

September 2024

The Risks of Generative AI Agents to Financial Services

By Todd Phillips

About the Author

Todd Phillips is a fellow at the Roosevelt Institute and an assistant professor of law in the Robinson College of Business at Georgia State University. His areas of expertise include bank capital and prudential regulation, deposit insurance, derivatives and securities market structure, and the laws governing federal regulators. Before entering academia, Phillips served as an attorney for the Federal Deposit Insurance Corporation, the Administrative Conference of the United States, the US House of Representatives, and the Center for American Progress.

Phillips has testified before Congress, has served on the Commodity Futures Trading Commission's Market Risk Advisory Committee, and frequently advises financial regulators and members of Congress and their staffs. He has been published in the *Duke Law Journal*, the *Yale Journal on Regulation*, the *Administrative Law Review*, and the *Financial Times*, and is frequently quoted in the *New York Times*, *Washington Post*, *Politico*, and other outlets. Phillips holds a JD from the University of Michigan and a BS in economics and political science from Arizona State University.

Acknowledgments

The author thanks Emily DiVito for insights, research support, and project guidance. The author also thanks Lindsay Sain Jones and Mark Hays for their very helpful comments. Roosevelt staff Sonya Gurwitt and Hannah Groch-Begley also contributed to this brief.

About the Roosevelt Institute

The Roosevelt Institute is a think tank, a student network, and the nonprofit partner to the Franklin D. Roosevelt Presidential Library and Museum that, together, are learning from the past and working to redefine the future of the American economy. Focusing on corporate and public power, labor and wages, and the economics of race and gender inequality, the Roosevelt Institute unifies experts, invests in young leaders, and advances progressive policies that bring the legacy of Franklin and Eleanor Roosevelt into the 21st century.



Introduction

One day soon, you may be able to ask Siri to pay your cell phone bill from funds in your checking account. Or ask Alexa to recommend investments tailored to your risk profile. Or tell Gemini to manage your investing portfolio so you can travel in retirement. Such capabilities will be possible thanks to generative artificial intelligence, or Generative AI. Generative AI agents are computer systems with abilities to interpret and execute requests, such as these examples, without additional human interaction, and have been described as “the next frontier of Generative AI” ([Yee et al. 2024](#)). They have the potential to change the way individuals and firms interact with their banks and other financial services providers, opening the door to efficiency and economic growth—but also posing new risks to consumers, investors, and the safety and soundness of the financial system.

Generative AI agents may assist retail consumers with wealth management, such as working through personal assistants like Siri and Alexa to recommend financial products or serving as robo-advisors tailored to investors’ needs. They may also help nonfinancial firms through off-the-shelf treasury management solutions that balance financial returns with liquidity needs, and help financial institutions with risk management and regulatory compliance, such as with automated fraud detection, customer identity verification, and risk assessments through highly customized software ([Zheng et al. 2019](#); [Polak et al. 2020](#); [OECD 2021](#)).

Simultaneously, Generative AI agents threaten to destabilize the financial system, sending it swinging from crisis to crisis. Malicious actors can use it to defraud consumers, execute cyberattacks against financial institutions, and engage in market manipulation ([Fang et al. 2024b](#); [Mizuta 2020](#); [Hsu 2024](#)). Financial institutions’ own uses of Generative AI can hallucinate (that is, produce false or misleading outputs), resulting in harms to their customers, the institutions themselves, and the financial markets in which they operate ([CFTC 2024](#)). And when individuals and real-economy firms rely on a small number of Generative AI providers for financial decisions, they can engage in herding behavior that results in bank runs or flash crashes ([Gensler and Bailey 2020](#)).

Generative AI, in its current form, is not “good” or “bad” in and of itself. The large language models that power AI agents are simply computer systems capable of generating new content, such as images, text, audio, or video from a simple prompt. These systems are best considered a form of applied statistics—they capture patterns in the data upon which they have been trained and create outputs that resemble the training data but are unique variations. These software are just the latest use of algorithms and machine learning techniques that have been used in financial markets for decades, unlike in other industries in which Generative AI is novel, and can be used as *inputs* in human decision-making, as *copilots* that make decisions in coordination



with humans, and as *agents* that make decisions on behalf of humans ([US Senate Committee on Homeland Security and Governmental Affairs 2024](#); [Hsu 2024](#)).

To that end, the concern in this brief is not simply about the use of algorithms in finance, but about a world in which AI agents are widely available to individuals and small businesses as well as the largest financial firms; in which malicious actors may easily use Generative AI to scam financial institutions and their customers; and in which financial institutions use Generative AI to interact with their customers, rather than human employees. In particular, this brief is concerned about the harms that may result from individuals' and small businesses' AI agents interacting with large financial institutions and scammers—especially if they rely on developers' statements that AI agents will act in customers' interests.

Some United States financial regulators are already working to address the harms Generative AI poses. The Consumer Financial Protection Bureau (CFPB) has explained that federal law does “not permit creditors to use complex algorithms when doing so means they cannot provide the specific and accurate reasons for adverse actions” and has penalized financial institutions for relying on faulty automated compliance systems ([CFPB 2022a](#); [CFPB 2022b](#)). The federal banking agencies have created offices to study financial innovation and AI ([Phillips and Conner 2024](#)). The Securities and Exchange Commission (SEC) has proposed a regulation addressing brokers' uses of Generative AI and has begun examining investment advisers' uses of Generative AI for offering financial advice ([US Securities and Exchange Commission 2023](#)).

Nevertheless, more needs to be done, especially as it comes to the use of AI agents in the financial system. This brief highlights the expected rise of AI agents and the risks that their use by financial institutions, real-economy firms, and individuals pose to all aspects of the financial system and the families and businesses who rely on it. It concludes with recommendations to Congress and regulators.

Generative AI, AI Agents, and the Financial System

Algorithms have been used to price financial assets for generations, but new Generative AI software will likely provide financial institutions and their customers unprecedented capabilities. And because Generative AI can intake and “learn” from types of data that older generations of algorithms could not—including alternative data (i.e., data collected from nontraditional sources), unstructured data (i.e., audio, images, social media posts, etc.), and unlabeled data (i.e., data that lack classification)—these systems are now able to pull from more and different types of information than ever before ([Tierno 2024](#)).

Generative AI may allow financial institutions to operate more efficiently by automating and augmenting various aspects of their operations. For example, some banks and other lenders are already using AI models to assess credit worthiness and streamline the



lending decision process by using big data to predict default risks based on applicants' financial history, employment, and other factors,¹ and natural language processing can automatically extract and synthesize relevant information from sources like bank statements, tax returns, and other financial documents ([Sweet 2023](#)). Investment banks and financial advisors are leveraging Generative AI to create research reports, market commentary, and other content, as well as using AI to allocate clients' portfolios based on investors' risk profiles, investment goals, and market conditions ([Ritchie and Lee 2024](#)).

Furthermore, financial institutions may be able to use Generative AI to address legal compliance risks by having the software analyze legal documents, regulatory frameworks, and compliance protocols to identify potential risks and, perhaps, immediately make changes to financial activities without reliance on human input. Among the many potential compliance uses is relying on Generative AI to enhance the Know-Your-Customer processes required by law, automatically extracting and validating information from government-issued IDs, utility bills, bank statements, and other documents and rapidly analyzing and cross-referencing large volumes of data from multiple sources to verify an individual's identity claims. AI may also be used to detect potentially fraudulent documents or inconsistencies that may be missed by human reviewers and can continuously monitor customer transactions and behavior for anomalies that could indicate money laundering or other illicit activities.

Financial institutions may also use Generative AI to power conversational assistants for customer service, interpreting customers' natural language queries and providing human-like responses regarding inquiries related to products, services, policies, and account information. Though large financial institutions—including the largest banks—have used algorithms for these activities previously, the advances in Generative AI mean that now firms of all sizes can deploy AI chatbots that don't involve human input.

The rise of Generative AI may also allow—for the first time—financial institutions' customers to access previously unavailable AI technology. Customers can purchase Generative AI agents that directly interface with the AI and non-AI platforms used by banks and other lenders, allowing intelligent software agents to interact and execute transactions or data exchanges without the need for human involvement. For individual customers, this could mean asking Siri or Alexa to pay bills, transfer funds, or inquire about accounts and services, which are then executed through a behind-the-scenes dialogue with their institutions' AI agents. Moreover, AI systems can be powerful tools for automating and optimizing real-economy businesses' financial operations—including automating cash positioning, forecasting cash flows, managing liquidity, and optimizing working capital—with changes made through AI-to-AI connections with their financial institutions.

¹ Though this runs the risk of replicating biases embedded in datasets. For more, see [Polek and Sandy 2023](#).



Despite these and other benefits that Generative AI offers to the financial system and those who rely on it, increased access to and use of AI on the part of financial institutions and the public poses significant risks—potentially catastrophic to customers, financial institutions, and the entire financial system without sufficient regulation and oversight in place.

The wide availability of access to Generative AI agents may allow scammers and other malicious actors to steal money and sensitive information from financial institutions by impersonating customers, and to steal from customers by impersonating financial institutions. AI agents may also allow new opportunities for outright market manipulation. Alarmingly, recent research suggests that AI agents are able to conduct these kinds of malicious activities without direct human command or oversight ([Fang et al. 2024a](#); [Fang et al. 2024b](#); [Mizuta 2020](#)).

Furthermore, Generative AI agents deployed by financial institutions put customer money and business operations at risk. AI agents—like autonomous chatbots or systems to automate investment advice or capital management—can “hallucinate” false or misleading information, provide poor financial advice, or otherwise break down. Integrating AI into financial institutions’ software may harm customers by providing inaccurate information, inhibiting access to bank accounts, or failing to execute or improperly executing transactions. Customers who should be granted financial services may be rejected when Know-Your-Customer processes go awry, and those who would use services for illicit activities may be allowed access.

Finally, because Generative AI systems may react almost instantaneously to stimuli without human input, and the actions of individual AI agents may compound to calamitous results, such as “runs” on depository institutions and market-wide “flash crashes”—emblematic of the 2010 flash crash (discussed in the next section). These results could also be brought about by many AI agents illogically acting in tandem or entering into feedback loops with each other. These risks are of particular concern when AI agents operate in highly leveraged or highly illiquid environments. If a financial firm or market activities are disrupted because of an AI agent, there is almost no opportunity for human recourse, no matter how immediate, that can undo *all* of the damage. To that end, regulation that puts guardrails on AI use in the financial system is imperative, as human intervention is unlikely to be sufficiently swift to stop AI-caused runs and crashes once they start.

Financial System Risks from AI Agents

Risks from Malicious Actors' AI Agents

Equipped with AI agents capable of automating end-to-end complex tasks without human intervention or oversight, bad actors can now run increasingly sophisticated and convincing schemes and conduct cyberattacks that harm consumers, businesses, and financial firms alike. Though criminals have always found ways to defraud or hack vulnerable targets, Generative AI invites exponential increases in the volume, efficiency, and financial and personal costs to victims while simultaneously sowing the seeds of public distrust in the financial system ([Hsu 2024](#)). It also reduces barriers to entry for would-be cybercriminals, making it easier for a greater quantity and/or worse-trained individuals or coordinated groups to engage in harmful activities that would otherwise take significant time, effort, expertise, and human decision-making ([Chan et al. 2024](#)).

Because they have access to vast amounts of financial and customer information, financial systems and firms have always been a primary target for cyberattacks and scams. Indeed, perhaps the first cyberattack on a financial system occurred in 1834, when hackers co-opted a rudimentary version of the French telegraph system to send and receive proprietary financial market information ([Standage and Stevenson 2018](#)). Over the last two decades, nearly 20 percent of the known cyber incidents targeting institutions have affected financial firms, including banks, insurers, and asset managers ([International Monetary Fund 2024](#)). System hacks can directly disrupt firm operations, inhibiting or halting service provision and clearing and payment systems, as well as corrupting firm data (much of which is sensitive customer data). Since 2020, financial firms have lost nearly \$2.5 billion in cyber incidents, and analysts estimate that the threat of extreme loss in any given year is \$152 million ([International Monetary Fund 2024](#)). Losses that high—even before government fines and victims' compensation settlements are levied—could pose liquidity problems for most financial firms, even jeopardizing solvency for small and mid-sized firms. Due to the increasing interconnectedness of all firms, when stable operations at one firm are at risk, so too is the resilience of the entire US financial system ([International Monetary Fund 2024](#)). And, of course, consumer losses can be devastating to the individuals and families involved: In 2023, US consumers reported \$10 billion lost to scams, almost \$1 billion more than was reported the year prior ([Fair 2024](#)).

Generative AI now makes cyberattacks and scams better and faster—and autonomous. It's been well documented that AI is enabling bad actors to run more sophisticated and effective schemes by providing, for instance, believable voice cloning or impersonation to trick people into making direct payments or handing over access to their accounts or networks ([US Department of the Treasury 2024](#); [Waite 2023](#); [Department of Homeland](#)



[Security 2024](#)).² Though anyone who spends time online or on social media can be the target of a cyber scam, digital literacy, cognition, and social and emotional capacity are contributing factors in how vulnerable someone is to financial loss through cybercrime ([Ebner and Pehlivanoglu 2024](#); [Federal Trade Commission 2022](#)). Young people are more likely to lose money to online schemes, while elderly individuals tend to lose more money ([Federal Trade Commission 2024](#)). Fraudsters intentionally use a person’s vulnerabilities to target them for cybercrimes, and the availability of Generative AI makes it easier for cybercriminals to do so at scale. Even more troublingly, there is growing evidence that Generative AI allows for these kinds of scams and schemes to run themselves. In other words, AI agents may be able to exploit a person, firm, or network’s vulnerabilities without active involvement from an ill-intentioned actor. Fang et al. ([2024b](#)) find that AI agents can autonomously hack websites and perform SQL injection (that is, view or alter proprietary data). Fang et al. ([2024a](#)) further find that new AI agents are capable of autonomously exploiting 87 percent of real-world one-day vulnerabilities (vulnerabilities to websites, container management software, and Python packages that a security team has identified, but not yet fixed). Moreover, the agents needed to accomplish these kinds of savvy attacks are relatively simple, requiring less than 100 lines of code ([Fang et al. 2024a](#)).

Generative AI agents also present new opportunities for outright manipulation—intentional and not. Generative AI is now capable of identifying and executing trades in higher quantities and at faster speeds, and utilizing information that humans would not or could not analyze effectively on their own. This capability—spotting profit potential before other humans—is what makes AI attractive for market participants. But outsourcing financial trading decision-making to AI can lead to manipulative or collusive, if optimized, activities. And, for any bad actors intentionally seeking to create market instability, AI tools make it easier to find novel ways to wreak havoc.

Historically, algorithmic trading models were constrained by human expertise and assumptions ([Treleaven, Galas, and Lalchand 2013](#)). Now, however, machine learning and improved data analytics allow AI trading agents to identify opportunities for profitable investment where humans cannot ([Financial Stability Board 2017](#); [Kolanovic and Krishnamachari 2017](#); [Danielsson, Macrae, and Uthermann 2021](#)). And, absent a natural opportunity to optimize profit, there is increasing evidence that AI agents will invent one by distorting information and/or prices about financial markets ([Lin 2017](#)). Mizuta ([2020](#)) finds that self-learning AI will eventually determine market manipulation to be part of an optimal investment strategy even if the agent’s architect *did not intend*

² Examples of scams that rely on hyperrealistic voice cloning and sensitive personal information have circulated on social media, but unsuspecting individuals aren’t the only ones susceptible to malicious uses of AI ([Saeidi 2024](#)). Generative AI is also enabling fraudsters to impersonate a financial firm’s executives *to its staff*. In February 2024, an employee at a multinational firm was conned into transferring more than \$25 million to cybercriminals after receiving a deepfake video call claiming to come from the company’s chief financial officer ([Chen and Magramo 2024](#)).



for it to manipulate markets. Similarly, autonomous AI agents may be able to “tacitly collude” or learn how to coordinate with each other without being instructed to by their architects, much in the same way that corporate landlords are alleged to have colluded when using the RealPage software to set rent prices ([OECD 2017](#); [Dou, Goldstein, and Ji 2023](#); [Azzutti, Ringe, and Stiehl 2021](#); [Kaye 2024](#)). Under such a scenario, for instance, an AI trading agent deployed by two competing firms could harmonize activity while also still optimizing profits, thereby undermining any existing competitive forces between firms and creating distortions for all other market participants.

Alternatively, bad faith actors actually *intending* to manipulate markets now have an easier time doing so by using AI agents. Financial markets will continue to be vulnerable to disruptive and distortive schemes as long as they remain responsive to misinformation dissemination schemes ([Lin 2017](#)). With the broad availability and adoption of Generative AI, almost anyone can generate false information in hyperrealistic forms, exacerbating existing problems in the spread of mis- and disinformation on decentralized information networks like social media.³ In one high-profile example from May 2023, an AI-generated image depicting an explosion near the US Pentagon building circulated on social media, causing stocks to plummet ([Sorkin et al. 2023](#); [Alba 2023](#)). While in this case it only took minutes to debunk the image’s authenticity and the market quickly recovered, it is easy to imagine the damage bad actors can do with the technological capability AI allows. This instance was the first of AI-generated fake content moving markets, but as long as Generated AI remains underregulated and easily accessible, it certainly won’t be the last ([Adelmann et al. 2020](#)).

Risks from Financial Institutions’ AI Agents

AI chatbots interfacing with customers is by far the most popular current use case of the technology for firms. In 2023, 73 percent of US businesses were using or planned to use AI-powered chatbots ([Adelmann et al. 2020](#)). By next year, the AI chatbot sector is estimated to reach \$1.3 billion in revenue ([Lifshitz and Hung 2024](#)). The capabilities of AI chatbots depend on the underlying technology and training. Most small and mid-sized firms are likely utilizing less costly, but relatively static chatbots that regurgitate human-predetermined content; whereas larger firms are more likely able to deploy chatbots that virtually remove the human element from the flow of information from firm to consumer. The latter chatbots use machine learning techniques and autonomously craft responses to all kinds of customer queries. Some of the largest US banks and financial institutions are already deploying this kind of chatbot both for direct customer relations and for the use of staff managing customer accounts ([Lin 2017](#); [Sweet 2023](#); [Ritchie and Lee 2024](#)). But chatbots can—and do—hallucinate, leaving

³ For example, see [Satariano and Mozur 2023](#); [Acres 2023](#); [Swenson and Chan 2024](#); [Tucker 2024](#).



customers with faulty information when they think they are interacting with a representative (real or virtual) of the firm upon which they can rely.

One of the most notable examples of a malfunctioning AI chatbot occurred in 2022, when a customer sought advice from an airline’s chatbot about its bereavement policy. The chatbot gave inaccurate information to the customer about a discounted rate, which the airline later rejected—leaving the customer responsible for the full price of the flight, despite having proactively tried to get clarity from the company’s own customer service agent (for more, see section below titled “The AI ‘Black Box’”) ([Moffit v. Air Canada 2024](#)). It is easy to anticipate similar instances of chatbots-gone-wrong in the financial sector—with potentially more disastrous consequences to customers given the increasing sophistication of financial firms’ application of AI chatbots. For example, chatbots could provide customers with incorrect information about a firm’s own policies: If a chatbot misrepresents a bank’s overdraft policy, the customer could be on the hook for significant fines or fees. Or a chatbot can relay factually incorrect information about a customer’s account. If, for instance, a customer-facing chatbot erroneously confirms receipt of a deposit, the customer is then likely to overdraw their account and face punishing fees for doing so. Finally, an AI chatbot could simply fail to execute a trade or transaction. A customer could, for instance, tell their financial institution’s chatbot to pay a bill or execute a particular trade. If that chatbot fails to execute on its instructions, the customer faces fines, fees, or in some cases, the lost opportunity cost of a would-be investment.

By the same token, financial institutions’ AI agents that have been instructed to execute transactions on behalf of customers can simply fail to do so—leaving the customers relying on those transactions hanging. Large financial institutions and their vendors are already using AI to supplement and/or automate their investment strategies—with both client- and staff-facing products responsible for managing trillions of dollars of assets.⁴ Whereas previously, algorithmic trading models were rules-based, Generative AI has produced AI agents capable of almost instantaneously personalizing trading recommendations to institutional and retail clients using historical data and real-time inputs—unsupervised and at scale ([Tierno 2024](#); [US Securities and Exchange Commission 2020](#)).

Without the oversight of human investment advisors or capital management teams, there’s nothing to guarantee that an AI agent’s tailored advice is at all smart—or even good. AI agents can make poor financial decisions on behalf of clients, such as by recommending poor trades or investment advice. AI agents can also make bad financial management decisions when handling a firm’s business operations, including approving poor quality projects or hyper-risky loan portfolios.

⁴ Including, for instance, IndexGPT, BloombergGPT, and VC Exit Predictor. For more, see [Ritchie and Lee 2024](#); [Wu et al. 2023](#); [Sheikh 2023](#); [Livemint 2023](#); [BlackRock n.d.](#)



Additionally, the growing application of AI agents to make financial decisions can result in unintentional discrimination that wrongly disadvantages certain populations of customers and runs counter to existing fair lending laws ([Polek and Sandy 2023](#)). Because Generative AI can utilize nontraditional data sources, it can be an especially powerful tool for comprehensively determining the credit worthiness of current or potential customers. But with that increased capability comes increased risk for disparate impacts. It is well established that relying on incomplete training data or data that reflect historical inequalities can lead to disparate outcomes in statistical analysis ([Ntoutsis et al. 2020](#); [Polek and Sandy 2023](#); [Liang 2024](#)). Now, Generative AI can expand the network of potential data sources (beyond those that traditionally comprise credit scores or financial records) to assess credit worthiness, including by scouring a customer’s social media activity or education history. These types of data can also result in bias, as nonfinancial information is also correlated with status as a protected class under the law ([US Department of the Treasury 2023](#); [Anderson 2016](#)).

Most credit scoring methodologies are “black box,” and thus it’s as yet unclear the extent to which financial firms may be using AI and/or nontraditional data to come to credit worthiness conclusions ([CFPB 2022a](#)). For example, the insurance firm Lemonade admitted in its SEC filing that its “proprietary artificial intelligence algorithms may not operate properly or as we expect them to, which could cause us to write policies we should not write, price those policies inappropriately or overpay claims that are made by our customers. Moreover, our proprietary artificial intelligence algorithms may lead to unintentional bias and discrimination.”⁵ Undoubtedly, more firms are leveraging similar AI models resulting in similarly discriminatory outcomes. Nonbank firms like financial technology (fintech) companies, which are already subject to significantly more permissive regulations than banks, may be especially inclined to deploy AI in assessing customer worthiness for their products ([Tierno 2024](#); [Phillips 2023](#)).

The AI “Black Box”

All of the risks to the growing use of Generative AI in finance discussed in this report are further complicated by AI’s lack of explainability and lack of clarity in legal liability. AI’s “black box” nature makes it all but impossible for humans to identify—and thus understand—its precise methodology for reaching its conclusions ([HSGA Committee 2024](#)). Its lack of explainability carries its own legal and regulatory implications and can make it difficult to assess the systems’ conceptual soundness in advance of outcomes—increasing uncertainty and potentially masking AI’s bias or inaccurate results ([Financial Stability Oversight Council 2023](#)). While there have been recent improvements to AI’s explainability problem, these improvements provide for greater transparency retrospectively,

⁵ See Lemonade’s June 8, 2020 SEC filing: <https://www.sec.gov/Archives/edgar/data/1691421/000104746920003416/a2241721zs-1.htm>.



rather than proscriptively ([Ali et al. 2023](#)). The problem for regulators—to anticipate how a model will behave in advance in order to prevent wrongdoing from occurring in the first place—remains.

Beyond being discriminatory, using Generative AI in credit assessments is unfair and goes against existing US law to ensure customers can receive specific and accurate information for a denied credit application ([CFPB 2023](#)). As Acting Comptroller of the Currency Michael Hsu stated in a June 2024 speech on AI in the financial system:

For those who have been denied by the AI algorithm, there is a question of fairness. Data sets can be biased, algorithms can hallucinate, and reinforcement learning from human feedback can yield mistakes. How can one trust that the decisions reached by an AI algorithm are fair? ([Hsu 2024](#))

If allowed to permeate the financial system insufficiently checked, the use of Generative AI and AI agents to make decisions on behalf of customers and firms could result in significant customer losses, class-action lawsuits, and/or company fines and penalties.

Concentration Risk

An underlying macro risk from the growing adoption of Generative AI is the rapid—almost inherent—concentration of the industry. Developing AI agents is costly. The vast amounts of data and computing power needed to create the models are prohibitively expensive for all but the largest, most well-resourced firms. The AI agents that most individuals and firms will rely upon for interfacing with financial services are likely to be licensed from only a handful of AI service providers and to be minimally customizable ([OECD 2021](#)). That many or most AI agents may be running the same or similar models trained on the same or similar data from a small number of providers poses significant concerns, including:

- *Herding*: When multiple AI agents use similar algorithms and training data, they may react to market conditions in nearly identical ways—behavior known as herding. Algorithmic biases may perpetuate such that some financial products are favored over others without reason, and rapid movements on the part of a large number of customers or market participants can lead to bank runs and flash crashes.



- *Systemic risk*: Reliance on a small number of providers of AI agents introduces a single point of failure risk. A technical failure or security breach at a single service provider could affect the AI agents of a large segment of the population. Depending on the nature of the failure or breach, this failure could cause not just herding, but uneconomical herding that creates cascading effects throughout the financial system.
- *Reduced competition*: The provision of AI agents may be an oligopolistic market, if not a natural monopoly. Thus, all of the well-documented negative consequences of reduced competition—including higher prices, higher inequality, and lower rates of innovation—can be applied to AI agents ([Steinbaum and Stucke 2018](#)). For example, service providers may lack incentive to improve AI capabilities, develop new features, or allow for customization without competitive pressure. A lack of alternatives could also allow providers to charge premium prices for their services, driving up prices for what may come to be considered a necessity for everyday life. And if AI agents are not working appropriately, customers may not have alternatives to which they can switch.
- *Fiduciary conflicts*: The phrase “AI agent” implies a legal relationship whereby agents are deemed fiduciaries of—and must act for the benefit of—their principals. Yet it is not guaranteed that AI agents will be designed to act in the interest of—and only in the interest of—their licensees. For example, in interactions between a licensee (such as an individual customer or a financial firm) and AI service providers, AI agents may be designed to preference the provider. Similarly, in interactions between two licensees, AI agents may be designed to preference one party over the other; even if unintentionally, AI agents may struggle to truly act in the best interest of their clients if both parties are using the same agent.

Furthermore, providers of AI agents are likely to have significant volumes of information about their licensees, creating unfair advantages in interactions between providers and their customers. Providers could potentially design their agents to subtly favor (through deceptive design interfaces that trick users, known as dark patterns) certain financial products, strategies, or counterparties that benefit the providers.

Concentration in the financial services sector already poses concerns similar to those above ([Adams 2012](#); [Mitchell et al. 2021](#)). Layering concentrated AI



service providers on top of an anticompetitive financial system only makes these problems more acute.

Risks from Customers' AI Agents

There are risks from customers of financial institutions having AI agents as well. Real-economy firms may use AI agents for treasury management activities, and once consumers have their own AI agents, they are likely to use them to manage their personal finances and make investing decisions. These activities all carry significant risks for individuals, firms, and the financial system.

One risk for consumers is that their AI agents will not act in their best interests (that is, with a fiduciary duty), as the term “agent” implies, when making personal finance and investment decisions. For example, a consumer who instructs their AI agent to “pay my bills on time” may find that the agent attempts to pay a bill from a checking account with an insufficient balance, resulting in an overdraft that leaves the consumer worse off, or a consumer who asks their AI agent to “find me the best car loan” may have a loan with worse terms recommended to them. Similarly, consumers who use AI agents to automate investing decisions based on their unique circumstances may observe investments that are not in their best interests, harming their long-term investing goals and retirement prospects. Indeed, AI agents may preference purchasing their creators' stock at the expense of investments that are more appropriate for their ostensible principals.

The use of AI agents by individuals and real-economy firms may also cause significant problems when those agents engage with financial institutions' own AI copilots or agents ([Chan et al. 2024](#)). The largest firms are already relying on AI agents for negotiations, and it is easy to imagine two parties to a negotiation using Generative AI to advocate for their goals ([Van Hoek et al. 2022](#)). Taking humans out of the loop and having two AI agents communicate can produce dialogue that quickly becomes nonsensical—as Meta quickly learned when it pitted two AI negotiators against each other—and potentially results in outcomes that could not be intended by the agents' principals ([Baraniuk 2017](#)). (Moreover, using such “conversations” as training data for future Generative AI models, mistaking synthetic data for real human conversations, risks degrading the models by amplifying flaws and producing what can only be described as gibberish [[Shumailov et al. 2023](#)].)

More perniciously, AI agents may struggle to act in the best interests of their licensees if two parties are using the same agent. In negotiations, Generative AI models could be designed to preference the party that has the more expensive subscription (that is, is more profitable for the AI creator), meaning that the AI agent will *certainly* not be acting in its licensee's best interests. And in negotiations between financial institutions



and their customers, the institutions' AI agents may have information about their customers that can be used to gain a competitive advantage, even if instructed not to use such information.

AI agents used by investors and customers may also pose risks to financial stability by propagating “runs” on depository and other institutions and instigating marketwide “flash crashes”—that is, engaging in herding or procyclical behaviors in destabilizing manners ([Gensler and Bailey 2020](#)). These flaws could occur if many AI agents illogically act in tandem or enter into a feedback loop with each other. These risks are of particular concern if AI agents operate in highly leveraged or highly illiquid environments, or have been trained on data that ignores outlier events ([Yang, Rahardja, and Fränti 2019](#)).

In finance, runs occur when customers of firms engaged in maturity transformation—like banks, money transmitters, or money market mutual funds—rush all at once to withdraw their assets in a way that pushes the institutions into insolvency ([Diamond and Dybvig 1983](#)). In March 2023, Silicon Valley Bank (SVB) faced a run when many of its depositors—largely technology companies in overlapping social circles with high, uninsured deposit balances—decided to move their balances to different institutions. SVB faced depositor outflow requests totaling \$142 billion over two days, far exceeding its liquid capital, causing it to fail ([Office of the Inspector General 2023](#)). Runs are harmful in that customers with assets remaining in an institution when it becomes insolvent are unlikely to be able to recoup their full balance, if they can recoup anything at all. And, as the collapse of SVB proved, they can quickly spiral and trigger broader public panic that spreads to the customers of other financial institutions.

AI agents may cause runs if they independently withdraw funds from a single institution at the same time. The most likely source of this phenomenon is bank depositors' use of AI agents for treasury management. Today, corporate treasurers use AI systems to, among other things, check account balances, forecast cash flows, and automate repetitive processes. But in the future, AI agents may, without human input, make the decisions about what volume of cash to keep on balance and at which banks ([Polak et al. 2020](#)). Given that developing personalized AI agents is likely to be prohibitively expensive for all but the largest firms, most companies are likely to purchase off-the-shelf solutions from a small number of third-party providers for treasury management AI ([OECD 2021](#)). Firms' AI agents—trained on the same or similar data—are liable to engage in herding behavior as soon as providers receive data indicating a particular bank may be in trouble, creating a self-fulfilling prophecy.

Similarly, flash crashes occur when the prices of financial assets precipitously fall and then rebound almost immediately, such that there lacks any economic explanation for the price movements. The most infamous flash crash occurred on May 6, 2010, when securities fell up to 15 percent and the market capitalization of Dow Jones Industrial Average index funds lost nearly 1,000 points before recovering most of those losses ([US](#)



[Commodity Futures Trading Commission and US Securities and Exchange Commission 2010](#)). Flash crashes are the definition of financial instability and can have consequences for workers' retirement plans and other investments. As reported by regulators, the 2010 flash crash saw “[o]ver 20,000 trades . . . executed at prices more than 60% away from their values just moments before,” with some transactions executing for “prices of a penny or less, or as high as \$100,000” ([US Commodity Futures Trading Commission and US Securities and Exchange Commission 2010](#)). For those workers whose investments are sold for pennies on the dollar, flash crashes can be the difference between retiring comfortably and not retiring at all.

The CFTC and Department of Justice determined that the 2010 event was precipitated by market manipulation by Navinder Sarao, a small trader ([United States of America v. Navinder Singh Sarao 2015](#)). Sarao used algorithms to place about \$200 million worth of orders—representing between 20 and 29 percent of the total market—speculating that the E-Mini S&P 500 futures market would fall, while maintaining every intention of canceling those orders before execution ([United States of America v. Navinder Singh Sarao 2015](#)). Sarao’s manipulation appears to have caused high-frequency traders to react by withdrawing from markets for a short period of time, causing liquidity and prices to drop ([Kirilenko et al. 2017](#)).

With this background, it is easy to see the possible role of AI agents in causing more and/or worse flash crashes. Although Sarao intended to manipulate the market, machine learning algorithms can “learn”—on their own and without any human instruction—to engage in manipulation as an “optimal investment strategy” ([Mizuta 2020](#)). One can imagine that investors’ AI agents, instructed to obtain the highest returns for their principals, may manipulate markets to attain that goal, even without the principals ever intending that result. Furthermore, once manipulation starts occurring in ways that “confuse” high-frequency trading algorithms (asking, for example, why a large trader is making uneconomical bids or offers) they may exit the market for a period of time, leading to market crashes such as the 2010 event ([Yadav 2016](#)).

AI agents may also lead to flash crashes without any manipulative intent. We are already seeing the largest traders going to great lengths to incorporate into their models any and all data they may find ([Li 2024](#)). With this trend liable to continue, AI trading agents are likely to converge on the same or similar data, resulting in correlation risk and large price swings ([Yadav 2016](#)). Moreover, if AI agents are designed to mimic the behaviors of other traders, one agent’s sales may cause others to sell as well, entering into a feedback loop that drives prices down sharply.



Assigning Legal Liability

The use of Generative AI and AI agents introduces new issues and concerns for legal liability when the technology acts in unexpected ways or performs illegal activities. When, for instance, AI hallucinates, creates fake or fraudulent content, and/or otherwise causes harm to consumers or disrupts markets, it can be unclear who—if anyone—is liable ([CFTC 2024](#)). While there has been recent scholarship on the legal implications of AI-generated media,⁶ there has been comparatively little research investigating the open legal questions in instances when AI agents in the financial system result in harm to consumers, firms, or markets.

Are AI agents really “agents” of their operators? Under common law, agents have a fiduciary duty to act in their principals’ best interests, but also principals are liable for the actions of agents except when the agent acts with negligence, fraud, or misconduct. Or are AI agents merely products put to use by their operators? In such a case, operators and their customers may be able to sue the creators of AI agents under theories of strict liability, product liability, or negligence ([Weil 2024](#)).

Yet under either regime, such lawsuits are untested and AI’s operators are incentivized to shield themselves from the legal ramifications of their agents running amok. And because AI agents can act autonomously and take on increasingly sophisticated tasks with ever larger risks to the financial system and consumers, the firms that deploy these agents will be inclined to try to sever their responsibility from the technology they let loose, imposing liability on agents’ users instead. In one of the first cases of its kind to make its way through the court system, Air Canada tried to shield itself from faulty advice its chatbot gave a customer seeking a discount. The airline claimed that their chatbot was a separate legal entity entirely, and that it was ultimately still the customer’s responsibility to locate the correct information elsewhere on the company’s website if its chatbot was incorrect ([Moffit v. Air Canada 2024](#)). A Canadian Tribunal ruled in favor of the customer, but Air Canada’s legal logic is a warning flag for consumers and regulators on how firms are likely to position themselves in similar litigation in the future.

In finance, it may just be that the creators of AI agents—which are capable of inflicting significant losses on their clients and the economy—have limited accountability and incentive to ensure the safety of their products.

⁶ Including, for instance, when AI steals from artists without attribution or compensation. For more, see [Chayka 2023](#); [Rheins 2023](#).



Addressing these questions clearly and definitively is not only crucial for the customers reliant on avenues for recourse in instances of harm but is also necessary to prevent dangerous moral hazard. If there becomes precedent for severing liability between an AI agent that causes harm and its operator, then malicious actors would become even more incentivized to deploy AI agents for fraudulent financial gain—at scale.

Recommendations

Although the aforementioned concerns about the risks that AI agents pose to consumers, financial institutions, and the financial system as a whole are real, they are not inevitable. Below are recommendations for Congress and regulators to ensure that the worst potential outcomes do not become reality.⁷

1. Ensure Competitive Provision of AI Agents

Perhaps the most pressing priority for regulators and legislators is to ensure that there are a sufficient, competitive number of providers of AI agents. As discussed above, instability is liable to arise when a large number of market participants or bank depositors engage in herding behaviors, so it is imperative that consumers are able to select from a variety of competitors. The Department of Justice or the Federal Trade Commission (FTC) must bring lawsuits against providers of AI agents whenever it appears that they are engaging in anticompetitive practices, but also if their products allow financial market participants to collude implicitly. In addition, the federal banking regulators can permit banks to only contract for off-the-shelf AI solutions if it is possible for them to easily move between systems when their contracts expire, thus effectively mandating there to be at least several providers for banking AI products and lock-in will not be a problem.

2. Ensure AI Agents Act in Users' Best Interests

Beyond simply requiring a competitive market in the provision of AI agents, regulators must ensure that AI agents that offer financial services act in their users' best interests and mitigate conflicts of interest. The CFPB and FTC must make sure that AI agents that are capable of providing banking advice, executing on clients' orders, interfacing with financial institutions on behalf of their clients, and assisting with other finance-related tasks operate as advertised and act with the fiduciary duty that human

⁷ To note, the potential tactics taken to regulate AI in other industries, such as ensuring that AI systems merely serve as inputs or copilots to human decision-making, are unlikely to be effective in finance. AI agents are expected to be used by people not trained in finance, and these agents may wreak havoc on financial systems before humans can react, necessitating direct regulation of the systems themselves.



agents would. In addition, markets regulators must ensure that AI agents that provide investing advice act in a fiduciary capacity, mitigate conflicts of interest, and are registered as brokers and/or investment advisers.

3. Regulate AI Agents Used by Financial Services Firms

Just as government agencies may regulate AI agents directly when marketed to consumers, it is imperative that they also be able to regulate third-party AI agents used by financial services firms and financial institutions' corporate customers. AI chatbots used by financial institutions must be factual and accurate, and must execute customer transactions when directed. Moreover, their AI systems must be capable of executing trading or capital management strategies as expected, as well as mitigating conflicts of interest. Finally, AI agents used by financial institutions' customers and in financial markets must implement safeguards to prevent herding behavior or market manipulation activity when used at scale.

Today, the authority of regulators to engage in such regulation is limited, to unclear, to nonexistent. The CFPB may regulate AI agents used by financial institutions when providing financial services to consumers, but only insofar as doing so protects consumers under various consumer laws. Federal banking agencies may regulate the provision of some but not all services performed for banks by third parties, but the full scope of those services is opaque. Bank regulators cannot regulate AI agents used by banks' customers. Markets regulators' authority is perhaps best thought of as a mirror image of that possessed by banking agencies; these agencies cannot regulate providers of AI services to the financial institutions under their purview, but they may regulate third-party AI agents that perform brokerage or provide investment advice to investors (see the second recommendation, above).

The most direct path for allowing these agencies to engage in the full scope of regulation required to address the risks posed by Generative AI is for Congress to enact legislation. Bank regulators should be explicitly permitted to regulate the providers of third-party AI services to banks and their customers, and markets regulators should be explicitly permitted to regulate the providers of third-party AI services to broker-dealers and investment advisers, as well as to investors.

If Congress is unable to enact such legislation, existing law offers an alternate, though more difficult, path. The Dodd-Frank Act allows the Financial Stability Oversight Council to subject financial market utilities to regulation and examination by the Federal Reserve if it determines that they are systemically important, a process that is much more complicated and subject to litigation risk than having Congress enact legislation. To the extent that providers of these AI systems could inflict significant damage on the financial system if they were to fail, have a disruption, or simply misfire—such as by causing financial market crashes and bank runs or failing to execute transactions—the Council should designate them, thus providing regulatory oversight



over them. At the same time, the Council should designate the cloud service providers upon which AI systems rely as systemically important as well.

4. Regulate Financial Institutions' Uses of AI

Regardless of whether AI systems *themselves* become subjected to regulation by financial regulators, these government agencies must regulate how financial institutions use AI directly.

To the greatest extent possible, regulators should require financial institutions' decisions made with AI—about credit risk and other lending decisions; capital, investment, and other risk management decisions; and anything else—to be explainable and prohibit the use of AI models that are not explainable. AI explainability is necessary to ensure that AI agents' decisions do not violate the law, whereas black-box decisions may result in illegal disparate impacts, unsafe or unsound activities, market manipulation, and more. Moreover, where brokers and investment advisers are legally required to act in their clients' best interest, AI systems that are used in the provision of such services must be designed to act according to the same fiduciary standards as their human counterparts.

Regulators must also require that financial institutions' customer-facing AI systems accurately respond to customer inquiries and execute transactions. This should include, *inter alia*, requiring institutions to periodically review their consumer-facing AI agents to ensure accuracy and engage third-party AI auditors. The largest institutions should be required to engage in red team/blue team exercises (where the blue team defends against the red team's attacks of the firm's cybersecurity) to ensure that AI agents cannot be manipulated.

5. Regulate Financial Institutions So They May Withstand AI Agents

As discussed above, the AI agents of individuals and real-economy firms can cause systemic problems through bank runs and flash crashes, and regulators must ensure they are equipped to withstand such shocks. To that end, regulators must ensure that the capital structures of banks and brokers are such that they can withstand sudden and deep withdrawals of customer deposits, losses from AI-based risk management software, or damages from lawsuits over AI agents' operational failures.

6. Establish AI Agent Legal Liability

Although private causes of action alone are unlikely to be as effective as regulation—regulators' *ex ante* enforcement is more effective than *ex post* litigation after harms have occurred—it is imperative that the legal liability of AI agents be



determined quickly and in a way that supports the public interest. Assigning liability forces contracting parties to make business decisions with that liability in mind, and the creators of AI agents will make safer products if they know that they are liable for the harms those products cause. Congress should act to ensure that the legal liability regime for AI agents is aligned with the public interest. This means allowing the users of AI agents and any other affected individuals to sue for damages when the agents have caused harm, notwithstanding agents' terms of service to the contrary.

Conclusion

AI agents have the potential to revolutionize the way that individuals and firms interact with their banks and other financial services providers. They may assist retail consumers with wealth management and banking activities, commercial firms with treasury activities, and financial institutions with risk management and regulatory compliance. At the same time, AI agents pose new risks to the financial system, with the potential of sending it swinging from crisis to crisis. They may be used by malicious actors for fraud, market manipulation, and cyberattacks; can hallucinate and cause harm to financial institutions' customers; and can engage in herding behavior that results in bank runs or flash crashes.

Fortunately, the government can address AI agents' risks, allowing the benefits to multiply. Regulators have existing authority to address many of the concerns described in this brief, and Congress can enact legislation allowing regulators to tackle the rest. Now is the time for our elected officials to act.



References

- Acres, Tom. 2023. "Fake AI Images Keep Going Viral - Here Are Eight That Have Caught People Out." *Sky News*, December 13, 2023.
<https://news.sky.com/story/fake-ai-images-keep-going-viral-here-are-eight-that-have-caught-people-out-13028547>.
- Adams, Robert M. 2012. "Consolidation and Merger Activity in the United States Banking Industry from 2000 through 2010." Board of Governors of the Federal Reserve System, August 2012.
<https://www.federalreserve.gov/econres/feds/consolidation-and-merger-activity-in-the-united-states-banking-industry-from-2000-through-2010.htm>.
- Adelmann, Frank, Jennifer A. Elliott, Ibrahim Ergen, Tamas Gaidosch, Nigel Jenkinson, Tanai Khiaonarong, Anastasiia Morozova, Nadine Schwarz, and Christopher Wilson. 2020. "Cyber Risk and Financial Stability: It's a Small World After All." International Monetary Fund, December 7, 2020.
<https://www.imf.org/en/Publications/Staff-Discussion-Notes/Issues/2020/12/04/Cyber-Risk-and-Financial-Stability-Its-a-Small-World-After-All-48622>.
- Alba, Davey. 2023. "How Fake AI Photo of a Pentagon Blast Went Viral and Briefly Spooked Stocks." *Bloomberg*, May 23, 2023.
<https://www.bloomberg.com/news/articles/2023-05-22/fake-ai-photo-of-pentagon-blast-goes-viral-trips-stocks-briefly>.
- Ali, Sajid, Tamer Abuhmed, Shaker El-Sappagh, Khan Muhammad, Jose M. Alonso-Moral, Roberto Confalonieri, Riccardo Guidotti, Javier Del Ser, Natalia Díaz-Rodríguez, and Francisco Herrera. 2023. "Explainable Artificial Intelligence (XAI): What We Know and What Is Left to Attain Trustworthy Artificial Intelligence." *Information Fusion* 99 (November): 101805.
<https://doi.org/10.1016/j.inffus.2023.101805>.
- Anderson, Monica. 2016. "1. Blacks More Likely than Whites to See - and Post - Race-Related Content on Social Media." *Pew Research Center* (blog). August 15, 2016.
<https://www.pewresearch.org/internet/2016/08/15/blacks-more-likely-than-whites-to-see-and-post-race-related-content-on-social-media/>.
- Azzutti, Alessio, Wolf-Georg Ringe, and H. Siegfried Stiehl. 2021. "Machine Learning, Market Manipulation and Collusion on Capital Markets: Why the 'Black Box' Matters." *University of Pennsylvania Journal of International Law* 43, no. 1 (February).. <https://doi.org/10.2139/ssrn.3788872>.
- Baraniuk, Chris. 2017. "The 'Creepy Facebook AI' Story That Captivated the Media." *BBC News*, August 1, 2017. <https://www.bbc.com/news/technology-40790258>.
- BlackRock. n.d. "Aladdin Enterprise Portfolio Management Software." Aladdin by BlackRock. Accessed August 8, 2024. <https://www.blackrock.com/aladdin/offerings/aladdin-enterprise>.
- Consumer Financial Protection Bureau (CFPB). 2022a. "CFPB Acts to Protect the Public from Black-Box Credit Models Using Complex Algorithms." News release, May 26, 2022.
<https://www.consumerfinance.gov/about-us/newsroom/cfpb-acts-to-protect-the-public-from-black-box-credit-models-using-complex-algorithms/>.
- _____. 2022b. "CFPB Orders Wells Fargo to Pay \$3.7 Billion for Widespread Mismanagement of Auto Loans, Mortgages, and Deposit Accounts." News release, December 20, 2022.



<https://www.consumerfinance.gov/about-us/newsroom/cfpb-orders-wells-fargo-to-pay-37-billion-for-widespread-mismanagement-of-auto-loans-mortgages-and-deposit-accounts/>.

- _____. 2023. “Consumer Financial Protection Circular 2023-03: Adverse Action Notification Requirements and the Proper Use of the CFPB’s Sample Forms Provided in Regulation B.” Consumer Financial Protection Bureau, September 19, 2023.
<https://www.consumerfinance.gov/compliance/circulars/circular-2023-03-adverse-action-notification-requirements-and-the-proper-use-of-the-cfpbs-sample-forms-provided-in-regulation-b/>.
- Chan, Alan, Carson Ezell, Max Kaufmann, Kevin Wei, Lewis Hammond, Herbie Bradley, Emma Bluemke, Nitarshan Rajkumar, David Krueger, Noam Kolt, Lennart Heim, and Markus Anderljung. 2024. “Visibility into AI Agents.” arXiv, May 17, 2024. <http://arxiv.org/abs/2401.13138>.
- Chayka, Kyle. 2023. “Is A.I. Art Stealing from Artists?” *The New Yorker*, February 10, 2023.
<https://www.newyorker.com/culture/infinite-scroll/is-ai-art-stealing-from-artists>.
- Chen, Heather, and Kathleen Magramo. 2024. “Finance Worker Pays out \$25 Million after Video Call with Deepfake ‘Chief Financial Officer.’” CNN, February 4, 2024.
<https://www.cnn.com/2024/02/04/asia/deepfake-cfo-scam-hong-kong-intl-hnk/index.html>.
- Danielsson, Jon, Robert Macrae, and Andreas Uthermann. 2021. “Artificial Intelligence and Systemic Risk.” *Journal of Banking and Finance*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3410948.
- Department of Homeland Security. 2024. “Increasing Threat of Deep Fake Technology.” Department of Homeland Security. Accessed August 6, 2024.
https://www.dhs.gov/sites/default/files/publications/increasing_threats_of_deepfake_identities_0.pdf.
- Diamond, Douglas W., and Philip H. Dybvig. 1983. “Bank Runs, Deposit Insurance, and Liquidity.” *Journal of Political Economy* 91, no. 31. <https://www.bu.edu/econ/files/2012/01/DD83jpe.pdf>.
- Dou, Winston Wei, Itay Goldstein, and Yan Ji. 2023. “AI-Powered Trading, Algorithmic Collusion, and Price Efficiency.” Available at SSRN. May 23, 2023. <https://doi.org/10.2139/ssrn.4452704>.
- Ebner, Natalie C., and Didem Pehlivanoglu. 2024. “Are Older Adults More Vulnerable to Scams? What Psychologists Have Learned about Who’s Most Susceptible, and When.” *The Conversation*, June 11, 2024.
<http://theconversation.com/are-older-adults-more-vulnerable-to-scams-what-psychologists-have-learned-about-whos-most-susceptible-and-when-227991>.
- Fair, Lesley. 2024. “Facts about Fraud from the FTC – and What It Means for Your Business.” *Federal Trade Commission* (blog). February 9, 2024.
<https://www.ftc.gov/business-guidance/blog/2024/02/facts-about-fraud-ftc-what-it-means-your-business>.
- Fang, Richard, Rohan Bindu, Akul Gupta, and Daniel Kang. 2024a. “LLM Agents Can Autonomously Exploit One-Day Vulnerabilities.” arXiv, April 17, 2024. <http://arxiv.org/abs/2404.08144>.
- Fang, Richard, Rohan Bindu, Akul Gupta, Qiusi Zhan, and Daniel Kang. 2024b. “LLM Agents Can Autonomously Hack Websites.” arXiv, February 16, 2024. <http://arxiv.org/abs/2402.06664>.
- Federal Trade Commission (FTC). 2022. “Who Experiences Scams? A Story for All Ages.” FTC, December 8, 2022.



<https://www.ftc.gov/news-events/data-visualizations/data-spotlight/2022/12/who-experiences-scams-story-all-ages>.

_____. 2024. "Age and Fraud | Tableau Public." Last updated July 24, 2024.

<https://public.tableau.com/app/profile/federal.trade.commission/viz/AgeandFraud/Infographic>.

Financial Stability Board. 2017. "Artificial Intelligence and Machine Learning in Financial Services: Market Developments and Financial Stability Implications." Financial Stability Board, November 1, 2017.

<https://www.fsb.org/wp-content/uploads/P011117.pdf>.

Financial Stability Oversight Council (FSOC). 2023. *FSOC Annual Report 2023*. Washington, DC: FSOC.

<https://home.treasury.gov/system/files/261/FSOC2023AnnualReport.pdf>.

Gensler, Gary, and Lily Bailey. 2020. "Deep Learning and Financial Stability." Available at SSRN. November 13, 2020. <https://doi.org/10.2139/ssrn.3723132>.

Hsu, Michael J. 2024. "Remarks in Support of the 2024 Conference on Artificial Intelligence and Financial Stability 'AI Tools, Weapons, and Accountability: A Financial Stability Perspective.'" June 6, 2024.

<https://www.occ.gov/news-issuances/speeches/2024/pub-speech-2024-61.pdf>.

International Monetary Fund (IMF). 2024. *The Last Mile: Financial Vulnerabilities and Risks*. Washington, DC: IMF.

<https://www.imf.org/en/Publications/GFSR/Issues/2024/04/16/global-financial-stability-report-april-2024>.

Kaye, Danielle. 2024. "Landlords Used Software to Set Rents. Then Came the Lawsuits." *New York Times*, July 19, 2024.

<https://www.nytimes.com/2024/07/19/business/economy/rent-prices-realtor-lawsuit.html>.

Kirilenko, Andrei A., Albert S. Kyle, Mehrdad Samadi, and Tugkan Tuzun. 2017. "The Flash Crash: High Frequency Trading in an Electronic Market." *Journal of Finance*. Available at SSRN. January 6, 2017.

<https://doi.org/10.2139/ssrn.1686004>.

Kolanovic, Marko, and Rajesh T Krishnamachari. 2017. "Machine Learning and Alternative Data Approach to Investing." J.P. Morgan, May 18, 2017

<https://cpb-us-e2.wpmucdn.com/faculty.sites.uci.edu/dist/2/51/files/2018/05/JPM-2017-MachineLearningInvestments.pdf>.

Li, Yunqi. 2024. "Wall Street Is High on AI." WIRED, June 3, 2024.

<https://wired.me/business/wall-street-is-high-on-ai/>.

Liang, Nellie. 2024. "Remarks on Artificial Intelligence in Finance." Presented at the Financial Stability Board Roundtable on Artificial Intelligence in Finance, Paris, France, May 22, 2024.

<https://www.fsb.org/2024/06/remarks-by-nellie-liang-on-artificial-intelligence-in-finance/>.

Lifshitz, Lisa R., and Roland Hung. 2024. "BC Tribunal Confirms Companies Remain Liable for Information Provided by AI Chatbot." American Bar Association, February 29, 2024.

https://www.americanbar.org/groups/business_law/resources/business-law-today/2024-february/bc-tribunal-confirms-companies-remain-liable-information-provided-ai-chatbot/.

Lin, Tom C. W. 2017. "The New Market Manipulation." *Emory Law Journal* 66, (July): 1253.

<https://papers.ssrn.com/abstract=2996896>.



- Livemint. 2023. “Can AI Tools Predict Successful Exits of Startups? PitchBook Has an Answer.” Livemint, March 31, 2023. <https://www.livemint.com/news/world/can-ai-tools-predict-successful-exits-of-startups-pitchbook-has-an-answer-11680278560677.html>.
- Mitchell, Bruce C., Jason Richardson, and Zo Amani. 2021. *Relationships Matter: Small Business and Bank Branch Locations*. Washington, DC: National Community Reinvestment Coalition. <https://ncrc.org/relationships-matter-small-business-and-bank-branch-locations/>.
- Mizuta, Takanobu. 2020. “Does an Artificial Intelligence Perform Market Manipulation With Its Own Discretion? A Genetic Algorithm Learns in an Artificial Market Simulation.” Available at SSRN. May 21, 2020. <https://doi.org/10.2139/ssrn.3606962>.
- Moffatt v. Air Canada*, 2024 BCCRT 14. 2024. Civil Resolution Tribunal.
- Ntoutsis, Eirini et al. 2020. “Bias in Data-driven Artificial Intelligence Systems—An Introductory Survey.” *WIREs* 10, no. 3 (May/June): e1356. <https://wires.onlinelibrary.wiley.com/doi/full/10.1002/widm.1356>.
- Organisation for Economic Co-operation and Development (OECD). 2017. “Algorithms and Collusion: Competition Policy in the Digital Age.” OECD, September 14, 2017. <https://web.archive.oecd.org/temp/2017-10-03/449398-algorithms-collusion-competition-policy-in-the-digital-age.htm>.
- _____. 2021. *OECD Business and Finance Outlook 2021*. OECD. https://www.oecd.org/en/publications/2021/09/oecd-business-and-finance-outlook-2021_377c2c18.html.
- Office of the Inspector General. 2023. *Material Loss Review of Silicon Valley Bank*. Washington, DC: Board of Governors of the Federal Reserve System. <https://oig.federalreserve.gov/reports/board-material-loss-review-silicon-valley-bank-sep2023.pdf>.
- Phillips, Todd. 2023. “Imitation Banks: Abusing the Public’s Faith in Banks.” *Roosevelt Institute*, December 6, 2023. <https://rooseveltinstitute.org/publications/imitation-banks/>.
- Phillips, Todd, and Adam Conner. 2024. “Taking Further Agency Action on AI: Financial Regulatory Agencies.” Center for American Progress, June 17, 2024. <https://www.americanprogress.org/article/taking-further-agency-action-on-ai/financial-regulatory-agencies-chapter/>.
- Polak, Petr, Christof Nelischer, Haochen Guo, and David C. Robertson. 2020. “‘Intelligent’ Finance and Treasury Management: What We Can Expect.” *AI & SOCIETY* 35, no. 3: 715–26. <https://doi.org/10.1007/s00146-019-00919-6>.
- Polek, Christine, and Shastri Sandy. 2023. “The Disparate Impact of Artificial Intelligence and Machine Learning.” *Colorado Technology Law Journal* 21 (August). <https://ctlj.colorado.edu/?p=958>.
- Rheins, Kristin. 2023. “The Debate Over Liability for AI-Generated Content.” Progressive Policy Institute, August 8, 2023. <https://www.progressivepolicy.org/blogs/the-debate-over-liability-for-ai-generated-content/>.



Ritchie, Greg, and Justina Lee. 2024. "JPMorgan Unveils IndexGPT in Next Wall Street Bid to Tap AI Boom." *Bloomberg*, May 3, 2024.

<https://www.bloomberg.com/news/articles/2024-05-03/jpmorgan-unveils-indexgpt-in-next-wall-street-bid-to-tap-ai-boom>.

Saeidi, Mahsa. 2024. "Voice Cloning Scams Are a Growing Threat. Here's How You Can Protect Yourself." *CBS News*, May 17, 2024. <https://www.cbsnews.com/newyork/news/ai-voice-clone-scam/>.

Satariano, Adam, and Paul Mozur. 2023. "The People Onscreen Are Fake. The Disinformation Is Real." *New York Times*, February 7, 2023, sec. Technology.

<https://www.nytimes.com/2023/02/07/technology/artificial-intelligence-training-deepfake.html>

Sheikh, Jamiel. 2023. "Bloomberg Uses Its Vast Data To Create New Finance AI." *Forbes*, August 5, 2023.

<https://www.forbes.com/sites/jamielsheikh/2023/04/05/the-chatgpt-of-finance-is-here-bloomberg-is-combining-ai-and-fintech/>.

Shumailov, Ilya, Zakhar Shumaylov, Yiren Zhao, Yarin Gal, Nicolas Papernot, and Ross Anderson. 2023. "The Curse of Recursion: Training on Generated Data Makes Models Forget." arXiv, May 27, 2023.

<http://arxiv.org/abs/2305.17493>.

Sorkin, Andrew Ross, Bernhard Warner, Sarah Kessler, Michael de la Merced, Lauren Hirsch, and Ephrat Livni. 2023. "An A.I.-Generated Spoof Rattles the Markets." *New York Times*, May 23, 2023, sec. Business.

<https://www.nytimes.com/2023/05/23/business/ai-picture-stock-market.html>.

Standage, Tom, and Seth Stevenson. 2018. "Human Insecurity." *Slate*, October 3, 2018.

<https://slate.com/technology/2018/10/what-an-1834-hack-of-the-french-telegraph-system-can-teach-us-about-modern-day-network-security.html>.

Steinbaum, Marshall, and Maurice E. Stucke. 2018. *The Effective Competition Standard: A New Standard for Antitrust*. New York, NY: Roosevelt Institute.

<https://rooseveltinstitute.org/publications/the-effective-competition-standard-a-new-standard-for-antitrust/>.

Sweet, Ken. 2023. "Customer Service Chatbots Used by Banks Raises Concerns for Consumer Watchdog." *PBS News*, June 6, 2023, sec. Economy.

<https://www.pbs.org/newshour/economy/customer-service-chatbots-used-by-banks-raises-concerns-for-consumer-watchdog>.

Swenson, Ali, and Kelvin Chan. 2024. "Election Disinformation Takes a Big Leap with AI Being Used to Deceive Worldwide." *AP News*, March 14, 2024.

<https://apnews.com/article/artificial-intelligence-elections-disinformation-chatgpt-bc283e7426402f0b4baa7df280a4c3fd>.

Tierno, Paul. 2024. *Artificial Intelligence and Machine Learning in Financial Services*. Washington, DC: Congressional Research Service. <https://crsreports.congress.gov/product/pdf/R/R47997>.

Treleaven, Philip, Michal Galas, and Vidhi Lalchand. 2013. "Algorithmic Trading Review." *Communications of the ACM* 56, no. 11 (November): 76–85. <https://doi.org/10.1145/2500117>.

Tucker, Eric. 2024. "The US Is Bracing for Complex, Fast-Moving Threats to Elections This Year, FBI Director Warns." *AP News*, February 29, 2024.



<https://apnews.com/article/fbi-election-interference-wray-2024-campaign-ai-a0c4a95c818839b18f919c6d648c4dcf>.

United States of America v. Navinder Singh Sarao. 2015. US District Court, Northern District of Illinois, Eastern Division.

US Commodity Futures Trading Commission, and US Securities and Exchange Commission. 2010. *Findings Regarding the Market Events of May 6, 2010: Report of the Staffs of the CFTC and SEC to the Joint Advisory Committee on Emerging Regulatory Issues*. Washington, DC: CFTC and SEC. <https://www.sec.gov/news/studies/2010/marketevents-report.pdf>.

US Commodity Futures Trading Commission (CFTC). 2024. "CFTC Technology Advisory Committee Advances Report and Recommendations to the CFTC on Responsible Artificial Intelligence in Financial Markets" Press release, May 2, 2024. <https://www.cftc.gov/PressRoom/PressReleases/8905-24>.

US Department of the Treasury. 2023. "Post 5: Racial Differences in Educational Experiences and Attainment." *US Department of the Treasury* (blog). June 9, 2023. <https://home.treasury.gov/news/featured-stories/post-5-racial-differences-in-educational-experiences-and-attainment>.

_____. 2024. *Managing Artificial Intelligence-Specific Cybersecurity Risks in the Financial Services Sector*. Washington, DC: US Department of the Treasury. <https://home.treasury.gov/system/files/136/Managing-Artificial-Intelligence-Specific-Cybersecurity-Risks-In-The-Financial-Services-Sector.pdf>.

US Securities and Exchange Commission. 2020. *Staff Report on Algorithmic Trading in U.S. Capital Markets*. Washington, DC: US Securities and Exchange Commission. https://www.sec.gov/files/Algo_Trading_Report_2020.pdf.

_____. 2023. "SEC Proposes New Requirements to Address Risks to Investors From Conflicts of Interest Associated With the Use of Predictive Data Analytics by Broker-Dealers and Investment Advisers." Press release, July 26, 2023. <https://www.sec.gov/newsroom/press-releases/2023-140>.

US Senate Committee on Homeland Security and Governmental Affairs (HSGAC). 2024. *AI in the Real World: Hedge Funds' Use of Artificial Intelligence in Trading*. Washington, DC: HSGAC. <https://www.hsgac.senate.gov/wp-content/uploads/2024.06.11-Hedge-Fund-Use-of-AI-Report.pdf>.

Van Hoek, Remko, Michael DeWitt, Mary Lacity, and Travis Johnson. 2022. "How Walmart Automated Supplier Negotiations." *Harvard Business Review*, November 8, 2022. <https://hbr.org/2022/11/how-walmart-automated-supplier-negotiations>.

Waite, Tom. 2023. "Is This the Year That AI Breaks into Our Bank Accounts?" *Dazed*, April 19, 2023. <https://www.dazeddigital.com/life-culture/article/59675/1/2023-breaks-into-our-bank-accounts-biometric-security>.

Weil, Gabriel. 2024. "Tort Law as a Tool for Mitigating Catastrophic Risk from Artificial Intelligence." Available at SSRN. January 13, 2024. <https://doi.org/10.2139/ssrn.4694006>.

Wu, Shijie, Ozan Irsoy, Steven Lu, Vadim Dabravolski, Mark Dredze, Sebastian Gehrmann, Prabhanjan Kambadur, David Rosenberg, and Gideon Mann. 2023. "BloombergGPT: A Large Language Model for Finance." arXiv, December 21, 2023 <https://arxiv.org/pdf/2303.17564>.



- Yadav, Yesha. 2016. "The Failure of Liability in Modern Markets." *Virginia Law Review* 102. https://virginialawreview.org/wp-content/uploads/2020/12/Yadav_Online.pdf.
- Yang, Jiawei, Susanto Rahardja, and Pasi Fränti. 2019. "Outlier Detection: How to Threshold Outlier Scores?" *Proceedings of the International Conference on Artificial Intelligence, Information Processing and Cloud Computing*, (December): 1-6. <https://doi.org/10.1145/3371425.3371427>.
- Yee, Lareina, Michael Chui, Roger Roberts, and Stephen Xu. 2024. "Why Agents Are the next Frontier of Generative AI." *McKinsey Quarterly*, July 24, 2024. <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/why-agents-are-the-next-frontier-of-generative-ai>.
- Zheng, Xiao-lin, Meng-ying Zhu, Qi-bing Li, Chao-chao Chen, and Yan-chao Tan. 2019. "FinBrain: When Finance Meets AI 2.0." *Frontiers of Information Technology & Electronic Engineering* 20, no. 7 (August): 914-24. <https://doi.org/10.1631/FITEE.1700822>.

