstreaming of GenAI alongside earlier AI/ML techniques.


Communications & Disclosures

Firms have deployed AI in customer-facing communications, chiefly through virtual assistants and chatbots. *FINRA* reports that several broker-dealers have already deployed conversational chatbots on their websites and apps, some are experimenting with third-party voice platforms like *Alexa* and *Siri*, and a few have gone as far as letting virtual assistants accept and process trade orders within set thresholds.

Bank of America's Erica chatbot has surpassed 3 billion customer interactions and now serves roughly 50 million clients as of 2025, and one industry estimate (based on a 2017 *Juniper Research* forecast) quantified chatbot savings at roughly \$0.50 to \$0.70 per customer interaction compared with a human agent.

But this is largely low-stakes, informational communication.

Institutions have stayed noticeably cautious about using AI for anything resembling a disclosure decision. According to the *GAO*, financial institutions currently limit generative AI to activities where a lower level of explainability is acceptable, and firm representatives have been explicit that they hold back where reliability or regulatory uncertainty is high. Meanwhile, *IIF-EY's* survey data show that internal use cases lead adoption, with 92% of



institutions using AI for internal knowledge access and 89% for process automation, while external customer-facing tools such as chatbots (49%) and sales and marketing applications (33%) trail behind.

Advice Generation


FINRA notes that digital investment platforms (robo-advisors) largely use rules-based models rather than generative AI to produce recommendations. Furthermore, firms have taken a cautious approach to letting AI offer investment advice directly to retail customers, partly because letting AI act more autonomously in advice-giving makes it harder to fulfill and demonstrate compliance with their fiduciary duties under the *Investment Advisers Act of 1940*.

Where AI *is* used in the advisory process, it is typically feeding a human advisor rather than replacing one. Some large firms are incorporating AI-based tools to build real-time, holistic customer profiles from account data, spending patterns, and prior communications. However, registered representatives use that output to augment their own judgment rather than deliver it to the client unfiltered.

The *GAO's* findings echo this. Robo-advisers are cited mainly for their cost advantage - offering investment advice at lower fees and with smaller account minimums - not for delivering more sophisticated or personalized advice than a human would. Current advice generation in the industry is about accessibility and cost, not about AI making autonomous financial judgments on a client's behalf.

Fraud Detection

Fraud detection and surveillance are, by a clear margin, the most mature institutional use case for AI. *FINRA* describes broad adoption of AI-based surveillance tools that monitor emails, social media, and text messages across an enterprise. This is moving compliance



teams beyond simple keyword searches and toward risk-based review that can pick up on tone, slang, or coded language.

An April 2018 *Chartis Research* report on AI in risk and compliance, produced in collaboration with *IBM Watson Financial Services*, highlighted early adoption of AI techniques in these functions. *Chartis* surveys found that around 70% of financial institutions were deploying AI in some capacity across risk and compliance, and that adoption has only accelerated since. AI is now also embedded in know-your-customer and anti-money-laundering programs – using machine learning, natural language processing, and biometrics to track suspicious activity more efficiently.

The *GAO* adds a credit-side data point worth noting alongside fraud: one AI provider told *GAO* its model cut the time and resources needed to make credit decisions by up to 67%, and credit unions using its AI model saw a 40% increase in credit approvals for women and people of color. AI is also reducing false positives in anti-money-laundering alerts, which the *International Monetary Fund* notes make up the vast majority of such alerts, freeing compliance staff to focus on genuinely suspicious cases.

Personalization

FINRA reports that a number of broker-dealers are exploring AI tools to analyze customers' investing behavior, website and app activity, and past inquiries in order to deliver customized educational content, market research, and sentiment analysis. Some firms already send this content via email or directly through their apps. At the same time, *FINRA* describes this largely as exploratory, with firms still determining whether individuals would be interested in specific services based on browsing history, rather than a fully scaled personalization engine.

The *GAO*'s example is instructive: one credit union uses AI to personalize customer service by recommending a member's frequent tasks, like transferring funds, when they interact with a chatbot.




Where Institutions Are Focusing Their Energy

Human-in-the-loop mechanisms are the single most relied-upon safeguard built into AI applications according to *IF-EY*, ranked first by the largest share of institutions among the safeguard types tracked, ahead of alert and log-monitoring systems, performance monitoring, and kill switches. Seventy-nine percent of executive committees are directly involved in establishing risk management frameworks for AI, and data confidentiality and privacy remains the top-cited AI risk at 72%.

Taken together, the industry benchmark shows firms moving fastest where AI reduces internal cost and risk, moving more cautiously wherever AI output reaches the client directly, and treating human oversight as the primary safeguard rather than a secondary one.


Where Expectations and Practice Diverge

With the industry benchmark now established, the divide between what consumers expect and what institutions actually deliver becomes clear when examined by function. In some areas the two sides are closely aligned. In others, the gap is wide enough to explain much of the trust deficit documented earlier.



The Pattern

Across all four functions, a single pattern holds: **institutions move fastest and consumers respond most positively wherever AI operates invisibly, protects rather than judges, and requires no new disclosure burden.** Everywhere AI becomes visible, makes a judgment about the customer, or depends on more personal data, consumer trust drops and institutional deployment slows – as in communications, advice generation, and personalization.



The gap in this report is not a simple case of institutions moving too fast or too slow. It is that institutions have concentrated their AI investment in the one place where trust was never really in question, while under-investing in the disclosure, transparency, and consent infrastructure needed to earn trust in the places where AI is most visible to the customer.



Appendix

This appendix consolidates every data source referenced throughout the report, organized by category, along with DALBAR's proprietary research base.

DALBAR Proprietary Research

DALBAR, "AI in Financial Services" consumer survey – fielded July 23 to August 2, 2026.
Internal, unpublished dataset that serves as the evidentiary core of this report.

Regulatory and Government Sources

- [FINRA, "AI Applications in the Securities Industry"](#)
- [FINRA, Rule 2210 FAQ, Advertising Regulation](#)
- [FINRA, Regulatory Notice 24-09](#)
- [FINRA, Regulatory Notice 21-29](#)
- [FINRA, Key Topics: Artificial Intelligence](#)
- [U.S. Government Accountability Office, GAO-25-107197, "Artificial Intelligence: Use and Oversight in Financial Services" \(May 2025\)](#)
- [International Monetary Fund, "AI and Regtech" \(October 2021\)](#)
- [Venable LLP, "AI in Financial Services: Popular Use Cases" \(February 2026\)](#)
- [Utah AI Act](#)
- [Colorado Consumer Protections for Artificial Intelligence](#)
- [California ADMT Regulations](#)

Industry and Institutional Surveys

- [Institute of International Finance and EY, "2025 IIF-EY Annual Survey Report on AI Use in Financial Services" \(October 2025\)](#)
- [KPMG Board Leadership Center, "2025 Survey: Boardroom Lens on Generative AI", 93 US board members surveyed January 21 to February 23, 2025](#)
- [Chartis Research and IBM, 2018 study on AI adoption among risk and compliance professionals, cited via FINRA's industry review](#)



Consumer Research and Public Opinion Studies

- [Consumer Reports, "Artificial Intelligence in Financial Services", nationally representative survey](#)
- [TD Bank, second annual U.S. AI Insights Report](#)
- [Wells Fargo, 2026 Money Study](#)
- [Attest, "AI in Financial Services: Consumer Attitudes"](#)
- [RFI Global, "Bridging the AI Trust Gap: The Key to Consumer Confidence in Financial Services"](#)
- [The Financial Brand, "AI in Banking: What Accountholders Want"](#)
- [YouGov, "Americans Still Don't Trust the Banking Sector's AI Use"](#)

News and Legal Coverage

- [Reuters, "ChatGPT's Advice Kept Man From Seeking Treatment for Dangerous Condition, Lawsuit Says" \(July 2026\)](#)
- [CBS News, coverage of the Scott Winters v. OpenAI lawsuit](#)
- [Reuters, "OpenAI Faces Lawsuit in California Court Claiming Chatbot Gave Advice That Led to..." \(May 2026\)](#)
- [The Guardian, "US Man Develops Rare Condition After Following ChatGPT's Health Advice" \(August 2025\)](#)

Historical and Background Sources

- [CB Insights, Betterment and Wealthfront Robo-Advisor Teardown](#)
- [Introducing Wealthfront, The Online Financial Advisor](#)
- [CBOT, "First Global Examples of Banking Chatbots"](#)
- [AI Hallucination Cases – Damien Charlotin](#)