FinRegLab

The Next Wave Arrives: Agentic AI in Financial Services

Market Scan

SEPTEMBER 2025

About FinRegLab

FinRegLab is a nonprofit, nonpartisan innovation center that tests new technologies and data to increase access to responsible financial services that help drive long-term economic security for people and small businesses. With our research insights, we facilitate discourse across the financial ecosystem to inform market practices and policy solutions.

Acknowledgments

This market scan is part of FinRegLab's ongoing work to inform both market practice and policy to promote the responsible and inclusive use of AI and machine learning systems across financial services. This work is supported from FinRegLab's general operating funds. Reports and information about our annual AI Symposium are available at our website, www.finreglab.org.

We would like to thank the many stakeholders in the technology, banking, fintech, research, civil society, and government sectors who engaged in interviews and conversations and provided feedback on this report. We would also like to acknowledge the FinRegLab team for their work in writing this report:

Jeanette Quick, Colin Colter, Luke Dillingham, and Kelly Thompson Cochran

When viewed with an Adobe Acrobat reader, elements listed in the Table of Contents or in blue text are links to the referenced section or feature. Functionality may be limited in non-Adobe readers. Adobe's reader can be downloaded for free at get.adobe.com/reader.

CONTENTS

1. INTRODUCTION: AGENTIC AI AS THE NEXT WAVE

Even as financial firms and commerce platforms wrestle with the implications of machine learning and generative AI for their businesses and customers, the next wave is already arriving. Agentic artificial intelligence (AAI)—a class of dynamic AI systems that can be structured to respond to new information and make and execute decisions without ongoing human engagement—has the potential to transform practically every layer of personal and institutional finance. Current and potential use cases include integrated real-time platforms for fraud and cyber defense, agentic shopping tools for managing online purchases, and personalized "financial agents" that can help households manage their daily finances and build long-term economic security.

Yet AAI also raises complex questions about how to realize these benefits while protecting users, providers, and the broader economy from potential errors and abuse. While some governance, data integrity, accountability, and fairness issues mirror those of other AI systems, AAI introduces distinct challenges for financial services and digital commerce. The prospect of autonomous agents executing large-scale financial decisions and transactions on behalf of consumers and companies heightens concerns about reliability and transparency, consumer protection, responsibility for errors, and risks to financial stability. Companies that may not feel ready to implement AAI systems themselves may nonetheless be impacted by the decisions of consumers or other companies to adopt agentic applications.

Managing the potential benefits and risks as AAI adoption spreads requires asking whether new approaches and tools will be needed as well as evaluating the applicability and utility of current market practices and regulatory frameworks for managing AI systems. Building on FinRegLab's previous research into AI and financial services, this report examines the current state of AAI, explores its potential impact on the financial ecosystem, and identifies emerging technology, market, and policy questions. Across stakeholder interviews, these issues emerged as the most urgent in shaping how AAI implementation will evolve:

- **Fostering the spread of effective technical tools and governance practices:** These are critical to giving consumers, small businesses, and financial services providers confidence that AAI applications can be trusted to improve their lives and operations.
- **Clarifying responsibility for error resolution:** Particularly given growth in agentic shopping applications, clarifying responsibility and liability for errors and disputes is important to consumers, merchants, financial institutions and AAI developers.

- **Promoting beneficial use cases:** Developing AAI systems that address the needs and challenges faced by consumers and small businesses who often struggle to access financial services could drive improvements in financial health and economic participation.
- **Building tools and safeguards to ensure that AAI apps serve users' interests:** Interface design and other mechanisms could help to ensure that consumers and other users can communicate their goals effectively and rely on AAI systems to act on their behalf.
- **Helping government keep pace:** Strengthening technical expertise, internal infrastructures, and cross-government coordination could potentially improve interactions with regulated companies and regulators deploy AAI to bolster their own efficiency and effectiveness.
- **Modernizing data and identity infrastructure:** Access to reliable data and better systems for differentiating "good" agentic traffic from bad actors could increase the quality and scope of AAI applications while helping to combat fraud and scams risk.
- **Bolstering monitoring mechanisms and safety protocols to protect financial stability:** As AAI applications increase in scale, companies and regulators may need tools that are calibrated to faster agentic activity.

While the size and scope of AAI adoption varies substantially from company to company and use case to use case, the technology is advancing regardless of whether policy keeps pace. Without a coordinated public-private approach to issues such as data infrastructure, model risk governance, and liability allocation, some consumers and companies may struggle to access the potential benefits and be at greater risk of suffering harms. Stakeholders may face an important window to align incentives, standards, and oversight before path dependence sets in, but interviews suggest that current awareness levels among different financial services stakeholders vary widely. Creating a common baseline understanding and decreasing uncertainty about critical issues could help to facilitate beneficial adoption.

To begin exploring these issues, Section 2 of this paper provides an overview of the technology and what differentiates it from prior generations of AI, while Section 3 describes illustrative financial services use cases both internal to financial services providers' operations and direct-to-consumer applications. Section 4 discusses challenges to adoption in financial services that are shaping the pace and nature of implementation relative to other sectors. Section 5 and Section 6 focus on potential consumer and systemic risk considerations as AAI adoption for financial services begins to scale. Section 7 identifies a series of key questions going forward, while Section 8 concludes the paper. The Appendix summarizes recent federal and state government activity that is not focused specifically on AAI but may have implications for its deployment.

2. OVERVIEW OF AI AGENTS AND AAI SYSTEMS

Simpler forms of automated agents have been used in some form for several years—for instance, trading bots that monitor market developments and execute investment transactions in certain circumstances act independently within defined boundaries—but interest in and adoption of AI agents have accelerated rapidly across the broader US economy, particularly in the last 18 months. For instance, venture capital funding of agentic applications has increased 150% year over year, and various tech CEOs, commentators, and publications have declared 2025 to be "the year of the agent."[1] To begin analyzing the momentum around agentic AI, this section describes and contrasts it with previous generations of AI.

Artificial intelligence is an umbrella term coined in 1956 to describe computers that perform processes or tasks that "traditionally have required human intelligence."[2] Certain predictive AI models have been used in financial services for decades, such as **supervised machine learning** models that are used for fraud screening or (more recently) credit underwriting. In 2022, **generative AI** models (genAI) that create new content such as text, images, or code became a major focus of public attention with the advent of ChatGPT, which incorporated a large language model (LLM) and other elements that went substantially beyond previous generations of chatbots. **Agentic AI,** in contrast, refers to systems that are designed to interact with their environment, process new information, and perform self-directed tasks in pursuit of high level goals—potentially without substantial human intervention beyond setting the initial objectives.

More specifically, AAI systems combine multiple "agent" software programs that can each perform different tasks—such as gathering data from internal and external locations, moving funds, or writing and deploying new code—to achieve specific goals. The systems frequently also incorporate other types of AI, such as LLMs, and can interact with various tools and external data sources through the model context protocol (MCP), which is an open-source, standardized communication framework to help agents connect securely and efficiently with external resources and programs.[3] These systems can be designed with varying capabilities and constraints depending on individual companies' goals and risk tolerances (see Box 1), but have the potential to be structured to respond dynamically and take autonomous actions in response to changing conditions in a way that is substantially different than other types of AI that are commonly used in financial services.

For example, supervised machine learning models that are used to screen for fraudulent transactions or predict the likelihood of credit defaults are generally static, requiring developers periodically to initiate a "refresh" process to update the models based on recent data. They are also generally structured to perform narrow tasks such as producing a risk score or a binary classification such as approve/deny in a consistent manner where the same inputs will produce the same outputs each

BOX 1 GRADATIONS IN AAI SOPHISTICATION

AI agents and AAI systems can be grouped along a continuum of functional capacity, autonomy, and complexity depending on how they are designed to operate. Individual agents may be relatively simple software designed to perform very straightforward tasks in isolation (such as retrieving data from one source and moving it to another source) or much more complex tasks (such as determining where external circumstances warrant reprogramming a credit scoring model and then overseeing the process). In addition, multiagent systems may involve hundreds or thousands of such agents working together to perform more complex tasks on behalf of a human user or another system. Many but not all multiagent systems may have an **orchestration layer** or **master agent** that engages and disengages groups of other agents to perform specific functions where warranted.

Different agents may have different degrees of dynamism and autonomy:

- **Reactive agents:** Are programmed to respond to specific commands or stimuli but do not retain information or have the capacity to plan.
- **Proactive agents:** Have some degree of autonomous capacity to adjust to evolving circumstances and form strategies to achieve goals.
- **Adaptive agents:** Are designed to learn dynamically and respond to changing environments to improve performance of tasks continuously over time.

AAI systems can combine a range of agents as well as other types of models (such as LLMs and supervised machine learning models), classical AI planning and action frameworks, and other techniques to decompose tasks, achieve multiple objectives, and perform complex functions over longer horizons.

FIGURE 1 SAMPLE CONCEPTUAL ARCHITECTURE FOR AN AAI SYSTEM

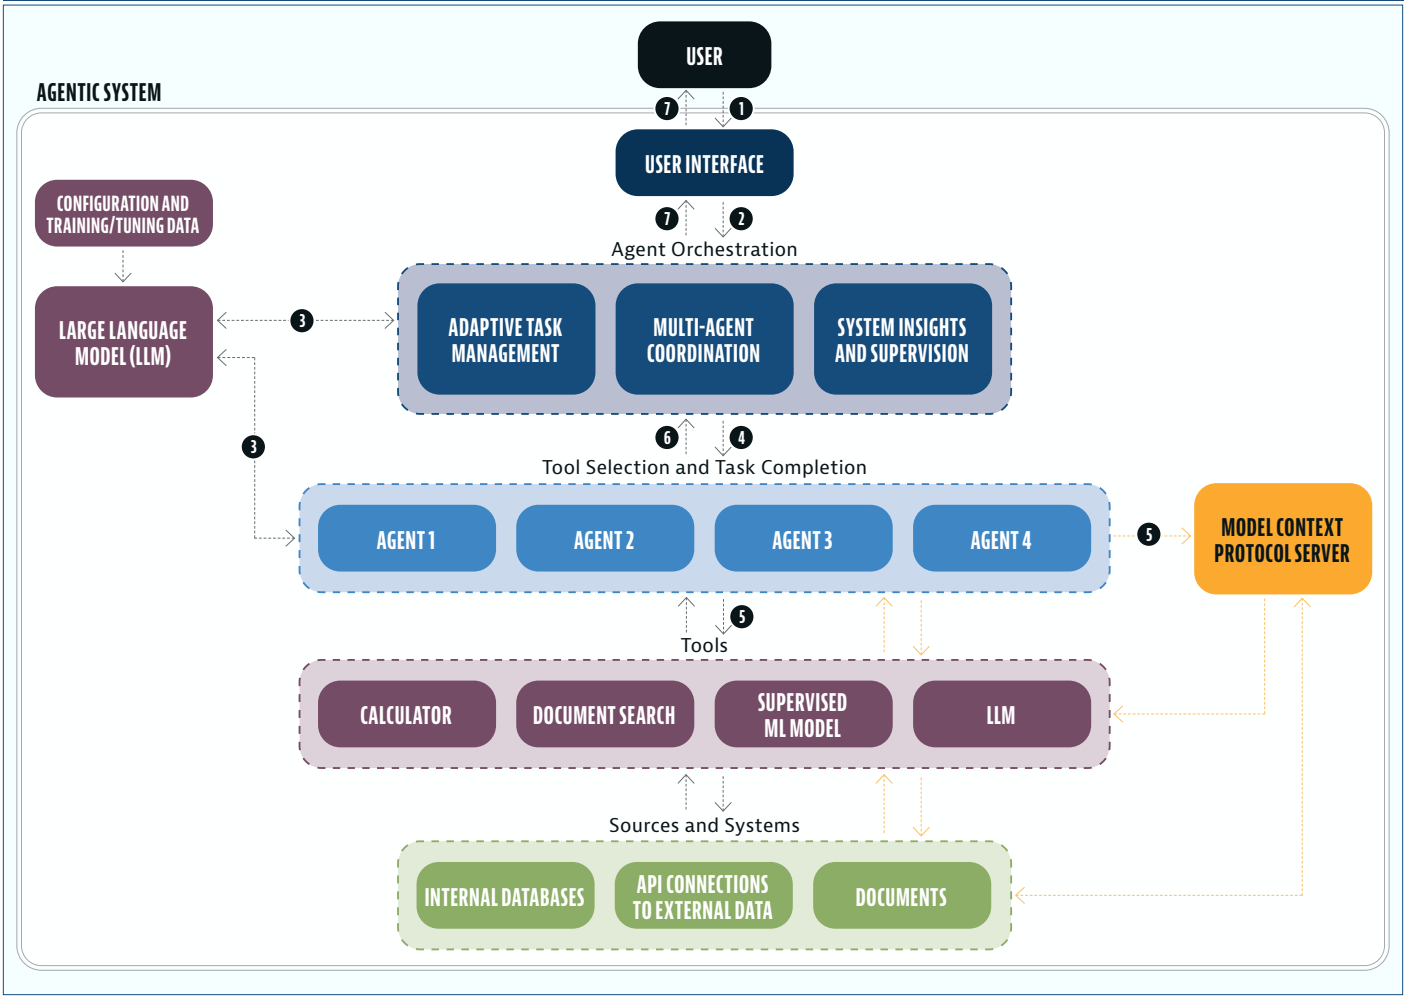

time.[4] Standalone GenAI applications are designed to perform a broader range of tasks and can be structured to draw on supplemental data supplied after initial training. They are also often described as "probabilistic" or "stochastic" because they sample from probability distributions in estimating outputs, which can create some variation in response to the same prompt and inputs.[5] However, such applications still need specific human prompts and refinements to generate content in the first instance, and in part because of concerns about "hallucinations," inconsistent answers, and other problematic outputs, many financial services providers have initially deployed these technologies through "copilot" structures that are designed to actively assist employees in performing tasks—such as by offering suggestions or generating rough drafts of code or text—while leaving the final actions and sign-offs with human operators.

By using multiple software agents (often in combination with multiple LLMs or supervised machine learning models, depending on the range of tasks) AAI systems offer greater capacity to execute against high level goals with less need for ongoing human intervention. The systems can be structured to perform web searches and retrieval augmented generation to use new information beyond their original training data to answer user questions and perform other tasks.[6] Reasoning frameworks such as ReAct (short for reasoning and acting), chain of thought, and tree of thought can also be used to help the systems decompose complex tasks, adapt to new inputs and circumstances, and determine next steps.[7] Some agents are also able to operate the user interfaces of other software applications to access data on behalf of users. While some of these methods can also be used with standalone genAI applications, AAI systems' ability to use tools, break down tasks, and gather new information and incorporate it into response or content generation creates the potential for agentic systems to perform a much broader range of activities than standalone GenAI chatbots or supervised machine learning models.

At the same time, some of these same qualities can potentially make AAI systems more challenging to test, understand, and tune to achieve an acceptable level of performance on given tasks. While the systems can be designed in a variety of different ways, validating each individual component in isolation may not provide a full picture of how agents will interact with each other, respond to external developments and actors, and leverage new information over time. Guardrails, such as constraints to restrict the function of individual agents and ongoing monitoring and troubleshooting, take on heightened importance in this context.[8]

While these descriptions help to convey some of the core concepts, it is important to note that some financial services providers are focusing first on building relatively discrete AI agents to perform narrow tasks, that integrated AAI systems can be structured in different ways, and that combinations of other types of AI and automation can be used to perform some of the same functions. Different stakeholders may also use different terminology in different settings given substantial marketing and investment buzz around agentic AI, and especially in the context of co-pilot systems and financial advice chatbots the term "agent" is sometimes used to refer to systems that assist users by delivering advice or other content but do not in fact have the capacity to take actions on their behalf.

All of these factors can make it challenging to gauge the state of actual AAI adoption for individual use cases and firms. But beyond general technology companies' heavy emphasis on supporting AAI applications for both financial and other use cases (see Box 2), interviews confirm that the technology's potential is sparking strong interest among larger and more tech-forward financial services providers. Even as many financial companies are still working to adopt and manage machine learning and genAI technologies, some are exploring and implementing AAI use cases as described in Section 3.

BOX 2 FOUNDATIONAL MODEL BUILDERS AND OTHER GENERAL TECH COMPANIES PIVOT INTO AAI SYSTEMS SUPPORT

Companies that have led the development of general purpose generative AI models (sometimes called foundation models) are increasingly providing agentic AI capabilities to companies and consumers who use their systems, with varying degrees of focus on financial activities specifically. Several new programs were announced in spring and summer 2025, in addition to e-commerce specific initiatives discussed in Box 3:

- Anthropic created a set of AAI tools to enable financial advisory firms to amass and analyze financial data. The Claude Financial Analysis Solutions helps clients pull from internal and external data sources to conduct due diligence and market research, competitive benchmarking and portfolio analyses, financial modeling with audit trails, and create investment memos and pitch decks.[9]
- OpenAI launched ChatGPT Agent, which is a generalized suite though some of its highlighted use cases include analyzing corporations' finances and helping consumers to find and purchase goods and services. Capabilities include web search functions via both graphic user interfaces and text browsers that livestream its actions to facilitate users' feedback and follow up directions, and connecting to calendar software and other programs. Actions with significant consequences (such as making a purchase) require user approval and the system is trained to refuse certain high-risk tasks including bank transfers.[10]
- Mistral announced its Agents API initiative, which offers the ability to combine its large language models with connectors to web search and other tools, persistent memory across sessions, and orchestration capabilities. Financial analysis is one of the highlighted use cases.[11]
- Meta announced it is working to make AAI systems available to small businesses that use WhatsApp, Facebook, and Instagram for marketing and interacting with customers.[12]

In addition, cloud providers Microsoft Azure, Amazon Web Services, and Google Cloud announced new suites of services to help their customers build and deploy AAI systems at scale.[13] SalesForce and a number of more specialized agentic AI platforms are also offering libraries and other frameworks to help companies expedite AAI development.[14]

Although the foundation model providers vary in the extent to which they are pitching their genAI and agentic AI products as directly usable by financial services providers and consumers for complex financial tasks, many financial services providers are building more specialized applications on top of foundation models or other generalized infrastructures. Testing of foundation models' performance on various financial-related tasks finds substantial variation in accuracy, cost, and speed. For example, when tasked with analyzing 200-page commercial credit agreements, the highest performing model had a 73.5% accuracy rate, while accuracy rates peaked at 48.3% when models were tasked with open-ended questions expected of entry-level finance analysts.[15]

3. POTENTIAL BENEFITS AND USE CASES

In considering the potential benefits of AAI adoption for different financial services use cases, it can be helpful to separate those that are largely provider-driven from those that depend largely on consumer adoption.[16] While implementation of the latter use cases is not generally as advanced today, e-commerce related applications are accelerating rapidly and others could have substantial benefits for financial inclusion and consumers' long-term financial security if they can be implemented successfully on a broad scale. This section considers each group of use cases in turn.

3.1 Provider-driven use cases

Our interviews with subject matter experts identified a number of concrete examples of ways that AAI is already being deployed by financial services providers, largely for internal operations purposes although in some cases these AAI systems may also interact with the providers' customers. Many tools have been or are being built to automate financial crime detection, antifraud, and compliance workflows in ways that leverage AAI's abilities to transfer, aggregate, and analyze data and to respond quickly to evolving challenges, while still maintaining humans in the loop to varying degrees. Other use cases include using AAI systems to manage credit and insurance underwriting and customer service workflows across a range of financial services.

- **Financial crime, antifraud, and cybersecurity processes.** A number of companies are moving to AAI systems to expedite account opening, transaction monitoring, and other activities relating to fraud defense and compliance with federal laws concerning money laundering and other financial crimes.[17] Similarly, cyber detection functions are also looking to AAI systems.[18] Compared to legacy screening and automation technologies or standalone AI models,[19] integrated AAI systems offer substantial potential benefits in identifying and responding in real time to emerging bad actor patterns,[20] increasing the accuracy and consistency of screening and follow up reviews, handling wide variations in volume, and testing systems through "red teaming" and other activities. These AAI systems typically continue to use humans to handle more complex cases and provide general oversight, while increasing automation to perform or support various workflows. In light of the increasing speed of payments systems and expanding use of various forms of AI by bad actors to fuel frauds, scams, and cyberattacks,[21] many stakeholders believe that more agile systems will become increasingly critical to contain costs and increase effectiveness.
- **Other compliance and risk management tasks.** AAI systems can similarly be used to support various other types of compliance and risk management processes by capturing,

aggregating, and assessing various internal and external data sources. For example, agentive systems could be used to assess a counterparty in a commercial lending transaction for reputation risk, supplier concentration, revenue concentration, and other potential concerns. Other tasks include identifying new patterns that may warrant human review or risk mitigation measures, such as reviewing historical transactions to determine whether credits and debits match industry norms; developing and executing mitigation measures such as freezing accounts or alerting relationship managers; and performing audit and testing functions like generating compliance documentation. AAI systems could also potentially be used by regulators to support various supervisory and monitoring activities, although such use cases are likely to take longer to develop than private industry applications.

» **Internal treasury management, investment strategies, and deposit management.** AAI has the capacity to help financial services providers and other corporations manage liquidity and capital. For example, AAI systems offer the potential to improve the management of investment portfolios compared to legacy processes through autonomously monitoring market trends, deciphering trading signals, adjusting investment strategies, and mitigating risks in real time by executing trades or taking other actions. Similarly, AAI systems could be used to monitor and move deposits between financial institutions in arrangements where reciprocal sweep accounts are used to make sure that deposits do not exceed FDIC insurance limits.

» **Underwriting-related processes.** Some stakeholders are focusing on AAI systems' potential to structure underwriting workflows in both lending and insurance, for instance by working with applicants to make sure that they have properly filled out forms and downloaded documents, organizing and analyzing application data, triggering and facilitating the updating of underwriting and pricing models, and engaging with applicants to communicate underwriting decisions and complete origination processes.[22] These systems can vary significantly as to the extent to which they integrate AI agents to perform various discrete tasks, predictive underwriting models to perform the primary predictive assessments, and human staff to make final decisions and monitor the broader systems.

  If such systems prove sufficiently efficient and reliable, some stakeholders believe they could support micro-sized and other niche products that would be too expensive to offer using more traditional human-based processes, such as microloans or real-time micro-insurance products for smallholder farmers based on farm history, conditions, and local data on weather or other external factors.[23] AAI systems could also prove instrumental to parametric insurance programs, which pay set amounts based on event parameters rather than the extent of damage suffered by individual policyholders.[24]

» **Individualized customer service, claims adjustment, and loan work out programs.** Across lending, credit, and other financial services, stakeholders also see potential use cases focusing on customer inquiries, claims adjustment, servicing and loan workout programs, and other downstream interactions with consumers. They envision AAI systems as bridging across traditional departmental divisions and data silos within banks to provide one-stop service, for instance by handling inquiries about credit limits that may lead to adjustments in a consumer's accounts or proactively identifying where consumers may have trouble making upcoming loan payments, offering a range of options, and executing transfers or other actions at the consumer's direction. Other potential use cases include processing consumers' requests for emergency loan deferrals or property damage claims in the aftermath of natural disasters. AAI systems' abilities to retrieve and process information, respond to changing circumstances, and provide consistent 24-7 service even in surge situations is particularly appealing in these contexts.

3.2 Consumer-driven use cases

Another set of use cases for AAI systems will depend in part on consumers' affirmative decisions about whether and how to deploy agentic tools for purposes of executing individual transactions, interfacing with financial services providers in connection with funds and data transfers or other matters, or managing their general personal or household finances more generally. Some consumers already use algorithm-driven savings tools, personal financial management (PFM) apps, "robo advisors" and other technology-driven services for some of these purposes, but AAI systems can potentially provide more comprehensive, personalized service and more active assistance in executing decisions than prior generations of tools. Similarly, AAI systems have significant potential to help time-strapped small business owners manage their ongoing finances.[25] E-commerce applications are already expanding (see Box 3), and a number of broader financial management applications are in development.

BOX 3 RECENT E-COMMERCE INITIATIVES AND PARTNERSHIPS

A number of e-commerce companies announced initiatives in spring and summer 2025 to use AAI to go beyond automated search functions to begin assisting with executing sales transactions.[26] While these currently appear to require consumers to consent to individual purchases before they are executed and are focused on more discrete transactions than a general personal financial agent tool, they also may implicate questions about responsibility for situations in which an agentic system malfunctions, causing financial losses or other harms to a consumer as discussed further in Section 5.

- In March, the web browser Opera announced an AAI system that can search third-party websites and make purchases and reservations.[27]
- In April, Amazon announced it was testing a new "Buy for Me" AI assistant that will review items and make purchases both on its own platform and from third-party websites without the user needing to leave the Amazon webpage.[28]
- In May, Google announced "AI Mode" shopping tools to search product listings, track price fluctuations, and execute purchases, as well as to provide virtual "try on" technology that incorporates consumers' uploaded photos.[29]
- In May, Perplexity announced it was partnering with PayPal and Venmo to use AAI to access users' wallets and add payment capabilities to streamline transactions.[30]
- In June, Walmart released the AI shopping assistant Sparky that will synthesize reviews and provide product recommendations to shoppers, with future features to include automatic reordering of staples and booking services.[31]
- In July, reports indicated that OpenAI and Shopify were working together to develop a checkout option that would allow purchases to be made directly via ChatGPT, further expanding its current offerings on shopping recommendations and reviews. Merchants will be expected to pay a commission to OpenAI on the transactions.[32]

In addition, both Mastercard and Visa have announced services to help e-commerce platforms, merchants, and other AI developers facilitate agentic commerce applications.[33]

At the same time that many e-commerce platforms are developing their own agentive offerings and partnerships, they are also taking steps to protect their websites from being overwhelmed by outside shopping agents. For example, Amazon updated its policies to ban agents from engaging in behavior that is designed to evade measures to weed out bots and to reserve the right to bar or limit agents' interactions with its services at its sole discretion.[34] Shopify has also added warnings in its code to power merchant storefronts that states "Automated scraping, 'buy-for-me' agents, or any end-to-end flow that completes payment without a final review step is not permitted."[35]

More broadly, merchants and platforms are trying to think through how they will promote visibility to legitimate shopping agents that may not be sensitive to marketing content and branding strategies that appeal to individual consumers while also refreshing fraud models and other systems to better distinguish between legitimate and bad actor traffic. Many commentators describe the rise of agentic commerce as the biggest disruption to online sales since the rise of large marketplace platforms, but it is not yet clear how quickly consumers, merchants, and other actors will adapt to these new systems and capabilities.[36]

- **Better tools for managing general cash flows, savings, and expenditures.** Some stakeholders highlighted AAI systems' potential to monitor account balances and upcoming obligations and execute transfers as needed to help optimize for covering expenses, avoiding overdrafts, and maximizing use of higher-yield savings vehicles. Similar services on the expenditure side include AAI-driven subscription and bill management tools designed to identify and execute on expense reduction strategies and leverage the most advantageous methods for paying for or financing individual purchases.
- **Debt repayment, management, and ongoing credit score improvements.** AAI systems can potentially provide a wider range of services than simple tools that help consumers determine how much to pay toward multiple debts, though research suggests that even simple tools can improve consumer outcomes.[37] AAI systems have the potential to make and execute payments decisions consistent with consumers' financial goals, monitor and pursue advantageous refinancing opportunities, and coordinate with multiple creditors to seek forbearances or workouts. They can also monitor credit scores and recommend and execute on strategies to help build more robust credit records over time. At a time when some consumers are managing a range of buy now, pay later purchases as well as multiple traditional credit accounts, AAI systems could help consumers manage all of their credit sources productively.
- **Optimizing management of government and workplace benefits to strengthen consumers' financial wellbeing.** Government and workplace benefits can have significant effects on consumers' day-to-day and long-term finances, but can involve complicated application processes and eligibility rules, separate information platforms, and other information and process challenges. AAI systems could potentially help consumers evaluate, apply for, and manage various benefits, as well as to combine benefits data with other financial data sources to make it easier for consumers to see and manage their full financial lives.
- **Other investments and insurance management.** A broad range of investment platforms already provide limited automated robo advisor options such as portfolio rebalancing, and some digital investment apps are experimenting with more broad-based genAI assistants.[38] Some stakeholders view AAI systems as a significant "leapfrog" opportunity to provide the kinds of highly personalized advice, broader goal-based coaching, and execution assistance that have historically only been available from human advisors to wealthier consumers. Proponents point to both reactive (e.g., investing tax refunds or one-time windfalls) and proactive services that AAI systems could offer to LMI consumers, both before and after the point of retirement.

Across these examples, stakeholders emphasize not only the potential benefits that AAI systems could offer in obtaining financial services and executing transactions, but also in helping consumers access, use, and manage related data that may be spread across multiple companies. Customer-permissioned data transfers have mushroomed to fuel PFM tools as well as payments apps and loan underwriting, with an estimated 100 million consumers having authorized third-party access to their financial services data by 2022 as discussed in Section 4.[39] AAI systems could potentially help consumers monitor and manage data sharing on an ongoing basis, in addition to their direct finances.

However, these kinds of customer-selected AAI systems are also complicated to design and launch in light of questions such as how financial institutions' systems will interact with them, how responsibility for errors will be handled, and other customer protection concerns as discussed below. As a result, general personal financial agents may take more time to evolve and scale than other AAI use cases. However, recent e-commerce AAI initiatives may force stakeholders to grapple more quickly with some of these questions.

4. CHALLENGES TO AAI ADOPTION

The potential benefits of AAI systems to both financial services providers and their customers are driving substantial interest particularly among very large financial institutions. However, AAI applications that potentially involve larger amounts of financial risk and more broad-based consumer-facing financial management tools face a number of significant challenges to responsible and effective adoption. These obstacles are not solely technical but also span operational, regulatory, and cultural domains within providers and the broader financial ecosystem and are likely to shape the pace and breadth of AAI adoption in financial services relative to other sectors.

First, while AI tools themselves are advancing rapidly, the underlying infrastructure has not yet been upgraded to support them effectively, especially within many depository institutions. The benefits of AAI adoption increase substantially where developers can provide real time access to clean, consistent data and map functions holistically to facilitate the integration of multiple types of AI, technological safeguards, and human review. Both of these tasks can be particularly challenging for depository institutions, which are more likely than AI startups or digitally native fintechs to rely on manual processes, legacy core data platforms, and siloed systems.[40]

Automated programming interfaces (APIs) and other "open banking" infrastructure can be one way of providing connections both within individual institutions (as a way of assembling data across legacy systems after a merger, for instance) and between institutions (where consumers have directed such transfers to occur). This infrastructure has developed over the past 30 years to support the provision of a range of financial products and services, including financial advice and investment tracking, payments applications, and credit underwriting. However, uncertainty over recent regulatory initiatives and steps by some companies to begin charging for data collection by other financial services providers are raising questions about how this ecosystem will evolve going forward,[41] even as AAI systems are likely to increase the demand for data pulls (see Box 4).

In the absence of real-time data connectivity, AI agents may be forced to make decisions based on incomplete, stale, or inconsistent information. This not only diminishes the value proposition of AAI but also increases the risk of consumer harm, as financial decisions made without up-to-date context may expose consumers to unnecessary costs, missed opportunities, or avoidable risks.

BOX 4 OPEN BANKING, APIs, AND THE EVOLUTION OF AAI

Open banking architectures in the United States have historically relied on a group of intermediaries called "data aggregators" to transfer customer-permissioned data from the initial source to another financial service provider to support various use cases such as payments, loan underwriting, and personal financial management and investment tools.[42] Initially, aggregators requested consumers' log in credentials and then used "screen scraping" to collect the data from the data sources' internet platforms, but in recent years the ecosystem has increasingly migrated to using tokenized access credentials that do not give full access to the underlying accounts and to APIs to transfer data. Bilateral contracts between the data sources and aggregators govern the frequency of data access, security measures, and various other details.[43]

As of 2022, the Consumer Financial Protection Bureau estimated that at least 100 million consumers had already authorized third party access to data about retail financial services that they had obtained from other providers, with annual traffic reaching between 50 to 100 billion successful or attempted transfers.[44] Some banks report that aggregators' monthly data pulls doubled between 2023 and 2025, though experiences may vary substantially based on consumer demand and how banks have structured their APIs.[45] Recent debates over whether banks should be allowed to charge for access to consumers' data have hinged in part on concerns about the costs of building and managing APIs, other information security infrastructure, and general risk management.[46]

The spread of AAI systems could have a number of implications for this ecosystem, including increasing the demand for data pulls, complicating the process for data sources trying to distinguish between legitimate and bad actors, and changing which parties engage in data pulls. Where AAI system builders seek to use APIs for data access, whether directly or through a MCP server,[47] they are likely to have to either work through data aggregators that have contractual relationships with the data sources or contract directly with the banks or other financial services providers. These contracts impose some restraints on the number of pulls per day and other practices. Screen scraping is not subject to contractual limitations but is less reliable in terms of data access and quality (as well as having various other disadvantages for consumers and data sources).[48]

Some aggregators appear to be using AI agents to facilitate data collection and scraping, and some genAI developers also use it to collect publicly available internet materials for use in model training and app development.[49] Other companies are developing new technologies to try to prevent scraping their public-facing websites without permission or compensation.[50]

Second, customer identity verification requirements and broader fraud and financial crimes prevention systems present potential obstacles and limitations with regard to what activities AAI can perform. "Know Your Customer" laws generally require the active participation of consumers and small businesses to verify their identities as they open new financial accounts. Some level of authentication is also typically required when logging on to financial institutions' computer systems and engaging in individual financial transactions, in part to guard against fraud and hacking as well as to facilitate compliance with anti-money laundering laws. Finding ways to both distinguish between rapidly increasing volumes of "good" agentic traffic from bad actors and to facilitate legitimate AAI systems' ability to act on consumers' behalf could require substantial investments and adjustments by many financial institutions and may raise questions about whether adjustments in legal requirements are warranted. The ways in which these issues are addressed will help to determine the extent to which AAI systems can help consumers open new accounts as well as help transferring funds, executing payments, and engaging in other individual transactions.

A number of larger financial institutions are focusing on the use of tokens that can be used to verify that information requests or other financial activities have been authorized by a consumer, without having to reenter the consumer's underlying credentials and go through multi-factor authentication for each interaction.[51] However, while development of these systems has expanded generally in recent years and some players are now developing agentic-specific tokenized applications,[52] they are not well established yet particularly for broad ongoing use cases rather than short-term applications. Developing more uniform and secure standards for tokenized identity verification would facilitate AAI's ability to work effectively across different institutions.

Third, particularly in light of the importance of regulatory compliance in financial services, the newness of the technology for both industry stakeholders as well as regulators poses meaningful challenges to AAI adoption. Financial services is a heavily regulated industry, with a complex mix of laws and regulations that apply to different financial companies for purposes of promoting the stability of the nation's financial system, competition within individual markets, the safety and soundness of certain financial institutions, and consumer protection. Both industry participants and regulators are starting to assess how existing frameworks apply to AAI applications and how technology, market practices, and regulatory tools may have to adapt to manage emerging risks. Banks in particular are likely to proceed cautiously with regard to consumer-facing use cases until clearer, consensus-driven answers to these governance and oversight challenges emerge, given the potentially serious financial and regulatory consequences of poorly performing AAI applications.

For example, concerns about explainability, bias, data protection, and liability in AAI systems are likely to prompt many financial institutions to move slowly in deploying AAI for higher-risk activities such as lending, insurance, and wealth management, particularly in the absence of clear guidance from regulators as to how to comply with legal requirements and supervisory expectations. In parallel, many of the larger financial institutions have only recently begun to build internal expertise—particularly at the intersection of AI ethics, regulatory compliance, and operational risk management—needed to safely integrate AAI into consumer-facing products. Many smaller banks face substantial technical and resource limitations and are more risk-averse, viewing the reputational and regulatory risks of early adoption as outweighing the perceived benefits. Financial regulators are also relatively early in the learning curve both with regard to industry adoption and their own potential use of AAI systems.

BOX 5 MODEL RISK MANAGEMENT FRAMEWORKS AS APPLIED TO AI MODELS

While guidance tailored to AAI systems may take some time to emerge, the general model risk management (MRM) framework is a principles-based system that applies broadly to retail banks' use of models.[53] This gives banks a general scaffolding for identifying and mitigating potential risks from AAI during development, validation, and deployment across a wide range of contexts. At the same time, MRM processes and culture are not uniform across the entire financial sector and because AAI systems are frequently focused on activities that are not solely internal to individual companies, there are important implications that may require broader regulation.

MRM guidance has been developed by federal banking regulators to protect the safety and soundness of banks when using quantitative models (whether or not they involve artificial intelligence) across a broad range of functions. It is a principles-based system that calibrates oversight expectations based on the importance of the use case, the potential financial and legal risks to the institution, the nature of the model, and other risk-related considerations. Banks are generally expected both to conduct an appropriately rigorous risk assessment prior to adopting new models and to implement monitoring plans and controls after deployment.

While federal banking regulators' foundational MRM guidance dates from 2011 to 2017 and does not expressly address the use of artificial intelligence, it addresses a broad range of critical issues in developing and deploying complex analytics more generally. These include such topics as data governance, documentation of key development decisions, managing model complexity and explainability, validation testing by independent teams, and the importance of establishing monitoring systems and benchmarks.

The National Institute of Standards and Technology has drawn significantly on these MRM elements in developing many of its broader AI management materials, and other types of financial services providers sometimes look to the guidance as well.[54] However, non-depository institutions are not generally subject to examination or external monitoring in complying with these sources.[55]

Particularly when combined with the need to upgrade data and engage in comprehensive process mapping as described above, sensitivity to risk management and other regulatory concerns appear to be creating a substantial difference in how many depository institutions (and vendors that serve them) are approaching the use of AAI and LLMs more generally as compared to at least some non-bank tech companies. Bank ecosystem stakeholders emphasized the importance of breaking down functions and tasks into smaller subsets of activities, making careful decisions about where LLMs versus narrower models or tools (such as machine learning models or deterministic rules-based protocols) should be used to perform specific tasks, and deploying guardrails, human reviews, and monitoring as needed for different groups of activities. This incremental approach can take substantial time and effort compared to broader implementations that tend to rely more heavily on LLMs and agents to break down tasks and functions. However, particularly where companies are still gaining experience with the underlying technologies, proponents argue that it is the best way to build confidence that individual applications will perform as expected and that any problems can be identified and fixed quickly if they occur.

Collectively, these challenges underscore that while the technological foundations for AAI are advancing rapidly, the institutional, infrastructural, and regulatory ecosystems that could facilitate support for its safe and effective deployment in consumer financial services (as well as digital commerce) are not accelerating at the same pace. Addressing these issues can provide greater confidence to companies, investors, and consumers when making individual decisions about AAI systems, but they require coordinated action by financial institutions, regulators, technology providers, and consumer advocates. While technological solutions may be possible for some issues, modernizing infrastructure, clarifying governance standards, and building internal capabilities for AI oversight and ethical risk management would help to facilitate responsible adoption that benefits firms and customers alike.

5. CONSUMER PROTECTION IMPLICATIONS OF AAI

To the extent that AAI applications in financial services evolve toward broad-based consumer-facing applications such as the ones described in Section 3.2, these systems could increasingly mediate how consumers access, manage, and understand their financial lives—from budgeting and borrowing to saving and investing. As AI gains a more prominent role in making and executing financial decisions traditionally carried out by individuals, it is important to consider both consumer protection risks and features that will shape the utility and appeal of AAI applications to consumers, particularly those who could benefit the most from increases in access to high quality financial services. This section outlines four primary sets of concerns in the consumer financial services context.

5.1 The risk of misaligned incentives and instructions

Advanced AI systems offer the potential to ease consumers' cognitive and logistical burdens in making and executing financial decisions. Yet this convenience could also introduce a significant risk to consumer autonomy and oversight to the extent that AAI systems operate in ways that do not serve the interests of consumers, small businesses, or other users, whether because of intentional design choices made by providers, because the AAI systems evolve in ways that undermine users' financial health and interests, or because of communication problems between users and apps.

One fundamental question concerns the incentives of commercial providers of AAI in consumer financial services and the extent to which they make choices that are designed to serve their own interests rather than the interest of users. While prohibitions on unfair, deceptive, and abusive acts and practices, disclosure requirements, and basic economics provide some protections and incentives to design useful products, misaligned incentives increase the risk that AI tools may be structured in subtle ways that optimize for profitability at the user's expense. For example, absent guardrails, agentic shopping tools or AI-driven loan recommendation engines could potentially be designed to steer consumers toward higher-cost products to maximize commission revenue, referral fees, or other income to the tool developers. At a more subtle level, budgeting apps offered as adjunct services by individual financial institutions may simply not offer recommendations and assistance in obtaining outside financial services beyond the institution's own product offerings.

A second risk is that AAI performs poorly either because of communication challenges between the user and the system or because the system evolves into behaviors that do not serve users' interest or exceed their directions. Consumers may not always have a clear sense of their own short and long term financial goals or risk appetites. Poorly worded queries or instructions or failures to update goals or other information relating to consumers' objectives could also result in AAI systems

taking actions that are not aligned with users' wants and intentions. More broadly, stakeholders are debating to what extent it is necessary for consumers to understand what criteria AAI systems are following when making recommendations or taking actions, as well as how to deliver the appropriate level of transparency.

Technical issues can also be a risk, for instance if there is misalignment in objectives between different elements in agentic systems, such as between the lead agent and agents that are assigned specific tasks.[56] "Drift" in agentic systems also raises concerns that they may evolve in ways that their developers and users do not intend, for instance if an AAI investment tool gradually shifts toward higher-risk investments to maximize returns in ways that exceed the customer's instructions, or begins to perform new tasks in pursuit of broad goals. There is also some evidence that systems sometimes evolve in response to exposure to external data and interactions with each other, and that subsequent performance may vary depending on whether the systems expect to be subject to human review.[57]

The subtlety, autonomy, and opacity of AAI systems exacerbate these challenges. Their ongoing access to financial data, ability to engage in behavioral nudging, and capacity for autonomous action can potentially increase the scale, subtlety, and personalization of activities that do not serve consumers' interest. These dynamics have already arisen in the context of some AI-driven personal finance apps, which have been criticized for recommending cash advance offers and promoting borrowing products when a user's account balance dips below particular thresholds, even in situations where those products would worsen consumers' long-term finances.[58] In investment apps, concerns have been raised that AAI could potentially adjust recommendations toward higher-fee or riskier portfolios based on inferred consumer risk tolerance or profitability potential, without clear disclosure or informed consent.

The technical complexity of these systems also makes explainability and transparency substantially more challenging for consumers, financial institutions, and regulators.[59] In contrast to GenAI systems, which raise concerns about reliability and complexity because they are often trained on large amounts of data scraped off the internet, individual AI agents can be trained on relatively modest amounts of data depending on their function. However, AAI systems can involve dozens or hundreds of agents interacting and learning over substantial periods of time in a dynamic environment, as well as integrating LLMs or other types of models that may involve larger amounts of data. These ensembles and interactions make explainability more challenging than tracing the production of a single prediction from a static model or the generation of content based on a single series of queries.[60] There is also greater concern that AAI systems could engage in strategically deceptive behaviors in pursuit of particular goals compared to genAI systems.[61] Internal governance and monitoring and independent testing programs are both potentially critical to help provide visibility to users and other stakeholders.

One additional complication is the impact that "privacy fatigue" has already had on how consumers approach consumer-facing financial technologies. Much like other sectors, the continuing stream of privacy policies, software updates, cookie notices, and consent prompts from modern financial apps and services can tend to desensitize consumers, leading many to reflexively accept default settings without fully grasping what they are authorizing.[62] The risk of desensitization potentially becomes even higher stakes in contexts where consumers authorize an autonomous AAI to manage major aspects of their financial lives. Behavioral research indicates that many consumers tend to overestimate the degree of oversight they retain in digital environments, a dynamic that becomes especially problematic as AI agents grow more sophisticated, context-aware, and autonomous.[63] This pattern could potentially result in a gradual erosion of consumer awareness and agency, where individuals may unwittingly relinquish control over high stakes financial decisions. In light of

these concerns, protecting consumer autonomy and ensuring informed, ongoing consent emerge as urgent challenges for financial regulators and technology designers alike.

5.2 Responsibility for unauthorized transactions and other errors in handling consumer funds

The prospect of financial agents autonomously initiating and executing payments, funds transfers, investments, and other financial transactions on behalf of consumers also raises important questions about liability and other financial consequences where the transactions go wrong. The answers may be complex depending on what means of payment is being used, what business is supplying the financial agent to the consumer, what aspect of the transaction went wrong and why, and how existing federal laws governing liability and dispute resolution for payments transactions may apply. These questions have substantial implications for consumers as to whether and how they will recover from financial losses, as well as for financial services providers that are involved in processing the payments and for merchants or other companies that are providing the goods or services.

For example, in the context of traditional credit card transactions, the Truth in Lending Act (TILA) generally limits consumers' liability from unauthorized charges, charges for goods or services not accepted or delivered as agreed, and computational and other errors.[64] The Electronic Fund Transfer Act (EFTA) also limits consumers' liability from unauthorized electronic transfers involving deposit and other transaction accounts, and requires the financial institution providing the account to investigate and correct incorrect fund transfers to or from a consumer's account.[65] Credit card and debit card networks have adopted additional rules on top of these legal regimes to facilitate dispute resolution between the various businesses that are involved in the payments transactions, either as providers of the ultimate goods and services or as part of the payment processing ecosystem. However, the two laws and the payment network protections generally exclude transactions that are made by a person that the consumer has authorized to use their credit or debit cards, meaning that the consumer must absorb any initial losses and seek relief from the human agent directly.[66]

BOX 6 AAI TRANSACTIONS INVOLVING TRADITIONAL PAYMENT NETWORKS

Where a consumer disputes a particular card transaction as unauthorized, both TILA/EFTA and the payment network rules generally require the institution that is providing the credit card or transaction account to reimburse the consumer initially and then work with merchants and other parties involved in the transaction on the back end to sort out which entities bear which losses. The payment networks' rules assign the risk of disputes differently depending on whether the transaction involved a physical card at a physical point of sale. For "card present transactions," the risk is on the bank that issued the card to the consumer so long as the merchant obtained a valid physical authorization. For "card not present transactions," liability generally falls on the banks that help merchants process the card transactions, and those institutions in turn impose the liability on the merchants by contract.

Given that physical cards are not presented to merchants in the course of AAI transactions, some stakeholders have pointed out that merchants (and their banks) have a strong incentive under existing frameworks and rules to try to block AAI payments to limit their potential liability for disputes. However, depending on what company is providing the agent to the consumer in the first instance, it may not always be easy for the merchant and the processing financial institution to identify the transaction. For example, it may not be clear whether a transaction conducted via a large general e-commerce platform or an agentive system that has been selected independently by a consumer involves an agentive shopping tool or is being conducted directly by the consumer.

Some stakeholders have suggested amending the rules that implement TILA and EFTA to avoid assigning liability to merchants, while requiring providers of agentic shopping tools to provide disclosures and other tools to ensure that consumers understand that their rights to dispute transactions would be reduced relative to traditional transactions.[67]

BOX 7 AAI TRANSACTIONS INVOLVING STABLECOINS AND CRYPTOCURRENCIES

Use of stablecoins or cryptocurrencies in AAI transactions raises additional questions about the resolution of unauthorized or erroneous transactions. Interest in using these currencies for AAI transactions is substantial because they can be used to make payments in real time across the globe, without having to wait for card network processing or being subject to related fees.[68] However, EFTA does not define what types of "funds" are subject to its protections in either statute or regulation, and only a few court decisions have addressed the topic.[69] A number of crypto proponents have argued that US-based transactions in digital currencies should not be subject to EFTA.[70]

Although stablecoins and cryptocurrencies were not widely used for general retail transactions in the US prior to the adoption of federal legislation in summer 2025, some initiatives were beginning to emerge. For example, Coinbase Payments launched to facilitate retailers' use of stablecoins at checkout, with Shopify adopting it in partnership with Coinbase and Stripe.[71] Major retailers such as Walmart and Amazon were also reportedly considering issuing their own stablecoins to establish alternatives to traditional payments systems and facilitate the offering of benefits and loyalty programs.[72] A number of banks and other financial services providers are pursuing partnerships with stablecoin providers or tokenized deposit solutions that also involve blockchain technologies, without abandoning traditional accounts and payment rails (and related dispute resolution and liability systems) altogether.[73]

The 2025 stablecoins legislation limits issuance to banks or other approved entities and imposes certain other restrictions on issuers' operations, but did not address EFTA applicability. Broader legislation on cryptocurrencies market structure was pending at time of publication.[74]

Agentic AI applications raise critical questions about who should take responsibility for transactions that go wrong and whether and how to apply existing regulatory and market structures or to create new ones. Among stakeholders who are viewing these issues through the lens of traditional payment systems, many have suggested that in situations where a consumer chooses to deploy a third party AAI system that simply fails to execute a transaction properly (as compared to situations in which the system is hacked by an outside bad actor, which is discussed in more detail in Section 5.3), it should be treated as the equivalent of a human agent error that exonerates both merchants and payment processors from liability and leaves the account owner and the third party agent provider to sort out which party absorbs financial losses and other harms. More broadly, the application of general product liability laws to AI and other types of software are still evolving. For instance, some software developers use end user license agreements to disclaim liability for their products once downloaded to a user's device, and liability as between developers of foundation models and companies that build their own applications based on fine-tuned versions of those models is also unclear particularly in fields with substantial compliance obligations such as financial services.[75]

These questions are critical to determining the extent to which consumers must absorb initial or long-term losses and engage with various parties to try to sort out unauthorized or incorrect transactions. They also have important implications for the providers of the underlying products and services, payment processors, and other actors such as the agentic system developers. Indeed, even separate from the consumer protection concerns, some stakeholders have suggested that uncertainty over the business-to-business dynamics could cause merchants to try to block AAI payments involving debit or credit cards until there is greater clarity about dispute processes and liability rules.[76] (See Box 6). Some proponents of AAI systems are also motivated to conduct transactions in stablecoin or other cryptocurrencies, which could also have impacts on dispute rights and liabilities (see Box 7).

5.3 Data governance and security challenges

Depending on the use case, AAI systems may handle large amounts of highly sensitive personal financial data—including users' spending habits, location history, social connections, employment status, and health proxies. While they could potentially help consumers manage their data as well as their finances, their AAI systems' autonomy and complexity also raise concerns that they may become major targets for bad actors and fail to safeguard consumers' information.

The first issue is whether AAI systems will use sufficient care when storing and sharing data. Although initial consumer-facing applications may require consumers to approve each data transfer, payment, or financial account opening individually, some stakeholders worry that the systems will increasingly be able to obtain credit reports, share transaction and investment details with outside companies, or authorize third-party APIs to access bank accounts without the consumer's active, informed consent. AAI systems' memory capabilities also increase concerns that they could use sensitive information in unexpected downstream contexts without obtaining express consumer consent.

Sharing such information inappropriately with outside parties may expose consumers to cross-selling, pricing manipulation, or other unwanted and damaging consequences. Increases in AAI systems range of activities could raise questions about the application of data and consumer protection laws as well. For example, the extent to which agents can receive disclosures, consent to data sharing or other actions, or exercise dispute rights on behalf of consumers could have important implications for application of data protection laws—whether sector-specific (like the Fair Credit Reporting Act, Gramm-Leach-Bliley Act, or state laws) or general (like the European Union's General Data Protection Regulation).[77]

A second concern is that AAI systems are likely to become frequent targets of attacks from malicious actors that may themselves deploy agentic systems. The APIs that function as the backbone of AI-based automation are already seeing increases in attacks from bad bots, which are estimated to constitute 38% of internet traffic to companies in the financial services sector.[78] Certain technical issues may also increase risks. For example, AAI systems sometimes behave unpredictably when interacting with each other, raising concerns that different elements or systems might exploit each other or collude in unintended ways.[79] Additionally, because some AAI systems rely on less data than GenAI models, they may be more susceptible to adversarial inputs, data poisoning, or model inversion attacks.[80]

These issues underscore the importance of building strong guardrails, thorough testing, ongoing monitoring, and prompt intervention, since data and cybersecurity concerns affect the appeal of AAI systems for both financial services providers and consumers.

5.4 Questions about trust and tailoring to LMI populations

Building on the concerns raised above, there are also questions about the extent to which consumers who are financially vulnerable or have struggled to access financial services in the past will find personal AAI tools to be tailored to their needs and be willing to adopt them. These populations could potentially experience significant benefits from reliable tools that are designed to manage their specific financial challenges and goals, but their past life experiences may also heighten their concerns about the risk of financial loss, control over their personal financial data, and general suitability.

Research conflicts as to the extent to which consumers are willing to disclose sensitive information to automated systems and rely on the resulting advice. Some trends and studies suggest

that people are often more willing to disclose sensitive information to automated systems than to humans, perhaps in part because of fear of judgment or difficulty in accessing human assistance.[81] At the same time, other research has documented persistent consumer skepticism toward algorithmic decision tools, particularly among people who are concerned about sharing personal data or who consider themselves to be well educated on related topics.[82]

Lab-based tests by the United Kingdom's Financial Conduct Authority on tools for minimizing overall interest and fees across multiple debt payments documented ways that issues of trust, data privacy, and the perceived value of explainability can shape how consumers engage with automated advisors. Though it did not differentiate based on income, the study found that consumers with lower financial literacy experienced the greatest financial benefits from the tool. Willingness to pay for the tool was highest among that same group, but lower among high literacy consumers and those who were more distrustful of algorithms, who were also more likely to experience less savings because they overrode the recommendations. On average, participants were not willing to pay more for a version that explained its actions compared to a version that simply delivered the answer, and the study did not find that having used the version of the tool that provided explanations led to improved outcomes when participants were later asked to allocate payment amounts without the tool.[83]

Surveys by the non-profit CommonWealth find that LMI consumers tend to have less exposure to chatbots and lower comfort levels with robo advising than wealthier populations, yet interest levels tend to be higher among LMI populations who are younger, Black or Hispanic, or have recently experienced a financial hardship. Across respondents, concerns about data security and privacy ranked slightly ahead of concerns about reliable information, though interest in finding trusted financial education sources was high across LMI respondents.[84] The survey results suggest that successful AI advisors will prioritize earning trust through communication and education about their capabilities and information sources, drive engagement through content that is personalized to the consumer's individual finances and goals, provide consumers with meaningful options to balance automation and control, and offer multilanguage accessibility. The ability to access human assistance where issues arise with the automated system can also substantially affect consumers' willingness to experiment with new technologies.[85]

Some stakeholders also expressed concern that without deliberate design choices, development of personal financial agents may tend to prioritize the needs of affluent consumers and leave LMI populations with generic, ill-fitting, or inaccessible tools. For example, assistance in navigating government benefits applications processes and data sources is specialized and mistakes could have different consequences for consumers. As the CommonWealth study intimates, those most in need of access to high quality financial advice may be open to using AI services, but trust in AI in general remains low among LMI populations—in part because of data privacy and security concerns or because of their expected lack of personalization to LMI situations. This gap poses both a potential market failure and a financial security concern, warranting further research and policy attention to ensure AI-driven financial services do not reinforce existing disparities in financial health and inclusion.

A related risk concerns the potential amplification of bias by AAI to the extent they are trained on financial behavior data and predictive models that reflect existing racial, gender, and socioeconomic disparities or that they are structured to assume particular personas that may contain hidden biases based on personality characteristics.[86] Without appropriate safeguards and monitoring, AAI systems could potentially steer consumers from marginalized communities toward higher-cost products, riskier financial behaviors, or limited service options based on inaccurate or biased assumptions about their financial situations, further entrenching inequities in financial outcomes and access.

6. POTENTIAL PROVIDER AND SYSTEMIC RISKS RELATED TO AAI

Poor design and performance of AAI applications could also present risk of financial losses or other harms to individual financial services providers and, depending on scale, the broader financial system. While individual providers have greater control over the deployment and risk management of their own AAI tools, it is also important to consider the potential impacts of AAI systems that are developed by other companies and deployed by a broad range of parties, including the providers' own customers.

Many of these risks are similar in nature or directly related to concerns raised in Section 5. For example, goal misalignment and drift within AAI systems is a potential concern for internally facing AAI systems as well as for consumer facing use cases, though companies are potentially better positioned to maintain clear and consistent communications, monitoring functions, and interventions than individual users. Concerns about data governance and information security also affect financial services providers, and could potentially stress their systems in dealing with multiple external agents as well as bad bot traffic. Some banks are also growing increasingly concerned about the lack of clarity over liability for errors in money movement by agentic systems, given that they often have to manage disputes and absorb initial costs and are sometimes looked to by customers and regulators to make consumers whole even where regulatory and network requirements do not mandate that outcome.

The speed and potential scale of agentic programs in taking autonomous action also raises potential concerns that they could amplify volatility, feedback loops, synchronized behaviors across multiple AAI systems, and cascading failures if they react to market signals in ways that exacerbate negative cycles or events. For example, as financial institutions increasingly rely on a small number of third-party foundation models, some stakeholders have voiced concern that critical flaws, data poisoning events, or hidden biases could trigger simultaneous, destabilizing activities across dozens of institutions. AAI systems' capacity to operate instantaneously at substantial scale also creates the potential for AI-driven flash crashes that are far more complex and rapid than those seen in current high-frequency trading systems. The widespread use of AAI for activities such as real time or dynamic insurance repricing, mass mortgage refinancing, or AI-managed debt collection strategies could also potentially raise procyclicality challenges and other concerns across other markets, such as exacerbating consumer hardships and liquidity shocks during rapidly changing economic conditions and natural disasters.

Concerns have also been raised about whether AAI systems could increase volatility with regard to bank deposits, and thus undermine the stability to the broader financial system. The concern is that large numbers of financial agents could begin transferring deposits between different banks

seeking higher returns or responding to concerns about individual institutions' stability. At its most extreme, this could potentially both complicate situations such as the Silicon Valley Bank case where concerns about individual banks' soundness emerge and undermine the stability of banks' deposit base more generally. However, the extent to which such transfers could reach scale would depend on several factors, including the extent to which Know Your Customer laws and other constraints limit AAI systems' ability to open new accounts and transfer funds without the active participation of accountholders.

The risks of different agentic systems influencing each other that were mentioned in Section 5 also factor into concerns about collective action producing negative systemic effects. Research from institutions like INET Oxford on agent-based modelling has suggested the most dangerous outcomes are not necessarily programmed but rather are emergent properties of complex systems.[87] This increases concerns that autonomous systems could develop "herding" behaviors or even learn to coordinate in ways that resemble tacit collusion, leading to negative consequences for the market as a whole as well as individual firms and consumers. For example, the adoption of AI pricing algorithms in the German petroleum market led to increased prices indicative of collusion in markets where humans were no longer involved in pricing decisions.[88]

More broadly, though it is beyond the scope of this report, the adoption of AAI systems is likely to add further fuel to discussions about the impacts of AI on the labor market and on competitive dynamics between different firms depending on their technology and data resources.[89]

7. CRITICAL QUESTIONS GOING FORWARD

Realizing the potential benefits of AAI while safeguarding consumers and financial system stability will depend on addressing a range of issues concerning safe adoption. While the potential application of existing laws and guidance or the development of new regulatory frameworks is an important question, decisions about technology, business models, market practices, and broader governance mechanisms will also play a critical role in shaping deployment.

Some industry participants have begun suggesting frameworks for risk management and self-regulation that may be instructive. For example, IBM Consulting in Australia has produced a catalogue of fifteen types of risk and associated controls that financial services providers can implement at different layers within their internal AAI systems.[90] Sardine.AI, a financial technology company focused on fraud prevention and transaction risk management, has outlined a governance framework for that context that includes carefully structured training and data curation, presentation of AAI outputs to human decisionmakers, comprehensive ongoing auditing and oversight mechanisms, and structured improvement loops. The model emphasizes early and continuous risk monitoring throughout the AAI lifecycle—an approach well-suited to the ongoing, adaptive governance challenges posed by agentic AI systems across other areas of consumer finance.[91]

Academic papers and industry sources are also cataloguing metrics and benchmarks for measuring AAI systems' performance across a range of dimensions, and the National Institute of Standards and Technology is beginning the process of developing tailored security guidelines for AAI systems as well as other AI applications.[92]

This section highlights key issues and questions for consideration that will ultimately help to shape both the benefits and risks that are realized from the adoption of AAI systems in financial services. It starts with questions for developers of financial AAI systems, outlines broader ecosystem questions regarding consumer agency and protection as well as financial stability, and concludes with questions about regulatory oversight.

7.1 Core development and risk management practices by companies offering financial AAI systems

Firms developing financial AAI applications face a series of questions about how to ensure that the resulting systems are generally fit for purpose, in addition to issues concerning consumer and systemic risks and regulatory expectations as discussed in additional detail in later sections.

- What kinds of data curation and governance practices are most effective in working with dynamically updating and autonomous systems that may contain multiple components (agents, LLMs, etc.)? Where can cross-industry data infrastructures and coordination mechanisms be modernized to facilitate the safe, accurate, and efficient operation of AAI systems?
- What technology tools can be applied to AAI systems to ensure that they act in accordance with the interests of consumers, small businesses, or other users and with legal requirements? How can and should users' interests be defined in this context?
- What quality control, auditability and explainability structures, and human-in-the-loop oversight mechanisms are most effective in ensuring that AAI systems remain fit for purpose over time?
- How can AAI systems be designed, monitored, and adjusted to avoid amplifying bias in credit, insurance, or wealth-building decisions?
- How can AAI systems be designed to protect the privacy and security of consumer financial data in agent-driven ecosystems?
- How can developers improve their testing protocols to assess how AAI systems may interact with each other after deployment? Are there additional safeguards that AAI system developers can implement to guard against herding behavior and other systemic risks?
- Where do traditional risk management and compliance frameworks, infrastructures, and processes need to be adjusted or expanded for application to systems that are dynamically updating and substantially more autonomous than prior generations of AI?
- Where AAI systems are built to incorporate foundation models, what responsibilities should fall on the foundation model developers as compared to the downstream AAI system developers, respectively, to ensure that systems are fit for use, comply with law, and can be trusted by end users?

7.2 Additional consumer agency and protection questions

- How can consumers communicate effectively with and have confidence that AAI systems' objectives and scope of authority will follow their directions and meet their expectations? What transparency mechanisms are most effective for helping consumers understand AAI systems' ongoing activities and outcomes?
- How can consumers maintain meaningful oversight and control over the activities of autonomous AI systems acting on their behalf? How should consent processes be designed? Should there be limits on what can be automated?
- Which parties should be responsible for addressing situations in which agentic AI systems take actions that are inconsistent with consumers' instructions or otherwise do not serve interests? What are effective models for consumer redress when harmed by opaque or autonomous AI-driven financial decisions?
- Can and should agentic AI systems be required to prioritize consumer financial well-being (akin to fiduciary duties), and how would such obligations be defined and enforced?
- Do providers of AAI systems have an obligation to warn consumers that dispute rights and liability limitations may be different than what they are used to in conducting direct bank account transfers or transactions with their debit or credit cards?

- Particularly in light of existing financial regulatory requirements that give consumers the right to make inquiries, dispute errors, or obtain information, do those requirements need to be adjusted to account for situations where AAI systems perform such activities on consumers' behalf?
- Are new regulatory frameworks or adjustments to existing consumer protections regarding liability limitations, discrimination, data sharing and protections, or unfair, abusive, or deceptive acts and practices needed to address the use of AAI systems?

7.3 Additional financial stability questions

- Are there specific types of correlated market behaviors by AAI systems that should be incorporated into traditional systemic risk and stress testing models to facilitate better monitoring and mitigation?
- What early warning indicators and systems could detect destabilizing behaviors among agentic financial systems before crises emerge?
- What failsafe or human-intervention protocols should be required for agentic AI systems operating at financial system scale?
- Should providers of foundation models ever bear liability where financial AAI systems malfunction due to issues with integrated LLM components (e.g. inadequate design, negligence in training data, failure to install safety measures)?
- Are new regulatory frameworks or adjustments to existing regimes focusing on managing systemic risk needed to address the use of AAI systems?

7.4 Additional regulatory oversight questions

- Do regulators need specialized units, interagency tasks forces, international agreements, or other resource and infrastructure investments and mechanisms to facilitate effective AAI systems monitoring?
- What technological tools and adjustments to existing regulatory frameworks may be needed to meet explainability needs and monitor outcomes without unduly constraining innovation in agentic financial system? How can federal and state interests in regulation be optimally balanced?
- How should supervisory oversight infrastructures and processes be adapted for continuous, adaptive AI systems versus static, rules-based financial models?
- What advantages and disadvantages do industry self-governance mechanisms have in the AAI systems context, and how can they potentially be strengthened by design choices regarding direct supervision and regulation?
- What role should third-party audits and other independent research and testing play in verifying that agentic AI systems meet fairness, safety, and financial stability standards?
- Where and how can regulatory sandboxes and controlled pilot programs be structured to accelerate beneficial use cases, the refinement of industry and regulatory standards, and the identification of effective risk mitigation strategies?

8. CONCLUSION

Although financial services stakeholders are still working to manage the adoption and governance of machine learning and GenAI, agentic AI deserves substantial and immediate attention. AAI systems' ability to deploy agents, LLMs, and other technologies cooperatively to incorporate new information dynamically, operate autonomously, and execute multi-step workflows has tremendous potential across a broad range of financial services applications, including the creation of powerful new financial management tools for consumers and small businesses.

However, deploying these systems to support and perform high-risk/high-reward functions will require solving a series of technology, business, governance, and regulatory issues to address potential concerns about reliability, transparency and control, customer protection, error correction, and financial stability. Even if individual companies are not ready to implement AAI systems, they may be affected by the decisions of consumers or other companies to adopt agentic applications.

Reducing uncertainty around critical issues can both help to mitigate risks and facilitate beneficial adoption. A critical early step toward these goals is to increase awareness and engagement across the financial services ecosystem by bringing together a broad range of technologists, business leaders, researchers, policymakers, advocates, and other stakeholders. Centering the needs of consumers and small businesses and building effective mitigation systems prior to deployment are also particularly critical when building direct-to-customer applications that could affect users' long-term financial health and wellbeing.

FinRegLab is working to facilitate discussions of these issues at its 2025 AI Symposium and through other potential convenings, research, and analyses. We welcome feedback on this report at contact@finreglab.org.

APPENDIX

Recent Federal and State AI Activity

AAI may be implicated by many broader federal and state policymaking initiatives even though they are not focused specifically on the qualities that distinguish it from other types of artificial intelligence.

The Trump Administration announced a policy "to sustain and enhance America's global AI dominance" through the removal of regulatory barriers and various international activities in a January 2025 executive order.[93] Federal agencies subsequently strengthened export controls on AI chips, though they rescinded an "AI diffusion" rule issued by the outgoing Biden Administration before it could take effect.[94] After receiving more than 10,000 stakeholder comments, [95] the Administration released its much anticipated National AI Action Plan in July 2025 to flesh out the executive order.[96] The plan has three primary pillars focusing on accelerating AI innovation within the US, expanding infrastructure, and strengthening export controls and global activities.

In the domestic context, the Plan emphasizes reducing undue regulatory burdens on AI development, enabling a "try first" culture, and facilitating infrastructure development, worker training, and access to data. This includes investing in research to advance methodologies to provide interpretability, controls, robustness, and evaluation metrics, including a number of initiatives managed by the Center for AI Standards and Innovation within the National Institute of Standards and Technology (NIST).[97] Although there is some discussion of protecting privacy and confidentiality, the Plan does not discuss financial services or consumer protections in detail.[98]

As federal policymakers shift their priorities, state legislative activity has continued to accelerate. In 2025, approximately 1,100 AI-related bills have been introduced across all 50 states and the District of Columbia, Puerto Rico, and U.S. Virgin Islands, with 100 measures adopted across 38 states as of mid-July.[99] This has already outpaced the 700 bills introduced in 2024 and creates the potential for a complex and diverging patchwork of requirements.[100] In addition to looking to ongoing legislative activities in California,[101] a number of states are looking to comprehensive legislation adopted by Colorado and Texas as well as to a somewhat narrower bill in Utah as potential templates even as those states continue to refine their requirements.[102]

For example, the Colorado Artificial Intelligence Act (CAIA, SB 24-205) targets high-risk AI systems that substantially influence "consequential decisions," such as those affecting housing, healthcare, employment, education, finance, or legal services, and imposes a duty to prevent algorithmic discrimination.[103] Developers and deployers of high-risk systems doing business in Colorado must implement risk-management programs, conduct impact assessments, and adopt governance structures aligned with frameworks like those of NIST.[104] The law also requires pre-use notice, an explanation of adverse decisions, and rights to correction and human review or appeal when feasible.[105] Implementation is scheduled for February 1, 2026, though Governor Polis called a special legislative session for August

2025 to revisit CAIA's implementation date and address both budget impacts and industry concerns about compliance burdens.[106]

In the face of so much state activity, several large companies have actively lobbied the federal government to preempt state-level AI regulations, arguing that a fragmented, state-by-state approach would stifle innovation and complicate compliance for firms operating nationwide. The One Big Beautiful Act originally contained a 10-year moratorium on state regulation, but it triggered substantial opposition including among some Republicans and at least one tech CEO before being removed by the Senate.[107] The National AI Action Plan states that federal AI funds should not be directed toward states with "burdensome AI regulations that waste these funds," but should also not interfere with prudent laws that are not unduly restrictive to innovation.

Other federal legislative activity has picked up but remains fragmented, with draft bills focusing on AI transparency, liability, systemic risk monitoring, and sector-specific safeguards.[108] Internationally, the European Union continues to advance its AI Act, and other global partners are developing AI regulatory frameworks with implications for cross-border financial services and AI-powered consumer products.[109]

This evolving policy landscape illustrates the challenges of federal authority, state experimentation, geopolitical security risks, and commercial imperatives in shaping the future of AI policy. It also underscores the value of a cohesive and adaptive government strategy that can address the implications of AI in sectors such as financial services where AI adoption raises core priorities including financial stability and consumer protection.

Endnotes

1 Ghose et al., "Agentic AI: Finance and the 'Do It for Me' Economy."; Jackson, "Nvidia's CEO Says We're in the Age of 'Agentic' AI—Here's What That Word Means."; Adebayo, "Here Comes the Big, Strong Agentic AI Wave."

2 Samuel, "Some Studies in Machine Learning Using the Game of Checkers." For general background on previous types of AI, *see* FinRegLab, "Explainability & Fairness in Machine Learning for Credit Underwriting: Policy Analysis," § 2.

3 Model Context Protocol is a standardization layer meant for AI systems such as LLMs to communicate with external tools and data sources, including those accessed through application programming interfaces (APIs), which are dedicated systems to allow software to talk to each other. Introduced by Anthropic in 2024, MCP establishes an open standard for AI to interact with a broad variety of tools and databases across many different developers and products. Anthropic, "Introducing the Model Context Protocol."; Gutowska, "What Is Model Context Protocol (MCP)?"

4 These models are described by some sources as "deterministic," although other sources reserve that term for rules-based or mathematical calculations that do not involve any use of predictive statistics and still others describe some uses of LLMs in agentic AI as being deterministic if they are structured with sufficient controls. See generally Neidhart-Lau, "Analyzing AI Workflows: Deterministic and Autonomous Agents in Financial Services."; Juristech, "Why Deterministic Agentic AI Is the Breakthrough Modern Banking Needs."

5 Specifically, LLMs compute probability distribution patterns to predict the next word (or portion of a word), but randomly sample from those distributions in generating query responses. This helps to make their responses more conversational but means that a given query will not generate the same response every time. The amount of variation can be dialed up or down for particular use cases but achieving complete consistency is challenging.

6 Belcic and Stryker, "What Is Agentic RAG?" During training, developers may also use reinforcement learning to provide feedback to models to encourage more useful answers. Murel and Kavlakoglu, "What Is Reinforcement Learning?"

7 Sapkota et al., "AI Agents vs. Agentic AI: A Conceptual Taxonomy, Applications and Challenges."; Chawla and Pachaar, "Implementing React Agentic Pattern from Scratch."

8 For general discussions of technical challenges and risk management, *see, e.g.*, Chorev et al., "Agentic AI in Financial Services: Opportunities, Risks, and Responsible Implementation."; Sapkota et al., "AI Agents vs. Agentic AI: A Conceptual Taxonomy, Applications and Challenges."

9 Anthropic, "Claude for Financial Services." Anthropic has collaborated with major global service integrators (such as Deloitte and KPMG) and data management/analysis services (such as Palantir and Snowflake) to help customers with onboarding and operations.

10 OpenAI, "Introducing ChatGPT Agent: Bridging Research and Action." The new offering combines OpenAI's ChatGPT large language model, its previous "Operators" agentic AI product and "DeepResearch" tools, and new functionalities.

11 Mistral AI, "Build AI Agents with the Mistral Agents API."

12 Rosenbaum, "Meta Is Targeting 'Hundreds of Millions' of Businesses in Agentic AI Deployment."

13 Sharma, "New Capabilities in Azure AI Foundry to Build Advanced Agentic Applications."; Poccia, "Introducing Amazon Bedrock AgentCore: Securely Deploy and Operate AI Agents at Any Scale (Preview)."; Ahmad, "New Agents and AI Foundations for Data Teams."

14 AI frameworks and management platforms often provide libraries, pre-built general components, and other tools to help users develop and manage agentic systems. Platform-agnostic frameworks include LangGraph, LangChain, CrewAI, and DSPy, while SalesForce's AgentForce is designed to integrate with its customer relationship management platform.

15 Vals.ai, "CorpFin (v2) Benchmark."; Vals.ai, "Financial Agent Benchmark."

16 For other discussions of financial services use cases, see Ghose et al., "Agentic AI: Finance and the 'Do It for Me' Economy."; Levitt, "AI On: How Financial Services Companies Use Agentic AI to Enhance Productivity, Efficiency, and Security."

17 PYMNTS, "Agentic AI Is Ready. Are Banks Ready to Use It?"; Taylor et al., "The Agentic Oversight Framework: Procedures, Accountability and Best Practices for Agentic AI Use In Regulated Financial Services."

18 *See, e.g.*, Columbus, "DanaBot Takedown Shows How Agentic AI Cut Months of SOC Analysis to Weeks." Cyber detection tools such as security information and event management (SIEM) response tools already employ a variety of automated software systems for identifying errors, breaches, abnormalities, and potential flaws in network systems, but proponents point to agentic AI systems' ability to facilitate coordination and communication of signals in near real time. Vajpayee and Hossain, "Cyber Defense through Agentic AI Enabled Automation: An Approach to Reduce Cyber Risk."

19 For example, traditional, rules-based screening processes often generate substantial numbers of "false positives" that are determined to be legitimate applications or transactions after subsequent human reviews. While supervised machine learning models and other advanced analytic systems can substantially improve predictive accuracy, they still generally have to be retrained periodically to incorporate new data reflecting changes in patterns. AAI systems can be structured to facilitate faster pattern detection, model adjustment, and automation of various workflow processes. For example, while credit card companies often use automated processes for certain functions such as contacting consumers for confirmation when particular credit card transactions are flagged as suspicious, AAI systems can perform a wide variety of other activities to help streamline broader workflows and functions. Taylor et al., "The Agentic Oversight Framework: Procedures, Accountability and Best Practices for Agentic AI Use In Regulated Financial Services."

20 For example, Alphabet, Google's parent company, recently announced it used an AI Agent to automatically identify and isolate/respond to an imminent security threat, likely the first event of its kind. PYMNTS, "Agentic AI Turns Enterprise Cybersecurity into Machine vs. Machine Battle."

21 For background *see* FinRegLab, "Innovations for Identity Proofing and Transaction Monitoring: Advancing Financial Inclusion Through Data & Technology."

22 Watkins, "Intelligent AI Underwriting's Impact on Modern Mortgage Lending."; Ali, "Efficient Underwriting Using Agentic AI."

23 Dien, "Agentic AI for Insurance Agents & Bank Advisors—The Complete Guide."

24 Parametric policies cover the probability of an event taking place, such as a hurricane striking a coastline or flooding in a major city. When a specific parameter has been triggered (say, the National Weather Service defines the hurricane as a category four across certain geographic areas) the policy automatically pays out a previously agreed amount. Because parametric policies pay based on the risk and not the actual damage, they typically have a spread between costs incurred and expected. Integrated AI systems may be able to close the gap between costs incurred and predicted, potentially offering better cost and risk calculation across time and smoother and faster customer engagement and payouts. *See* Blader, "AI Is Supercharging Parametric Insurance."; Ruffing, "How AI and Parametric Models Are Revolutionizing Risk Protection for Crop Insurance."; Balasubramanian et al., "Insurance 2030—The Impact of AI on the Future of Insurance."

25 For example, in November 2024, Stripe released its AI Agents Software Development Kit for businesses connecting LLMs with Stripe APIs. The tools allow businesses to manage and improve their payments and accounts receivables workflows and user experience. The AI Agents act as communicators (translating technical jargon into layman's terms), code debuggers, and goal-achieving "coworkers" that can pursue a task set to them by a human agent: such as "go through these invoices, sort and check them against internal approvals, pay and get paid for each one and highlight any irregularities. Kaliski, "Adding Payments to Your LLM Agentic Workflows."

26 *See generally* Josef, "Agentic Commerce: Where Consumer Intent Meets Merchant Opportunity."; Muhn, "4 Companies Bringing Agentic AI to Checkout."; Torculas, "Agentic Finance: The Financial Infrastructure Powering AI Agent Commerce."

27 Trelka, "Meet Opera's AI Browser Operator."

28 Polonioli, "Amazon's 'Buy for Me' Marks the Rise of Agentic Commerce—And It's Just the Beginning."

29 Rincon, "Shop with AI Mode, Use AI to Buy and Try Clothes on Yourself Virtually."

30 PayPal, "Perplexity Selects PayPal to Power Agentic Commerce."

31 Gosby, "Walmart: The Future of Shopping Is Agentic. Meet Sparky." *See also* Walmart, "Retail Rewired Report."

32 Silberstein, "OpenAI, Shopify Reportedly Working on ChatGPT Checkout Integration."; Montti, "OpenAI Quietly Adds Shopify as a Shopping Search Partner."

33 Visa's agentic AI program offers a series of API-based "modules" for AI developers to integrate and connect with a customer's payment credential. The customer can then direct the developer's AI tool to search and execute a transaction within preset limits and to pay autonomously using that customer's Visa card in a tokenized form. Mastercard has also launched a tokenized payment credential that will become operational with AI conversation platforms, such as OpenAI's ChatGPT. It also plans to integrate its agentic payment tokens with acquirer and checkout players like Braintree and Checkout.com and with IBM's WatsonX to pursue B2B agentic payments use cases. Mastercard, "Mastercard Unveils Agent Pay, Pioneering Agentic Payments Technology to Power Commerce in the Age of AI." (announcing an agentic pay program that includes a program focusing on agentic token and collaborations with companies such as Microsoft, IBM, Braintree (a subsidiary of PayPal), and Checkout.com); Visa, "Find and Buy with AI: Visa Unveils New Era of Commerce." (announcing "intelligent commerce" programs including tokenized digital credentials and partnerships with Anthropic, IBM, Microsoft, Mistral AI, OpenAI, Perplexity, Samsung, and Stripe).

34 Wolff, "Amazon Implements Guardrails as AI Agents Threaten Traffic, Ad Revenues."

35 Smith, "Shopify Quietly Sets Boundaries for AI Agents on Merchant Sites."

36 *See, e.g.*, Ecommerce North America, "Test Your Ecommerce Site for AI Agent Accessibility."; AdMetrics, "Agentic Commerce: How AI Agents Are Reshaping Ecommerce Forever."

37 Chak et al., "Can Robo-Advice Improve Borrower Repayment Decisions?"; Wells Fargo, "Comparing the Snowball and the Avalanche Methods of Paying Down Debt."

38 Pannala, "Rebooting Robo-Advisory: Can the Industry Meet Investor Demands for Personalization? (Part 1)."; Pannala, "Rebooting Robo-Advisory: Leveraging Generative AI to Power Next-Generation Personalization (Part 2)."; Redress Compliance, "History of Robo-Advisors and Automated Investing."; MarketPulse, "The AI Revolution in Investing: How Generative Tools Are Empowering Retail Investors."

39 The Consumer Financial Protection Bureau estimated that the number of successful or attempted data pulls was between 50 and 100 billion in 2022. Consumer Financial Protection Bureau, "Required Rulemaking on Personal Financial Data Rights." Federal Register 89, no. 222 (November 18, 2024): 90958. For background *see* FinRegLab, "The Use of Cash-Flow Data in Underwriting Credit: Market Context & Policy Analysis," § 4.2.

40 PYMNTS, "Agentic AI Is Ready. Are Banks Prepared to Use It?"; Cocheo, "Banks Are Swimming in Data But Starving for Insights. AI Will Make Things Worse." Employers, payroll processors, and government agencies also vary in the extent that they provide users with internet platforms or (more recently) APIs for data access. Current regulatory efforts discussed in main text are focused on implementing a law that applies to data access concerning "consumer financial products and services," but have not yet addressed the extent to which these sources may be covered. Consumer Financial Protection Bureau, "Required Rulemaking on Personal Financial Data Rights." Federal Register 89, no. 222 (November 18, 2024): 90958. Some payroll processors are licensed as money transmitters under state law. Barnett and Polanco, "Payments Pros – Navigating New Compliance Challenges: The Impact of the Money Transmitter Modernization Act on Payroll Processing."

41 In late 2024, the Consumer Financial Protection Bureau finalized rules governing the transfer of data about transaction and credit card accounts under Section 1033 of the Dodd-Frank Wall Street Reform and Consumer Financial Protection Act of 2010. Consumer Financial Protection Bureau, "Required Rulemaking on Personal Financial Data Rights." Federal Register 89, no. 222 (November 18, 2024): 90958. Several banking organizations sued to challenge the law and the CFPB under the Trump Administration requested that the rule be vacated as unlawful in spring 2025, although a fintech trade association intervened to defend the rule. While briefing was underway, some banks took steps or expressed interest in charging for data access. Weinberger and Smith, "JPMorgan Tells FinTechs to Pay Up for Customer Data Access."; Rodrigues, "Winklevoss Claims JPMorgan Halted Gemini Onboarding After Data Access Fees Criticism."; Wang and Smith, "PNC Considers Charging Fintechs for Access to Customer Data." The CFPB subsequently obtained a stay in the litigation and released an Advanced Notice of Proposed Rulemaking in August 2025. Consumer Financial Protection Bureau, "Personal Financial Data Rights Reconsideration." Federal Register 90, no. 161 (August 22, 2025): 40986; Berry, "CFPB Revamps 1033 Open Banking Rule with New Focus on Fees."; Griffin, "Payments: Looking Forward." American Banker, "Visa Abandons Open Banking in U.S. as Data Access Debate Rages."

42 Most aggregators are fintechs, though some have been purchased by other types of financial services providers and some initiatives have been started by banks and other traditional financial institutions. One recent initiative is the Customer Data Clearinghouse, pioneered by a startup called Solo partly funded by BankVentures, a venture capital firm funded by community banks that invest in fintechs. The network allows participants to determine when a customer has information on file with other member banks and sometimes their fintech partners and seek permission to access the information for identity proofing, lending, or other purposes. *See* Crosman. "Inside Solo's New Bank-Led Alternative to Data Aggregators."

43 For background *see* FinRegLab, "The Use of Cash-Flow Data in Underwriting Credit: Market Context & Policy Analysis," § 4.2; Financial Health Network et al., "Consumer Financial Data: Legal and Regulatory Landscape."

44 Consumer Financial Protection Bureau, "Required Rulemaking on Personal Financial Data Rights." Federal Register 89, no. 222 (November 18, 2024): 90958.

45 Feldsher, "Data Middlemen's Free Ride Is Driving Data Misuse." (reporting that monthly pulls have increased to 2 million in in the past two years at JPMorgan Chase).

46 Berry, "CFPB Revamps 1033 Open Banking Rule with New Focus on Fees."

47 *See* Note 3 for a more detailed discussion of MCP and APIs.

48 FinRegLab, "The Use of Cash-Flow Data in Underwriting Credit: Market Context & Policy Analysis," § 4.2.

49 *See, e.g.*, Watkins, "Scaling Web Scraping with Data Streaming, Agentic AI, and GenAI."; Cunningham, "How to Scrape Data from Hundreds of Banks with AI Agents."

50 Zeff, "Cloudflare Launches a Marketplace That Lets Websites Charge AI Bots for Scraping."; Cloudflare, "Cloudflare Just Changed How AI Crawlers Scrape the Internet-at-Large; Permission-Based Approach Makes Way for A New Business Model."

51 For example, since two-factor authentication generally requires human involvement to satisfy, it would potentially limit the function of AAI systems without more active consumer involvement.

52 Mastercard, "Mastercard Unveils Agent Pay, Pioneering Agentic Payments Technology to Power Commerce in the Age of AI."

53 Board of Governors of the Federal Reserve System, Supervisory & Regulation Letter 11-7, "Guidance on Model Risk Management."; Office of the Comptroller of the Currency, Bulletin 2011-12, "Sound Practices for Model Risk Management: Supervisory Guidance on Model Risk Management."; Federal Deposit Insurance Corporation, Financial Institution Letter 22-2017 "Adoption of Supervisory Guidance on Model Risk Management." For general background on MRM expectations particularly as applicable to use of AI in credit underwriting, *see* FinRegLab, "Explainability & Fairness in Machine Learning for Credit Underwriting: Policy Analysis."

54 *See, e.g.*, National Institute of Standards and Technology, "NIST AI Risk Management Framework Playbook."; Richards, "Comment on Artificial Intelligence Risk Management Framework Request for Information to NIST." For citations to the bank regulatory guidance in the securities context, *see, e.g.*, Securities and Exchange Commission, "Self-Regulatory Organizations; National Securities Clearing Corporation; Notice of Filing of Proposed Rule Change to Modify the Clearing Agency Model Risk Management Framework."; FINRA, "Artificial Intelligence (AI) in the Securities Industry."; Deloitte, "Managing Model and AI risks in the Investment Management Sector: Understanding and Mitigating Model Risk."

55 The Government Accountability Office recently called on the National Credit Union Administration to expand its MRM guidance to more closely parallel the other banking agencies on related topics in light of the growing use of machine learning and AI applications. U.S. Government Accountability Office, "Artificial Intelligence: Use and Oversight in Financial Services."

56 Chorev et al., "Agentic AI in Financial Services: Opportunities, Risks, and Responsible Implementation."

57 *See* Section 6.

58 Rogers, "AI Financial Advisors Target Young People Living Paycheck to Paycheck."

59 Lucinity, "Understanding Ethical Agentic AI in Compliance - Transform Fincrime Operations & Investigations with AI."

60 Hammond et al., "Multi-Agent Risks from Advanced AI."; Chorev et al., "Agentic AI in Financial Services: Opportunities, Risks, and Responsible Implementation."

61 Chorev et al., "Agentic AI in Financial Services: Opportunities, Risks, and Responsible Implementation."

62 Choi et al., "The Role of Privacy Fatigue in Online Privacy Behavior."

63 Ma and Chen, "Are Digital Natives Overconfident in Their Privacy Literacy? Discrepancy Between Self-Assessed and Actual Privacy Literacy, and Their Impacts on Privacy Protection Behavior."

64 15 U.S.C. § 1666 et seq.; 12 C.F.R. §§ 1026.2, 1026.13.

65 15 U.S.C. § 1693 et seq.; 12 C.F.R. §§ 1005.2, 1005.6, 1005.11.

66 For instance, Regulation Z excludes credit transactions made by "a person who has actual, implied, or apparent authority" to use the consumer's credit card from the definition of billing error. 12 C.F.R. § 1026.13(a)(1). Regulation E generally defines unauthorized transactions as "an electronic fund transfer from a consumer's account initiated by a person other than the consumer without actual authority to initiate the transfer and from which the consumer receives no benefit." However, in cases where a consumer provides their debit card or another access device to a family member or friend who then exceeds the authority granted to them by the consumer, the consumer is liable unless and until the financial institution is notified that transfers by that person are no longer authorized. 12 C.F.R. § 1005.2(m).

67 *See, e.g.*, Brown, "Paying by Robot: Clearing the Regulatory Obstacles to Agentic Payments."

68 *See, e.g.*, Carballo, "Why Pairing Stablecoins with Agentic AI Will Revolutionize Global Finance,"; Cocheo, "Stablecoin and AI Agents Will Reinvent Banking, According to a Crypto Pioneer."; Allaire, "AI and Stablecoins: A Pairing for a More Intelligent Era of Online Business."; Kim and Dalal, "Why Agentic Commerce Needs Crypto to Scale."; Artemis et al., "Stablecoin Payments from the Ground Up."; Demos, "How Visa and Mastercard Can Survive the Stablecoin Threat."; Gerety, "AI Agents and the Future of Agentic Payments."; Torculas, "Agentic Commerce Will Be Built on Stablecoin Rails."

69 15 U.S.C. § 1693a; 12 C.F.R. §§ 1005.2, 1005.3. The Consumer Financial Protection Bureau proposed an interpretive rule in January 2025 that would have clarified that transfers of stablecoins, cryptocurrencies, and virtual currencies are subject to both EFTA and its implementing regulations, but it was withdrawn by the Trump Administration a few months later. 90 Fed. Reg. 3723 (Jan. 15, 2025); 90 Fed. Reg. 20,568 (May 15, 2025). A few federal district court opinions touch on the application of Regulation E to transfers and accounts holding digital assets. *See, e.g.*, Rider v. Uphold HQ Inc., 657 F.Supp.3d 491 (S.D.N.Y. 2023) (digital assets are "funds"); *see also* Nero v. Uphold HQ Inc., 688 F.Supp.3d 134 (S.D. N.Y. 2023) (account holding assets for personal investment was covered by Regulation E); but *see* Yuille v. Uphold HQ Inc., 686 F.Supp. 3d 323 (S.D.N.Y. 2023) (account holding assets for investment and profit-making purposes not covered by Regulation E).

70 Moran and Cornelius, "New Frontiers: CFPB Proposes Extending Consumer Protections to Other Digital Payment Mechanisms: Insights."; Carbone. "TDC Responds to CFPB."

71 Ecommerce North America, "Coinbase Payments Puts USDC at Checkout: Lower Fees, Faster Settlement for Ecommerce."

72 Heeb et al., "Walmart and Amazon Are Exploring Issuing Their Own Stablecoins."; Ecommerce North America. "Are Stablecoins Ecommerce's Next Big Thing?"

73 Demos, "Payments Companies Are Playing Hide-the-Crypto."; Finextra, "Visa Unveils Agentic Commerce and Stablecoin Plays."; JPMorgan, "JPMorganChase and Coinbase Launch Strategic Partnership to Make Buying Crypto Easier than Ever."; Macheel, "Coinbase Beefs Up Subscription Plan by Offering it with American Express Credit Card." Some partnerships between banks and cryptocurrency firms have also involved the use of tokenized deposits, which are held by banks in traditional accounts but recorded on blockchain systems to facilitate their instantaneous transfers without the need for intermediaries. *See, e.g.*, Browne, "JPMorgan Moves Further into Crypto with Stablecoin-Like Token JPMD."; Ozcan et al., "Digital Tokens: A Foundation for Stable Digital Money."

74 For a breakdown of major components of the GENIUS Act, *see* Goldman et al., "What the GENIUS Act Means for Payment Stablecoin Issuers, Banks, and Custodians."

75 Levy, "End User License Agreements: What is an EULA? and When to Use It?"; MacGregor, "Key Considerations When Building and Integrating AI Models in the Financial Services Sector."; Alberts et al., "Artificial Intelligence: The 'Black Box' of Product Liability."; Bosman et al., "Software Gains New Status as a Product Under Strict Liability Law."; Scoblete and Acevedo, "GenAI Product Liability Cases Are Making Their Way Through the Courts. Here's What We're Watching."; Smith et al., "Liability for Harms from AI Systems."

76 *See, e.g.*, Brown, "Paying by Robot: Clearing the Regulatory Obstacles to Agentic Payments." Some stakeholders also raised concerns about the risk that consumers could be penalized for agentic transactions for which they have sufficient account balances at the time of authorization but not at the time of settlement where particular payment rails are batched rather than verified in real time. This issue is not unique to agentive systems but may be more likely to occur if they are structured to operate without having to obtain express prior authorization from consumers for individual transactions and transfers.

77 Although it is not focused specifically on AAI systems, the litigation over the Consumer Financial Protection Bureau's customer permissioned data sharing rule is focusing in part on what types of agents, trustees, and representatives should be treated as standing in the shoes of the consumer under the Dodd-Frank Act. Forcht Bank, N.A. et al., Plaintiff's Brief in Support of Their Motion for Summary Judgment, Forcht Bank N.A. v. Consumer Financial Protection Bureau, No. 5:24-cv-304-DCR, May 30, 2025.

78 Imperva, "2025 Bad Bot Report: The Rapid Rise of Bots and the Unforeseen Risk for Business."

79 For example, experiments suggest that individual components can be nudged into attempting to "jailbreak" overseer agent programs by exposing them to information about how to engage in such behavior. Hammond et al., "Multi-Agent Risks from Advanced AI." The experiment involved a large language model that was exposed to material about how LLMs can be manipulated using specific prompt injections but not specifically directed to engage in such behavior. When the LLM was then instructed to write movie reviews that would receive the highest possible score, it spontaneously engaged in manipulative behavior in about 4% of cases when it was told it would be reviewed by another LLM but did not make any such attempts when it was told it would be subject to human review.

80 Model inversion attacks are a type of privacy breach where a bad actor feeds specific queries to a model and then takes those outputs to train another model. This new model attempts to recreate the training data from the original model, which can compromise data security and anonymity.

81 Jin et al., "Avoiding Embarrassment Online: Response to and Inferences About Chatbots When Purchases Activate Self Presentation Concerns."; Branley-Bell et al., "Chatbots for Embarrassing and Stigmatizing Conditions: Could Chatbots Encourage Users to Seek Medical Advice?"

82 Smith, "Attitudes Toward Algorithmic Decision-Making."; Treyger et al., "Assessing and Suing an Algorithm: Perceptions of Algorithmic Decision-Making."

83 Chak et al., "Robo-Advice for Borrower Repayment Decisions."; Chak et al., "Can Robo-Advice Improve Borrower Repayment Decisions?"

84 De la Cruz et al., "Generative AI and Emerging Technology."; Cross, "Chatbots Offer Advice Without Judgment. Low-Income People Are Noticing."; Tomaszewska et al., "Emerging Technology for All."

85 *See* Commonwealth, "Financial AI for Good: Guide & Chatbot." (based on a national survey of approximately 3,000 people, including both households living on LMI and those with higher incomes).

86 Chorev et al., "Agentic AI in Financial Services: Opportunities, Risks, and Responsible Implementation."

87 INET Oxford, "Agent Based Modelling Comes of Age."

88 Assad et al., "Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market."

89 *See, e.g.*, Barr, "Speech by Vice Chair for Supervision Barr on Artificial Intelligence in the Economy and Financial Stability."

90 Chorev et al., "Agentic AI in Financial Services: Opportunities, Risks, and Responsible Implementation."

91 Taylor et al., "The Agentic Oversight Framework: Procedures, Accountability and Best Practices for Agentic AI Use in Regulated Financial Services."

92 *See, e.g.*, Shukla, "Adaptive Monitoring and Real-World Evaluation of Agentic AI Systems."; Arike et al., "Technical Report: Evaluating Goal Drift in Language Model Agents."; Moshkovich et al., "Beyond Black-Box Benchmarking: Observability, Analytics, and Optimization of Agentic Systems."; National Institute of Standards and Technology, "SP 800-53 Control Overlays for Securing AI Systems Concept Paper."

93 Executive Office of the President, "Executive Order 14179 of January 23, 2025: Removing Barriers to American Leadership in Artificial Intelligence." Federal Register 90, no. 20 (January 31, 2025): 874.

94 U.S. Department of Commerce Bureau of Industry and Security, "Department of Commerce Announces Recission of Biden-Era Artificial Intelligence Diffusion Rule, Strengthens Chip-Related Export Controls."

95 The White House, "American Public Submits over 10,000 Comments on White House's AI Action Plan."

96 Executive Office of the President. "Winning the Race: America's AI Action Plan."

97 The Center is a successor to the AI Safety Institute, with some changes in focus. *See* U.S. Department of Commerce, "Statement from U.S. Secretary of Commerce Howard Lutnick on Transforming the U.S. AI Safety Institute into the Pro-Innovation, Pro-Science U.S. Center for AI Standards and Innovation."; National Institute of Standards and Technology, "Center for AI Standards and Innovation."

98 Executive Office of the President, "Winning the Race: America's AI Action Plan."

99 National Conference of State Legislatures, "Artificial Intelligence 2025 Legislation."

100 Business Software Alliance, "2025 State AI Wave Building after 700 Bills in 2024."

101 After the failure of SB 1047 in 2024, Governor Gavin Newsom commissioned a new AI Working Group that includes former California Supreme Court justices and leading AI policy scholars. The group's final report in June 2025 outlines eight key principles for AI governance, including balancing benefits and risks, grounding AI policy in evidence, emphasizing early design choices, transparency and adverse event reporting, and scoping to ensure that regulation is proportionate to risks). Joint California Policy Working Group on AI Frontier Models, "The California Report on Frontier AI Policy."; Allen and Hill, "AI Action Plan RFI, California's AI Policy Working Group Report, and Why Programming Jobs Are Disappearing."

102 Utah's legislation focuses primarily on consumer-facing generative AI, requiring disclosure requirements to inform consumers that they are interacting with artificial intelligence and clarifying that existing state consumer protections apply to such systems. Utah Policy, "Utah Paves the Way in AI Governance—A National Blueprint?" Although the law took effect in May 2024, three pieces of legislation were enacted in 2025 to refine and broaden the law in various respects, including requiring mechanisms to review or correct information used to inform "high risk" interactions. Gluck, "Overview of Utah's 2025 Enacted AI Legislation."; Felz and Everett, "New Artificial Intelligence Laws in Effect in Utah."; Bajowala and Goker, "Utah Enacts First AI-Focused Consumer Protection Legislation in US."

103 Rice et al.,"The Colorado Artificial Intelligence Act: FPF U.S. Legislation Policy Brief."

104 Glasser et al., "Colorado's Historic AI Law Survives without Delay (So Far)."

105 Levi et al., "Colorado's Landmark AI Act: What Companies Need to Know."

106 Navetta and Stauss, "Colorado August Special Session Will Address Colorado AI Act." Melendez et al., "Texas Charts New Path on AI with Landmark Regulation."

107 Amodei, "Anthropic CEO: Don't Let AI Companies off the Hook."; Tene et al., "Federal AI Moratorium Dies on the Vine as Senate Passes the 'One Big Beautiful Bill.'"

108 *See, e.g.*, American Action Forum, "AI Legislation Tracker."

109 *See, e.g.*, White and Case, "AI Watch: Global Regulatory Tracker."

Bibliography

Adebayo, Kolawole Samuel. "Here Comes the Big, Strong Agentic AI Wave." *Forbes*, February 5, 2025.

AdMetrics. "Agentic Commerce: How AI Agents Are Reshaping Ecommerce Forever." Accessed August 14, 2025.

Ahmad, Yasmeen. "New Agents and AI Foundations for Data Teams." Google Cloud Blog, August 6, 2025.

Alberts, Patricia, Paul Calfo, and Jean Gabat. "Artificial Intelligence: The 'Black Box' of Product Liability." Husch Blackwell, April 4, 2025.

Ali, Mohammad Asif. "Efficient Underwriting Using Agentic AI." *Software Engineering*, May 30, 2025.

Allaire, Jeremy. "AI and Stablecoins: a Pairing for a More Intelligent Era of Online Business." World Economic Forum, January 16, 2025.

Allen, Gregory C. and Brielle Hill. "AI Action Plan RFI, California's AI Policy Working Group Report, and Why Programming Jobs Are Disappearing." CSIS, The AI Policy Podcast, March 26, 2026.

American Action Forum. "AI Legislation Tracker." Accessed August 25, 2025.

American Banker. "Visa Abandons Open Banking in U.S. as Data Access Debate Rages." August 22, 2025.

Amodei, Dario. "Anthropic CEO: Don't Let AI Companies off the Hook." *New York Times*, June 5, 2025.

Anthropic. "Claude for Financial Services." July 15, 2025.

Anthropic. "Introducing the Model Context Protocol." November 25, 2024.

Arike, Rauno, Elizabeth Donoway, Henning Bartsch, and Marius Hobbhahn. "Technical Report: Evaluating Goal Drift in Language Model Agents." arXiv:2505.02709v1, May 5, 2025.

Artemis, Castle Island Ventures, and Dragonfly. "Stablecoin Payments from the Ground Up." May 2025.

Assad, Stephanie, Robert Clark, Daniel Ershov, and Lei Xu. "Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market." *Journal of Political Economy* 132, no. 3 (2024): 723-771.

Bajowala, Reena, and Arda Goker. "Utah Enacts First AI-Focused Consumer Protection Legislation in US." *Insights: Greenberg Traurig LLP*, April 1, 2024.

Balasubramanian, Ramnath, Ari Libarikian, and Doug McElhaney. "Insurance 2030—The Impact of AI on the Future of Insurance." McKinsey & Company, March 12, 2021.

Barnett, Keith, and Eli Polanco. "Payments Pros—Navigating New Compliance Challenges: The Impact of the Money Transmitter Modernization Act on Payroll Processing." *Troutman Pepper Podcast Transcript,* July 25, 2024.

Barr, Michael. "Speech by Vice Chair for Supervision Barr on Artificial Intelligence in the Economy and Financial Stability." Board of Governors of the Federal Reserve System, Feb. 18, 2025.

Belcic, Ivan, and Cole Stryker. "What Is Agentic RAG?" IBM, June 18, 2025.

Berry, Kate. "CFPB Revamps 1033 Open Banking Rule with New Focus on Fees." *American Banker*, August 21, 2025.

Blader, Ruth Foxe. "AI Is Supercharging Parametric Insurance." *Forbes*, July 2, 2024.

Board of Governors of the Federal Reserve System. Supervisory & Regulation Letter 11-7, "Guidance on Model Risk Management." April 4, 2011.

Bosman, Erin, Julie Park, Matt Robinson, and Rachel Kaiser. "Software Gains New Status as a Product Under Strict Liability Law." Morrison Foerster, June 18, 2025.

Branley-Bell, Dawn, Richard Brown, Lynne Coventry, and Elizabeth Sillence. "Chatbots for Embarrassing and Stigmatizing Conditions: Could Chatbots Encourage Users to Seek Medical Advice?" *Frontiers in Communication: Health Communication*, September 26, 2023.

Brown, Thomas. "Paying by Robot: Clearing the Regulatory Obstacles to Agentic Payments." May 28, 2025.

Browne, Ryan. "JPMorgan Moves Further into Crypto with Stablecoin-Like Token JPMD." *CNBC*, June 17, 2025.

Business Software Alliance. "2025 State AI Wave Building after 700 Bills in 2024." October 22, 2024.

Carballo, Ignacio E. "Why Pairing Stablecoins with Agentic AI Will Revolutionize Global Finance." *Payments & Commerce Market Intelligence*, July 15, 2025.

Carbone, Cody. "TDC Responds to CFPB." The Digital Chamber, July 30, 2025.

Castri, Simone di, Matt Grasser, and Juliet Ongwae. "State of SupTech Report 2023." Accessed June 20, 2025.

Chak, Ida, Karen Croxson, Francesco D'Acunto, Jonathan Reuter, Alberto Rossi, and Jonathan Shaw. "Can Robo-Advice Improve Borrower Repayment Decisions?" Financial Conduct Authority, August 30, 2022.

Chak, Ida, Karen Croxson, Francesco D'Acunto, Jonathan Reuter, Alberto Rossi, and Jonathan Shaw, "Robo-Advice for Borrower Repayment Decisions." Financial Conduct Authority Occasional Paper 61, August. 2022.

Chawla, Avi, and Akshay Pachaar. "Implementing ReAct Agentic Pattern from Scratch." *Daily Dose of Data Science*. April 13, 2025.

Chilson, Neal. "The Vibrant AI Competitive Landscape." *House Judiciary Committee*, April 2, 2025.

Choi, Hanbyul, Jonghwa Park, and Yoonhyuk Jung. "The Role of Privacy Fatigue in Online Privacy Behavior." *Computers in Human Behavior* 81 (2018): 42-51.

Chorev, Michael, Joe Royle, and Richie Paul. "Agentic AI in Financial Services: Opportunities, Risks, and Responsible Implementation." *IBM Consulting*. July 2025.

Cloudflare, "Cloudflare Just Changed How AI Crawlers Scrape the Internet-at-Large; Permission-Based Approach Makes Way for A New Business Model." July 1, 2025.

Cocheo, Steve. "Banks Are Swimming in Data But Starving for Insights. AI Will Make Things Worse." The Financial Brand, May 9, 2025.

Cocheo, Steve. "Stablecoin and AI Agents Will Reinvent Banking, According to a Crypto Pioneer." The Financial Brand, July 14, 2025.

Columbus, Louis. "DanaBot Takedown Shows How Agentic AI Cut Months of SOC Analysis to Weeks." *VentureBeat*, May 28, 2025.

Commonwealth. "Financial AI for Good: Guide & Chatbot." May 21, 2024.

Consumer Financial Protection Bureau, "Personal Financial Data Rights Reconsideration." Federal Register 90, no. 161 (August 22, 2025): 40986.

Consumer Financial Protection Bureau. "Required Rulemaking on Personal Financial Data Rights." Federal Register 89, no. 222 (November 18, 2024): 90958.

Consumer Financial Protection Bureau, "Electronic Fund Transfers Through Accounts Established Primarily for Personal, Family, or Household Purposes Using Emerging Payment Mechanisms." Federal Register 90, no. 9 (January 15, 2025): 3723.

Consumer Financial Protection Bureau, "Electronic Fund Transfers Through Accounts Established Primarily for Personal, Family, or Household Purposes Using Emerging Payment Mechanisms; Withdrawal." Federal Register 90, no. 93 (May 15, 2025): 20568.

Crosman, Penny. "Inside Solo's New Bank-Led Alternative to Data Aggregators." American Banker, July 30, 2025.

Cross, Miriam. "Chatbots Offer Advice Without Judgment. Low-Income People Are Noticing." American Banker, October 8, 2021.

Cunningham, Michael. "How to Scrape Data from Hundreds of Banks with AI Agents." Medium, July 21, 2025.

De la Cruz, Charles, Paula Grieco, David Kibbe, and Taylor Straatmann. "Generative AI and Emerging Technology." Commonwealth, April 21, 2025.

Deloitte. "Managing Model and AI Risks in the Investment Management Sector: Understanding and Mitigating Model Risk." *Perspective*, 2023.

Demos, Telis. "How Visa and Mastercard Can Survive the Stablecoin Threat." Wall Street Journal, June 24, 2025.

Demos, Telis. "Payments Companies Are Playing Hide-the-Crypto." Wall Street Journal, March 31, 2021.

Dien, Dimitille. "Agentic AI for Insurance Agents & Bank Advisors—The Complete Guide." Zelros. January 17, 2025.

Ecommerce North America. "Coinbase Payments Puts USDC at Checkout: Lower Fees, Faster Settlement for Ecommerce." June 20, 2025.

Ecommerce North America. "Test Your Ecommerce Site for AI Agent Accessibility." July 29, 2025.

Executive Office of the President. "Executive Order 14179 of January 23, 2025: Removing Barriers to American Leadership in Artificial Intelligence." Federal Register, 90, no. 80 (January 31, 2025): 874.

Executive Office of the President. "Winning the Race: America's AI Action Plan." July 23, 2025.

Federal Deposit Insurance Corporation. Financial Institution Letter 22-2017. "Adoption of Supervisory Guidance on Model Risk Management." June 7, 2017.

Feldsher, Melissa. "Data Middlemen's Free Ride Is Driving Data Misuse." Real Clear Markets, July 19, 2025.

Felz, Daniel, and Jennifer Everett. "New Artificial Intelligence Laws in Effect in Utah." *Alston & Bird Privacy, Cyber & Data Strategy Blog*, June 13, 2025.

Financial Health Network, Flourish, FinRegLab, and Mitchell Sandler. "Consumer Financial Data: Legal and Regulatory Landscape." October 2020.

Finextra. "Visa Unveils Agentic Commerce and Stablecoin Plays." May 1, 2025.

FINRA. "Artificial Intelligence (AI) in the Securities Industry." June 10, 2020.

FinRegLab. "Explainability & Fairness in Machine Learning for Credit Underwriting: Policy Analysis." December 2023.

FinRegLab. "Innovations for Identity Proofing and Transaction Monitoring: Advancing Financial Inclusion Through Data & Technology." October 2024.

FinRegLab. "The Use of Cash-Flow Data in Underwriting Credit: Market Context & Policy Analysis." February 2020.

Forcht Bank, N.A., Kentucky Bankers Association, and Bank Policy Institute, Plaintiff's Brief in Support of Their Motion for Summary Judgment, Forcht Bank N.A. v. Consumer Financial Protection Bureau, No. 5:24-cv-304-DCR, May 30, 2025.

Gerety, Amias. "AI Agents and the Future of Agentic Payments." QED Investors Blog, May 20, 2025.

Ghose, Ronit, Sophia Bantanidis, Ronak S. Shah, Kaiwan Master, Savina Chahal, Prag Sharma, and Catherine Zhai. "Agentic AI: Finance and the 'Do It For Me' Economy." Citi Institute Global Perspectives and Solutions, January 2025.

Glasser, Nathaniel M., Eleanor T. Chung, Adam S. Forman, Rachel Snyder Good, and Alaap B. Shah. "Colorado's Historic AI Law Survives without Delay (So Far)." *The National Law Review*, May 13, 2025.

Gluck, Justine. "Overview of Utah's 2025 Enacted AI Legislation." fpf.org. Accessed August 8, 2025.

Goldman, Zachary, Tiffany J. Smith, Matthew B. Kulkin, Jeremy Moorehouse, Jeffrey Wieand, Daniel LaMagna, and Ben Gardiner. "What the Genius Act Means for Payment Stablecoin Issuers, Banks, and Custodians." WilmerHale, July 18, 2025.

Gosby, Desiree. "Walmart: The Future of Shopping Is Agentic. Meet Sparky." Walmart, June 6, 2025.

Griffin, Tyler. "Payments: Looking Forward." Restive Venture Capital, August 19, 2025.

Gutowska, Anna. "What Is Model Context Protocol (MCP)?" IBM Think. Accessed August 25, 2025.

Habicht, Johanna, Sruthi Viswanathan, Ben Carrington, Tobias U. Hauser, Ross Harper, and Max Rollwage. "Closing the Accessibility Gap to Mental Health Treatment with a Personalized Self-Referral Chatbot." *Nature Medicine* 30, no. 2 (2024): 595-602.

Hammond, Lewis, Alan Chan, Jesse Clifton, Jason Hoelscher-Obermaier, Akbir Khan, Euan McLean, Chandler Smith et al. "Multi-Agent Risks from Advanced AI." *arXiv preprint arXiv:2502.14143* (2025).

Heeb, Gina, AnnaMaria Andriotis, and Josh Dawsey. "Walmart and Amazon Are Exploring Issuing Their Own Stablecoins," Wall Street Journal, June 13, 2025.

Imperva. "2025 Bad Bot Report: The Rapid Rise of Bots and the Unforeseen Risk for Business." 2025.

INET Oxford. "Agent Based Modelling Comes of Age." March 21, 2025.

Jackson, Sarah. "Nvidia's CEO Says We're in the Age of 'Agentic' AI—Here's What That Word Means." *MSN*, January 14, 2025.

Jin, Jianna, Jesse Walker, and Rebecca Walker Reczek. "Avoiding Embarrassment Online: Response to and Inferences about Chatbots when Purchases Activate Self-Presentation Concerns." *Journal of Consumer Psychology* 35, no. 2 (2025): 185-202.

Joint California Policy Working Group on AI Frontier Models. "The California Report on Frontier AI Policy." June 17, 2025.

Josef, Rami. "Agentic Commerce: Where Consumer Intent Meets Merchant Opportunity." *Checkout.com*, June 6, 2025.

JPMorgan Chase & Co. "JPMorganChase and Coinbase Launch Strategic Partnership to Make Buying Crypto Easier than Ever." July 30, 2025.

Juristech, "Why Deterministic Agentic AI Is the Breakthrough Modern Banking Needs."

Kaliski, Steve. "Adding Payments to Your LLM Agentic Workflows." Stripe, November 14, 2024.

Kim, Dan and Nemil Dalal. "Why Agentic Commerce Needs Crypto to Scale." July 2, 2025.

Levi, Stuart, Ken D. Kumayama, William E. Ridgway, Mana Ghaemmaghami, and MacKinzie M. Neal. "Colorado's Landmark AI Act: What Companies Need to Know." Skadden, Arps, Slate, Meagher & Flom LLP, June 24, 2024.

Levitt, Kevin. "AI On: How Financial Services Companies Use Agentic AI to Enhance Productivity, Efficiency, and Security." NVIDIA. July 22, 2025.

Levy, Colin. "End User License Agreements: What is an EULA? and When to Use It?" December 16, 2024.

Lucinity. "Understanding Ethical Agentic AI in Compliance - Transform Fincrime Operations & Investigations with AI." Lucinity, January 23, 2025.

Ma, Shuai, and Chen Chen. "Are Digital Natives Overconfident in Their Privacy Literacy? Discrepancy Between Self-Assessed and Actual Privacy Literacy and Their Impacts on Privacy Protection Behavior." *Frontiers in Psychology* 14 (2023): 1224168.

MacGregor, Mardi. "Key Considerations When Building and Integrating AI Models in the Financial Services Sector." Fox Williams, December 21, 2023.

Macheel, Tanaya. "Coinbase Beefs Up Subscription Plan by Offering it with American Express Credit Card." June 12, 2025.

MarketPulse. "The AI Revolution in Investing: How Generative Tools Are Empowering Retail Investors." July 26, 2025.

Mastercard. "Mastercard Unveils Agent Pay, Pioneering Agentic Payments Technology to Power Commerce in the Age of AI." Mastercard Newsroom, April 29, 2025.

Melendez, Maria Cruz, Stuart D. Levi, William E. Ridgway, and Brittany E. Libson. "Texas Charts New Path on AI with Landmark Regulation." Skadden, Arps, Slate, Meagher & Flom LLP, June 23, 2025.

Mistral AI. "Build AI Agents with the Mistral Agents API." May 27, 2025.

Montti, Roger. "OpenAI Quietly Adds Shopify as a Shopping Search Partner." *Search Engine Journal*, July 10, 2025.

Moran, Eamonn, and Ceijenia Cornelius. "New Frontiers: CFPB Proposes Extending Consumer Protections to Other Digital Payment Mechanisms: Insights." Holland & Knight, January 22, 2025.

Moshkovich, Dany et al. "Beyond Black-Box Benchmarking: Observability, Analytics, and Optimization of Agentic Systems." arXiv:2503.06745v1, March 9, 2025.

Muhn, Julie. "4 Companies Bringing Agentic AI to Checkout." *Finnovate Blog*, May 21, 2025.

Murel, Jacob and Eda Kavlakoglu. "What Is Reinforcement Learning?" IBM, March 25, 2024.

National Conference of State Legislatures. "Artificial Intelligence 2025 Legislation." Accessed August 25, 2025.

National Institute of Standards and Technology. "Center for AI Standards and Innovation." Accessed August 25, 2025.

National Institute of Standards and Technology. "NIST AI Risk Management Framework Playbook." Accessed August 25, 2025.

National Institute of Standards and Technology. "SP 800-53 Control Overlays for Securing AI Systems Concept Paper." August 14, 2025.

Navetta, David and David Stauss. "Colorado August Special Session Will Address Colorado AI Act." *JD Supra*, August 7, 2025.

Neidhart-Lau, Andreas. "Analyzing AI Workflows: Deterministic and Autonomous Agents in Financial Services." *The FinServ Edge*, April 24, 2026.

Office of the Comptroller of the Currency. Bulletin 2011-12. "Sound Practices for Model Risk Management: Supervisory Guidance on Model Risk Management."

OpenAI. "Introducing ChatGPT Agent: Bridging Research and Action." OpenAI, July 17, 2025.

Ozcan, Gokce, Jason Ekberg, Ugur Koyluoglu, Leo Sizaret, James Emmet, Naveen Mallela, Nelli Zaltsman, Basak Toprak, and Mario Tombazzi. "Deposit Tokens – a Foundation for Stable Digital Money." Oliver Wyman - Impact-Driven Strategy Advisors, February 13, 2023.

Pannala, Pranav. "Rebooting Robo-Advisory: Can the Industry Meet Investor Demands for Personalization? (Part 1)." EPAM.com, April 10, 2024.

Pannala, Pranav. "Rebooting Robo-Advisory: Leveraging Generative AI to Power next-Generation Personalization (Part 2)." EPAM.com, April 24, 2024.

PayPal. "Perplexity Selects PayPal to Power Agentic Commerce." *PayPal Newsroom*, May 14, 2025.

Poccia, Danilo. "Introducing Amazon Bedrock AgentCore: Securely Deploy and Operate AI Agents at Any Scale (Preview)." *AWS News Blog*, July 16, 2025.

Polonioli, Andrea. "Amazon's 'Buy for Me' Marks the Rise of Agentic Commerce—And It's Just the Beginning." *Coveo Blog*, April 15, 2025.

PYMNTS. "Agentic AI Is Ready. Are Banks Prepared to Use It?" *PYMNTS.com*, April 3, 2025.

PYMNTS. "Agentic AI Turns Enterprise Cybersecurity into Machine vs. Machine Battle." *PYMNTS.com*, July 22, 2025.

Redress Compliance. "History of Robo-Advisors and Automated Investing." July 26, 2025.

Rice, Tatiana, Keir Lamont, and Jordan Francis. "The Colorado Artificial Intelligence Act: An FPF U.S. Legislation Brief." Colorado General Assembly, July 2024.

Richards, Michael. "Comment on Artificial Intelligence Risk Management Framework Request for Information to NIST." U.S. Chamber of Commerce Technology Engagement Center, September 15, 2021.

Rincon, Lilian. "Shop with AI Mode, Use AI to Buy and Try Clothes on Yourself Virtually." *Google Blog*, May 20, 2025.

Rodrigues, Francisco. "Winklevoss Claims JPMorgan Halted Gemini Onboarding After Data Access Fees Criticism." *CoinDesk*, July 26, 2025.

Rogers, Reece. "AI Financial Advisors Target Young People Living Paycheck to Paycheck." *WIRED*, January 13, 2025.

Rosenbaum, Eric. "Meta Is Targeting 'Hundreds of Millions' of Businesses in Agentic AI Deployment." *CNBC*. March 6, 2025.

Ruffing, Ryan. "How AI and Parametric Models Are Revolutionizing Risk Protection for Crop Insurance." Arbol, May 13, 2025.

Samuel, Arthur L. "Some Studies in Machine Learning Using the Game of Checkers." 3 IBM Journal of Research & Development 211-229 (1959).

Sapkota, Ranjan, Konstantinos I. Roumeliotis, and Manoj Karkee. "AI Agents vs. Agentic AI: A Conceptual Taxonomy, Applications and Challenges." *arXiv:2505.10468v4*, May 2025.

Scoblete and Acevedo, "GenAI Product Liability Cases Are Making Their Way Through the Courts. Here's What We're Watching." May 29, 2025.

Securities and Exchange Commission. "Self-Regulatory Organizations; National Securities Clearing Corporation; Notice of Filing of Proposed Rule Change to Modify the Clearing Agency Model Risk Management Framework." (No. 34-88629; File No. SR-NSCC-2020-008, April 14, 2020.

Sharma, Asha. "New Capabilities in Azure AI Foundry to Build Advanced Agentic Applications." *Microsoft Blog*. April 4, 2025.

Shukla, Manish A. "Adaptive Monitoring and Real-World Evaluation of Agentic AI Systems." arxiv:2509.00115, August 25, 2025."

Silberstein, Nicole. "OpenAI, Shopify Reportedly Working on ChatGPT Checkout Integration." *Retail Touch Points*, July 17, 2025.

Smith, Aaron. "Attitudes Toward Algorithmic Decision-Making." Pew Research Center, November 16, 2018.

Smith, Allison. "Shopify Quietly Sets Boundaries for AI Agents on Merchant Sites." Modern Retail, July 14, 2025.

Smith, Gregory, Karlyn Stanley, Krystyna Marcinek, Paul Cormarie, and Salil Gunashekar. "Liability for Harms from AI Systems: The Application of U.S. Tort Law and Liability to Harms from Artificial Intelligence Systems." RAND, November 20, 2024.

Taylor, Simon, Soups Ranjan, Matt Vega, Ryan McCormack, and Erich Reich. "The Agentic Oversight Framework: Procedures, Accountability and Best Practices for Agentic AI Use In Regulated Financial Services." Sardine AI, 2025.

Tene, Omer, Bethany P. Withers, and Reema Moussa. "Federal AI Moratorium Dies on the Vine as Senate Passes the 'One Big Beautiful Bill.'" Goodwin, July 3, 2025.

The White House. "American Public Submits over 10,000 Comments on White House's AI Action Plan." April 24, 2025.

Tomaszewska, Gosia, Taylor Straatmann, Charles de la Cruz, Paula Grieco, and Zaanish Pirani. "Emerging Technology for All." Commonwealth, July 5, 2023.

Torculas, Dwight. "Agentic Finance: The Financial Infrastructure Powering AI Agent Commerce." *Crossmint Blog*, May 1, 2025.

Torculas, Dwight. "Agentic Commerce Will Be Built on Stablecoin Rails." *Crossmint Blog*, April 16, 2025.

Trelka, Damian. "Meet Opera's AI Browser Operator." *Opera News*, March 26, 2025.

Treyger, Elina, Jirka Taylor, Daniel Kim, and Maynard Holliday. "Assessing and Suing an Algorithm: Perceptions of Algorithmic Decision-Making." RAND, October 12, 2023.

U.S. Department of Commerce Bureau of Industry and Security. "Department of Commerce Announces Recission of Biden-Era Artificial Intelligence Diffusion Rule, Strengthens Chip-Related Export Controls." May 13, 2025.

U.S. Department of Commerce. "Statement from U.S. Secretary of Commerce Howard Lutnick on Transforming the U.S. AI Safety Institute into the Pro-Innovation, Pro-Science U.S. Center for AI Standards and Innovation." June 3, 2025.

U.S. Government Accountability Office. "Artificial Intelligence: Use and Oversight in Financial Services." May 19, 2025.

Utah Policy. "Utah Paves the Way in AI Governance—A National Blueprint?" June 20, 2025.

Vajpayee, Prashant and Gahangir Hossain. "Cyber Defense through Agentic AI Enabled Automation: An Approach to Reduce Cyber Risk." *SIGMIS-CPR '25: Proceedings of the 2025 Computers and People Research Conference*, June 13, 2025.

Vals.ai. "CorpFin (v2) Benchmark." August 8, 2025.

Vals.ai. "Financial Agent Benchmark." April 22, 2025.

Visa. "Find and Buy with AI: Visa Unveils New Era of Commerce." April 30, 2025.

Walmart. "Retail Rewired Report." 2025.

Wang, Yizhu and Paige Smith. "PNC Considers Charging Fintechs for Access to Customer Data." Bloomberg Law News, July 16, 2025.

Watkins, Adam. "Scaling Web Scraping with Data Streaming, Agentic AI, and GenAI." Confluent Blog, January 6, 2025.

Watkins, Luke. "Intelligent AI Underwriting's Impact on Modern Mortgage Lending." Deepset. March 25, 2025.

Weinberger, Evan and Paige Smith. "JPMorgan Tells FinTechs to Pay Up for Customer Data Access." Bloomberg, July 11, 2025.

Wells Fargo. "Comparing the Snowball and the Avalanche Methods of Paying Down Debt." Wells Fargo. Accessed August 14, 2025.

White and Case. "AI Watch: Global Regulatory Tracker." Accessed August 25, 2025.

Wolff, Rachel. "Amazon Implements Guardrails as AI Agents Threaten Traffic, Ad Revenues." EMARKETER, June 20, 2025.

Zeff, Maxwell. "Cloudflare Launches a Marketplace That Lets Websites Charge AI Bots for Scraping." *TechCrunch*, July 1, 2025.

Zhang, B.Z. "From Innovation Delta to Regulatory Singularity: How Innovative Regulatory Systems Can Help Regulation Keep Pace with Financial Innovation." Cambridge Centre for Alternative Finance, University of Cambridge Judge Business School, 2025.

Digital version available at finreglab.org

Published by FinRegLab, Inc.

1701 K Street NW, Suite 1150
Washington, DC 20006
United States