

January 2026

From Innovation to Impact

AI Adoption in U.S.

Financial Services

A Report with Key Findings from
the FinRegLab AI Symposium 2025



FinRegLab
AI SYMPOSIUM
AI 2025

What's In This Paper

- 3 Executive Summary
- 5 Introduction
- 7 **Theme 1:** AI adoption is already reshaping the U.S. economy, even as broader productivity gains and workforce impacts take time to develop.
- 11 **Theme 2:** As the financial sector implements AI, the technology is moving at lightspeed—yet the people, processes, and products are still catching up.
- 18 **Theme 3:** For people and small businesses, trust and personalization are critical to facilitate broader benefits and faster adoption.
- 22 **Theme 4:** Building better systems can help manage the risks and increase the benefits of AI adoption for individual companies, consumers, and the broader economy.
- 27 Conclusion

Acknowledgments

We would like to acknowledge the FinRegLab team who worked on various elements of this report and the FinRegLab AI Symposium 2025: Colin Colter, Melissa Koide, Chandra Lim, Kelly Thompson Cochran, and Elizabeth Vivirito. Thank you to the sponsors of the FinRegLab AI Symposium 2025 that also made this paper possible: Capital One, JPMorganChase, PNC Bank, Flourish Ventures, Skoll Foundation, Zillow Group, Synchrony Financial, and American Bankers Association. The findings and conclusions contained within are those of the authors and do not necessarily reflect the positions or policies of the sponsors.

About FinRegLab

FinRegLab is a D.C.-based research organization working at the intersection of financial innovation and public policy. Our work is grounded in a simple but powerful mission: to use technology and data to help people and small businesses manage their financial lives, weather economic shocks, and build long-term security. Learn more and subscribe here: finreglab.org.



About the FinRegLab AI Symposium 2025

On November 20, 2025, FinRegLab hosted its second annual AI Symposium. More than 300 leading experts and stakeholders from government, industry, and civil society assembled at the Kennedy Center REACH in Washington, D.C. to discuss and grapple with the challenges and opportunities AI brings to financial services. This report summarizes the key insights from the Symposium and is designed to guide leaders across financial services and policymaking as artificial intelligence alters our financial systems. Learn more here: aisymposium.finreglab.org.



Executive Summary

Artificial intelligence's (AI) promise to dramatically transform the financial services landscape is taking place before our eyes. Just a year after the inaugural [FinRegLab AI Symposium](#) highlighted the inflection point facing the sector and the broader U.S. economy, participants in the 2025 Symposium described concrete impacts, evolving risks, and expanding opportunities. Technology advancements that still seemed years away have progressed rapidly, further fueling excitement about AI's ability to advance financial access, consumer wellbeing, and broader economic growth and stability.

But both realizing AI's full benefits and managing its risks require rethinking financial firms' internal processes, building trust and value for customers, and driving interoperability among market actors. The sector's ability to achieve these goals will help to shape the broader impacts of AI investments on the U.S. economy, particularly because of the unique role that financial services play in household and small business economic participation and financial security.

The tools and approaches that the financial services sector develops could also be instructive for broader AI adoption as the United States works to build broad-based consensus about what we as a society want AI to do and not do. To facilitate this process, this report highlights key messages, ideas, and examples shared at FinRegLab's second annual AI Symposium on November 20, 2025.

Four Themes and Ten Key Insights from the AI Symposium

First Theme: AI adoption is already reshaping the U.S. economy, even as broader productivity gains and workforce impacts take time to develop.

1. Record levels of AI investments are substantially bolstering the broader economy but also raise concerns about the financial complexity and sustainability of those investments.
2. The rapid acceleration of agentic AI-enabled commerce could have substantial impacts on consumers, small merchants, and financial institutions.

"AI allows us to meet people exactly where they are with hyper-personalized services. It's not just about financial inclusion, it's about turning that first point of contact into an escalator of economic opportunity."

SHENA ASHLEY
Vice President of
Community Impact
& Investment
Capital One



3. Longer-term impacts on economic growth and employment will be shaped by the extent to which individual companies and households realize benefits from AI adoption.

Second Theme: As the financial sector implements AI, the technology is moving at lightspeed—yet the people, processes, and products are still catching up.

4. Particularly for legacy providers, leveraging AI's full benefits requires rethinking data, infrastructure, and talent, in addition to managing new models.
5. Deploying generative and agentic AI requires new tools, strategies, and guardrails for managing reliability and consistency.
6. As they gain more experience, financial services providers are shifting their sights toward more ambitious reimaging of products and workflows.

Third Theme: For people and small businesses, trust and personalization are critical to facilitate broader benefits and faster AI adoption.

7. Next-generation AI tools that offer personalized, proactive support in managing users' finances are especially promising.
8. Meaningful transparency, consent, and controls are critical to building trust, as customers decide how much autonomy they want to cede to AI systems.

Fourth Theme: Building better systems can help manage the risks and increase the benefits of AI adoption for individual companies, consumers, and the broader economy.

9. Cross-industry standards and processes can have multiplier effects by encouraging consistency, facilitating dispute resolution, and increasing trust among participants.
10. Clear and effective regulatory frameworks can also give stakeholders greater confidence in AI adoption and a more competitive marketplace.

We hope this report serves both as a summary of the major highlights of the Symposium and as a resource for future discussions, ideas, and innovations.

"I do think we will get to a place very soon where you have a smarter angel on your shoulder telling you things [about the risks and benefits of financial products]. The arbitrage of people making money on [consumers'] lack of attentiveness, lack of education, [and] lack of people to lean on [for advice] is going away very, very quickly. That's going to be awesome."

MAX LEVCHIN
Co-Founder and CEO
Affirm



Introduction

AI Investments and Uses Are Accelerating

The accelerating pace of AI adoption has made itself felt on both Wall Street and Main Street in 2025, marked by record investment activity, energy price impacts, and AI integrations into everything from office and accounting software to e-commerce platforms.

The financial services sector has played a critical role in these changes, both through helping to finance other industries' AI initiatives and adopting the technology directly. Financial firms have used certain forms of AI for decades, such as machine learning (ML) models for fraud detection. But attention has increasingly shifted to newer technologies, such as the ability of large language models (LLMs) to provide conversational interfaces in a broad range of settings, and the potential for "agentic AI" systems that combine software agents and other AI models to respond quickly to changing conditions and to automate a broad range of processes.

These newer technologies offer the chance for financial services providers to reimagine their business processes, deliver highly personalized solutions, and develop new ways to reach underserved populations. At the same time, open-source models and other tools have lowered the barrier for new entrants to build bespoke financial products from the ground up. Regardless of who uses them, these tools also raise new sources of uncertainty and concerns about risks—ranging from broader impacts on employment and financial stability to more granular questions about the reliability, complexity, and suitability of particular tools for particular tasks. The answers to these questions will prove especially consequential for low- and moderate-income (LMI) households as well as for small businesses. These groups often struggle to access financial services tailored to their unique needs and have little margin for error in adopting new products.

"The task for us as an industry, whether it's banks or fintechs or venture capital firms or the supervisory community, is to make sure that we internalize these technologies in ways that are sustainable, and we retain focus on our job in the real economy: to enable daily commerce and people's dreams."

BEN HOFFMAN
EVP, Chief Strategy Officer and Head of Consumer Products
Fifth Third Bank



Decisions We Make Today Will Shape Our Society and Economy

The ultimate impacts of AI deployment will depend on how people and companies will use AI to create value. Whether as consumers, employees, business leaders, or financial advisors, the people using AI tools will ultimately shape AI's future.

This report dives into the future of AI in financial services in the United States based on discussions during the FinRegLab AI Symposium 2025. The event featured over 30 speakers across 11 sessions, providing an unparalleled opportunity for dialogue and collaboration at the intersection of financial services, technology, public policy, and social and economic impact. Designed as a day for networking and building partnerships, the event enabled thought leaders to engage each other about their work to realize the benefits of AI while mitigating the risks.

Organization of the Report

We have organized the report into four major themes, each conveying key ideas, quotes from speakers, and examples of opportunities, challenges, and partnerships which were shared at the Symposium.



AI adoption is already reshaping the U.S. economy, even as broader productivity gains and workforce impacts take time to develop.

AI's short-term impacts on the U.S. economy have already accelerated in ways not fully anticipated a year prior. Record levels of investment and a burst of initiatives to adopt agentic AI systems in e-commerce have impacted critical sectors, even as AI increases pressure on energy systems and pricing in ways impacting a broad range of households and businesses.

Speakers emphasized the importance of both AI inputs and outputs over longer horizons. In its

first months in office, the Trump Administration published the [AI action plan](#) and issued a series of executive orders emphasizing infrastructure development, data sources, workforce training, and model-development priorities as a matter of global competition. At the same time, much of AI's long-term economic impacts will depend on the extent to which individual companies and households adopt AI applications that increase their productivity and provide other benefits.

“There isn’t really any place where AI isn’t going to touch us directly. It’s going to impact us in all facets of our lives. How do we respond to this? How do you regulate it? And how do you promote [it]?“

SEN. MIKE ROUNDS

Member, Senate Committee on Banking, Housing, and Urban Affairs



KEY INSIGHT #1

Record levels of AI investments are substantially bolstering the broader economy but also raise concerns about the financial complexity and sustainability of those investments.

Investments in AI have become so large that they are significantly driving gross domestic product (GDP) growth, said **Michael S. Barr, Member of the Board of Governors of the Federal Reserve System**, during morning remarks. This increase in capital wealth has likely sustained broader consumption in the economy through a mechanism economists call “the wealth effect”, since people who feel wealthier tend to spend more.

“The value of all stocks held by Americans is up \$8 or \$9 trillion in one year compared to this time last year,” **Mark Zandi, Chief Economist, Moody’s Analytics**, said. “If we hadn’t had that surprising AI supply shock this year, we might be in a very different place from a macroeconomic perspective.”

However, those impacts are not distributed evenly, since LMI households generally have less exposure to equity markets and are heavily impacted by price increases and other countervailing economic trends. This potentially exacerbates tensions from the K-shaped economic recovery from the COVID-19

pandemic, and longer-term trends toward more extreme differences in wealth distribution.

Other potential sources of concern include the increasing complexity and interconnectedness of financing structures, potential mismatches in investor expectations about the timing of productivity gains, and the speed at which current infrastructure investments become outdated in the face of additional technology developments. Compounding those concerns, these economic impacts will likely manifest much faster than prior technology cycles, even compared to the pace of internet adoption over the course of the 1990s.

“We’re not yet seeing [AI investments] that translate broadly into changes in productivity and the economy, changes in the labor market. It’s really the big, big investments in artificial intelligence and data centers.”

MICHAEL S. BARR

Member, Board of Governors of the Federal Reserve System



KEY INSIGHT #2

The rapid acceleration of agentic AI-enabled commerce could have substantial impacts on consumers, small merchants, and financial institutions.

A second area of rapidly growing impact is the integration with search engines, advertising, payments, and other aspects of e-commerce. Consumers, small merchants, and many banks may not be fully aware of these initiatives, some of which focus on infrastructure builds, industry standards, and limited pilot testing. But the flurry of activity from the providers of general LLMs, e-commerce platforms, search engines, payment networks, and other infrastructure providers could have significant impacts on large portions of the retail economy in coming months and years.

These efforts focus on developing agentic shopping tools that can search for products, monitor pricing changes, execute purchases, and select payment mechanisms on behalf of consumers, potentially with little or no active oversight. For merchants and shopping platforms, this raises questions about how to distinguish “good” from “bad” bot traffic, confirm what legitimate agents are authorized to do on behalf of their users, and market themselves to these agents.

For consumers (as discussed further below in the third theme), it raises questions about how much to trust systems that take action on their behalf. And questions about responsibility for handling disputes and errors if an agent makes mistakes or exceeds its authority remain critical for all ecosystem participants including payments providers.

While there will be transition costs, speakers emphasized opportunities to create new channels to reach customers. In addition to building protocols and other general infrastructure, several larger players have launched initiatives to help merchants prepare for the change. For **Ren Zhang, VP, Global Head of Product AI, PayPal**, the success of agentic commerce depends, in part, on driving success for smaller merchants. “For small businesses the key is about access. Access to those kinds of networks, access to the consumers... [by] making it visible to the customers through personalization, we give them a very good chance.”

KEY INSIGHT #3

Longer-term impacts on economic growth and employment will be shaped by the extent to which individual companies and households realize benefits from AI adoption.

As we look ahead to AI's broader impacts, speakers emphasized that long-term effects on productivity, job market changes, and economics growth will depend on firm implementation and household use of the technology. Changes in organizational culture and behavior have proven so substantial that productivity may decline in the short-term while companies work to translate these investments into concrete improvements.

The technology may also have varying impacts on job market dynamics. AI may completely displace some roles. In other roles, AI could drive productivity gains that enable employees to demand higher wages without triggering inflation, and AI could trigger the emergence of completely new occupations. These dynamics may already have led to slower hiring as employers wait to see how AI implementation plays out. While there is substantial reason for optimism about long-term effects, the speed and intensity of change could also produce different transition scenarios and challenges

going forward. The current levels of uncertainty demand taking all these possible scenarios seriously.

The Symposium focused on questions about how financial services providers, consumers and small business customers, and systems initiatives are adopting the technology, determining its benefits, and managing its risks. Given financial services' connection to other parts of the economy, the ways that stakeholders answer these questions will play a defining role in determining the broader ongoing impacts of AI on the U.S. economy.

"If you ask businesses, 'Is it a bubble?' they might have [different] opinions. But if you say, 'Okay, will you give the AI back?' they'll be like 'No, absolutely not. We're using that AI. It's very important to us now.'"

JOHN COLLISON
Co-Founder and President
Stripe



As the financial sector implements AI, the technology is moving at lightspeed—yet the people, processes, and products are still catching up.

In contrast to the past year's discussions about managing ML models for credit underwriting and early piloting of LLM applications largely in internal settings, participants in this year's Symposium highlighted a broader range of use cases and technologies. They emphasized that AI has begun to provide substantial productivity gains in diverse settings—increasingly including customer-facing contexts—as financial institutions gain more confidence in developing, deploying, and managing new models.

At the same time, there was broad agreement that the underlying technologies have often moved faster than people, processes, and product development in part because of the need to ensure appropriate levels of reliability, consistency, transparency, and consumer

protection in financial applications. Speakers outlined several factors shaping the pace and scale of adoption, including the need for broader infrastructure and process changes to fully leverage the benefits of AI, particularly among legacy services providers.

Another important factor is the need to develop tools, techniques, and guardrails when relying on generative and agentic AI models, given fundamental differences in their operation and complexity from other technologies. Nevertheless, while many early AI initiatives focused primarily on efficiency gains, some institutions have now started to reimagine processes and products in more ambitious ways.

Generative and Agentic AI Use Cases Broaden Across Financial Services

Since last year's Symposium, participants discussed the increasingly diverse contexts in which they are deploying AI tools:

Affirm is using generative AI applications to develop credit underwriting metrics that involve time series data and produce adverse action notices that use more natural language which can be tuned to match customers' tone.

Capital One has launched a consumer-facing agentic application called Chat Concierge that helps enhance the car buying experience for both auto dealers and consumers.

JPMorganChase has launched customer-facing AI applications for travel services and used AI models to facilitate proactive client outreach during recent market volatility.

SardineAI is providing banks and other financial institutions with agentic AI systems that automate Know Your Customer screening, transaction monitoring reviews, and fraud detection including detecting AI-generated deep fakes.

Stripe has developed its Payments Foundation Model to identify fraud patterns in more open-ended and sophisticated ways than traditional approaches.

“When I look at this sort of pace of change, it took 23 years since the web protocol was invented to achieve 90% internet adoption. It took nine months for that to happen in the LLM world.”

LAKHBIR LAMBA
President & CEO
Regional Management Corp.



KEY INSIGHT #4

Particularly for legacy providers, leveraging AI's full benefits requires rethinking data, infrastructure, and talent, in addition to managing new models.

While developing, validating, and deploying new types of AI models can prove a substantial technical and compliance challenge in its own right, financial services providers emphasized that these activities are only one aspect of what they must orchestrate to fully leverage the benefits of AI adoption. Change management starts with decisions about data sources, infrastructure, and talent.

Across all types of providers and use cases, speakers emphasized the importance of representative, well-curated, accessible data as the foundation for model building, training, and reinforcement. While LLMs can be useful in working with more unstructured sources, nevertheless, data provenance, supplemental training, and privacy safeguards are still paramount, especially when working with proprietary models from outside vendors. Data issues can be significantly challenging for companies that have evolved over decades, relying on multiple data silos, legacy systems,

and complex business processes.

Financial services providers also face fundamental strategic decisions regarding the nature of their technology infrastructure and whether and when to build their own models versus relying on vendors. To solve this problem, large banks are more likely to invest in building their own tech stacks to allow greater control over configurations and customization. Instead of committing to a single model or guardrail, they can swap models, tighten input and output guardrails, and add cyber protections as needed.

Another approach is to partner with vendors, which can have expertise and infrastructure that are difficult for smaller providers to match. However, relying on vendors also comes with challenges and tradeoffs, such as the degree of customization allowed and the ability to meet regulators' risk management expectations. In addition, large numbers of financial institutions relying on a small number of providers can create concentration risks.

KEY INSIGHT #4

In adjusting talent strategies, financial services providers need to compete for top-notch coding talent and to think through ways to bring their whole workforces along in understanding how different technology tools can make them more efficient. “A lot of it is that cultural shift of ‘Oh, I’m not sure what this is. Do I really need to learn this?’ to making sure our employees feel empowered and motivated to do so, understand why this is so important... and also bring their own ideas of the way AI can improve their own professional lives... [and] get to more meaningful work,” said **Kelly Shaw, Head of Synchrony Ventures.**

Though banks have tended to be more cautious in their initial AI deployments than fintechs and general technology companies, they are accelerating over time. **Prem Natarajan, EVP, Chief Scientist and Head of Enterprise AI, Capital One**, described going from a “build on Friday, deploy on Monday, and often roll back Wednesday” which he has seen in big tech to the bank environment. “What I have found is that when [we] integrate over a period of a year or a year and a half, we are actually able to move at a speed that I would say is at least comparable, if not faster, than the tech industry,” he said. “You just need to have a certain relentlessness while being well managed.”



KEY INSIGHT #5

Deploying generative and agentic AI across more settings requires new tools, strategies, and guardrails for managing reliability and consistency.

Addressing concerns about the reliability, consistency, and transparency of LLMs and agentic AI systems in specific products remains a fundamental challenge for both business and compliance reasons. While initial use cases often relied on co-pilot structures and “humans in the loop” for quality control, financial services providers are now adopting a greater mix of strategies and safeguards as they deploy these more complex systems in a growing range of use cases.

The challenge with these technologies is in part due to differences in their structure and

complexity. The kinds of machine learning models that have been used historically to predict fraud or credit risk are sometimes called “deterministic” because they produce consistent outputs in response to the same inputs, which are structured datasets such as credit report information. In contrast, LLMs are often trained on large portions of the internet. To generate responses to users’ queries, they sample from probability distributions and are often structured to include a certain degree of randomness. This makes them more flexible and “conversational” than prior chatbot technologies but increases the risk that the models will generate

“ *It’s ultimately about structure and modularity and control and how these things come together to let you understand what your system can and cannot do. There are multiple ways to get at it. But you cannot just unleash this general thing that is just as happy to write you a haiku about some medical system as it is to check your [account] balance.”*

DAN KLEIN

Professor of AI at UC Berkeley &
CO-Founder/CTO, Scaled Cognition



KEY INSIGHT #5

inconsistent and even incorrect responses (often called hallucinations).

The models' scale and complexity also create challenges relative to more traditional ML models, although technological advances and practical experience are helping financial service providers expand their toolkits for managing the models.

"Generative AI can do a lot of things. Then you start asking it more and more complex things, and suddenly you're less impressed," said **Prem Natarajan**. "[But] when you combine specialization with the power of reasoning, you can now suddenly take that complex thing that the plain old vanilla gen[erative] AI [couldn't handle]... and break it down into simpler things that the specialized agents are then able to execute and take action on."

Some companies use one AI model to check or monitor the operation of another AI model, though they emphasize the importance of careful structuring and deployment.

Soups Ranjan, CEO & Co-Founder, SardineAI, cautioned that "when we are building AI agents on top of those other LLMs, we have realized that you have to give them as narrow a scope as possible to prevent hallucinations and ensure repeatability." Other speakers also cautioned about this approach, noting that two non-deterministic models added together do not make a deterministic model, but emphasizing that there are a broad range of other tools and strategies.

For example, providers are making targeted decisions about when and how they use general LLMs versus when to use specialized and small language models or ML algorithms. While speakers continued to emphasize the importance of humans in the loop for final decision making or monitoring, they also highlighted strategies about LLMs' architecture, supplemental training with more curated data, and use of automated monitoring and alert systems to detect changes in performance once the models have been deployed.

KEY INSIGHT #6

As they gain more experience, financial services providers are shifting their sights toward more ambitious reimagining of products and workflows.

As financial service providers upgrade their infrastructure and gain more confidence in managing AI applications within existing frameworks, they have started to envision more strategic and ambitious products and workflows. **Richard Bynum, Chief Corporate Responsibility Officer, PNC Financial Services**, described recent discussions about applying AI to each stage of a 35-step business process before someone asked whether they needed all 35 steps in the first place. “Can we go from Step A to whatever the end is? What about using AI as an enabler to move more quickly through [the entire workflow]?”

Symposium participants had not yet mapped the entire path and corresponding guardrails to make these larger leaps toward “intelligent finance,” but they are starting to see glimmers of what the future state could look like for individual firms, customers, and the broader economy. As they envision that future, the tone is shifting toward more holistic planning processes. Speakers noted the tension between managing AI deployments in the here-and-now

versus taking more time to unleash its larger potential and building new frameworks to support confidence in the technology and the financial system.

Teresa Heitsenrether, EVP, Chief Data & Analytics Officer, JPMorganChase, predicted that the next wave of innovation will occur at a much deeper, more strategic level, as businesses think about the technology’s ability to effectively reduce the marginal cost of the next transaction to zero. She encouraged them to reimagine the way they work: “You have to know your domain, and you have to know what’s capable in the technology. But you also have to actually take a step back and think, ‘if I wasn’t hindered by the way things have always been, how would I think about it from start to finish? If I wasn’t constrained, what would I do?’” She hopes next year’s conversation is more strategic, focusing less on the technology itself and more about the transformation of business functions end-to end.

For people and small businesses, trust and personalization are critical to facilitate broader benefits and faster adoption.

As providers shift toward more ambitious product and process designs, Symposium participants emphasized the importance of centering user needs and impacts to produce larger improvements in users' financial health and security. **Richard Bynum** emphasized the importance of structuring AI implementation planning processes to consider the interests of "those who are not in the room, but [may] perhaps find themselves affected most by it, with the least amount of opportunity to affect that change." Participants also emphasized the importance of proactively designing products and services to meet the needs of and foster

greater trust among populations who have faced financial access barriers.

Speakers were expressly excited about the potential for AI-driven personal financial agentic (PFA) tools to markedly improve peoples' financial lives. Agentic AI has potential to reduce cognitive burdens on consumers and small businesses in managing routine and ongoing tasks. However, given the importance of building and maintaining trust, questions about transparency, control, and safeguards may be critical to users' willingness to adopt AI tools.

“ For our clients, managing their money is a very high stakes and complex endeavor, but up until quite recently, we were [only] able to offer the same set of solutions to everybody.”

LEIGH PHILLIPS
President and CEO
SaverLife



KEY INSIGHT #7

Next-generation AI tools that offer personalized, proactive support in managing users' finances especially promising.

Several participants described the prospect of personal financial agents as one of the most promising AI applications. Providing traditional human-based advice is expensive and difficult to scale, making it hard for LMI populations to access. But AI tools could potentially deliver highly personalized advice, help users navigate real-time decisions, and proactively support follow up activities.

Leigh Phillips described the personalized AI programs and chatbots offered by SaverLife, a nonprofit organization which has an app to help members save money and execute on other financial goals. "We wanted to know what happens if we can lean into the complexity of people's actual financial lives and start building things that can help them in a more personalized, individualized way." Early pilots achieved strong initial results, in some cases triggering a 200% increase in engagement and a rise in monthly savers, from approximately 60% to 90% over a 12-month period.

Proponents also view the agents as addressing a major concern with delivery of conventional financial education content, which research suggests largely fails unless it is delivered "in the moment" that a consumer faces a significant decision. In contrast, panelists emphasized that AI-enabled tools can deliver instantaneous content and even proactively anticipate issues or needs that users may want to consider.

That value proposition grows as these tools include information on a broad range of topics and community resources. For example, speakers discussed the idea of an "AI navigator" which legal aid attorneys, social workers, and other front line workers could use to not just give people financial advice but provide opportunities to take actions which improve their credit score, such as providing a list of places which make positive rent reporting to bureaus. **Clarence Wardell III, VP of Programs, NextLadder Ventures**, highlighted the opportunity to connect

KEY INSIGHT #7

initiatives to build personal financial agents with other types of tools to help LMI consumers navigate issues such as reentry from prison.

The use of agentic AI systems could help enable personal financial agents to execute those recommendations, not just give advice. This feature could especially benefit LMI consumers as they face tight time and capacity demands in managing their daily lives. Similarly, in the small business context, handling routine invoices and payments to known counterparties could

cut administrative time, improve cash-flow predictability, and free owners to focus on operations rather than bookkeeping.

At the same time, participants also warned against magical thinking. **Leigh Phillips** pointed out that such tools can't singlehandedly make up for lack of affordable childcare, health insurance, or other basic needs for LMI households' economic stability, though "I think it's going to give us a real fighting chance to do some really, really tremendous things.



KEY INSIGHT #8

Meaningful transparency, consent, and controls are critical to building trust, as customers decide how much autonomy they want to cede to AI systems.

Symposium participants also emphasized that it is not yet clear how comfortable consumers are in relying on agentic AI systems to act on their behalf in either the digital commerce or personal financial context. Speakers noted that many consumers enjoy certain activities such as planning travel or shopping for clothes and may not want to outsource them to an AI agent. Similarly, both industry and consumer advocacy speakers noted that many consumers do not want to turn shopping or financial decisions entirely over to autonomous AI systems without a more detailed sense of how they will perform.

In light of these dynamics, features that may influence consumers' willingness to test out agents include clear descriptions of the agents' capabilities and goals, data consents and protections, the ability to monitor the agent's actions, and whether the consumer can require the agent to obtain approval before taking particular actions. More broadly, participants emphasized that a gradual, trust-first approach to deployment starts with giving AI tools well-scoped tasks, and only delegating bigger responsibilities as the AI demonstrates reliability. By doing so, one can learn what a user wants and is comfortable with. Transparency

about capability and provenance, combined with incremental delegation, both builds user confidence and may help users identify the benefits of apps and product offerings.

More broadly, speakers pointed to the challenges in managing alignment in goals between customers, providers, and AI systems.

Delicia Reynolds Hand, Senior Director, Digital Marketplace, Consumer Reports, urged industry and regulatory stakeholders to take a proactive approach rather than waiting for a crisis to develop consumer guardrails. Instead, she urged "parallel tracks, where, alongside optimizing for all these commercial uses, we [are also] optimizing for consumer protection."

In the end, participants emphasized that the best tool won't produce long-term financial and economic benefits if consumers or small businesses don't trust it to solve real world problems. "You really can't outrun your customers," said **Tim Flacke, Co-Founder and CEO, Commonwealth**. "A lot of what we're talking about here today is ultimately dependent on end users. Real people deciding that what we are building is useful to them, and that delivers value that is relevant to their lives."

Building better systems can help manage the risks and increase the benefits of AI adoption for individual companies, consumers, and the broader economy.



As individual financial services providers, consumers, and small businesses navigate this rapidly evolving ecosystem to decide when and how to adopt AI-enabled applications, many participants emphasized the potential benefits of improving industry coordination mechanisms and regulatory frameworks. Though some characterize the latter as possibly chilling innovation, common systems can help to foster consistency, efficiency, and trust.

“ *This is begging for a protocol. You want to be able to interrogate what products merchants have available for sale. And then they need to be able to submit orders. And that's not going to be one with a proprietary solution.”*

JOHN COLLISON
Co-Founder and President
Stripe



KEY INSIGHT #9

Cross-industry standards and processes can have multiplier effects by encouraging consistency, facilitating more efficient dispute resolution, and increasing trust among participants.

Symposium speakers emphasized the value of industry standards. For instance, they discussed the release of payment-related protocols and other initiatives to facilitate the operation of AI shopping agents but noted that additional work to develop common practices and promote interoperability would facilitate broader participation and smoother transitions. Online transactions frequently involve not only buyers and sellers but also intermediaries in payments, e-commerce, and search engines. Each needs to be able to obtain information in order to execute a transaction. Interoperability and common standards help to facilitate all parties' participation and confidence in the other actors.

Speakers emphasized the value of being able to use existing systems where possible. For example, consumers' and merchants' banks can continue to conduct the primary Know Your Customer analysis, and web-searching AI agents can leverage existing web headers to communicate with merchants, thus minimizing the amount of new infrastructure those

merchants would need to adopt to prepare for agentic commerce. At the same time, the ecosystem needs to build consensus around processes for sharing the customer's identity information with the agent. That agent then needs to be able to securely move and convey to the merchant identification credentials and other necessary data elements, such as

"For the [agentic commerce] ecosystem to actually operate, each player in it needs to be able to look to all the agents and all the other players and say, 'We agree that these [accepted] standards are, in fact, moving with this agent, and when any particular agent presents a standards based credential, we can engage.' "

SHARDA CARO DEL CASTILLO
Chief Legal and Business Officer



KEY INSIGHT #9

documentation of the consumer's consent for the agent to share information with the merchant and take other actions.

Some speakers focused on the value of using existing payment options such as debit or credit cards, in part because those systems have substantial infrastructure for processing payments and resolving disputes. Others

described stablecoins as having synergies in the agentic commerce space because they may allow developers who want to sell to a worldwide audience to deploy a consistent payment instrument. More broadly, common standards concerning security, privacy, accountability, reliability, transparency, and recordkeeping would increase stakeholders' trust and confidence in engaging with agentic commerce.



KEY INSIGHT #10

Clear and effective regulatory frameworks can also give stakeholders greater confidence in AI adoption and a more competitive marketplace.

Though some participants cautioned against overregulation, they also cited several examples where clear regulatory guardrails could help to reduce uncertainty, foster trust among different market actors, and encourage consistent processes and standards. In the current moment, top-of-mind issues included traditional risk management frameworks as applied to generative and agentic systems, dealing with state regulations, and combatting fraud and scams. More broadly, participants emphasized the importance of risk-calibrated testing and monitoring as technologies and governance practices continue to evolve.

On governance frameworks, several speakers noted that the flood of AI applications is casting a new light on existing guidance and processes. While banking regulators' frameworks for vendor management and model risk management have helped establish principles-based frameworks for managing risks across different types of technologies and providers, speakers suggested they could potentially benefit from clarification as to whether and how they apply to generative and agentic models.

"Third-party risk management guidance was

never written for a world in which AI infuses every link in the supply chain," said **Stephanie Wake, Head of the Agentic AI Center of Excellence, Citigroup**. Similarly, she noted that more traditional quantitative models were the primary concern when banking regulators drafted model risk management guidance in 2011, such that clarifying the application of those frameworks to more advanced forms of AI could potentially prove useful to avoiding "square-peg, round-hole problems".

Ben Hoffman emphasized the value of ensuring that industry and regulators use a common vocabulary to determine what level of risk assessment and expectations apply. "When we say AI, what do we mean? When we say exposure, what do we mean? And which of the risk disciplines are impacted, in which ways?"

As states and the federal government navigate where the power to regulate AI should lie, speakers noted concerns that a deeply fractured regime could complicate scaling products nationwide or produce uneven standards of protection. At the same time, **Sen. Mike Rounds** defended the value of a dual track in which states act as laboratories of regulation

KEY INSIGHT #10

while Congress and federal agencies provide a national framework. “Let’s not disrespect what the states bring to the table,” he said. When states take the lead in regulating AI—notably when it comes to controversial issues such as the protection of minors—that is the states’ “messaging out there that says ‘Congress, you’re not doing a good job of regulating...that’s the shot over the bow saying, ‘get your act together.’”

Other speakers highlighted specific topics on which regulatory action could potentially help address issues that are difficult for financial service providers to manage individually. For instance, increases in fraud and scams activity are a drain on the economy – and can devastate families, particularly as AI-generated fakes continue to advance. Liability for unauthorized transactions or errors made by agentic AI is a particular concern in consumer-facing applications such as digital commerce and personal financial agents. Volatility in securities markets could also potentially increase if a small number of agentic AI systems begin interacting with each other in ways that are manipulative or collusive.

Beyond addressing immediate pain points, some speakers emphasized the importance of further strengthening industry-regulator dialogue,

providing sandboxes and other controlled testing environments before scaling high-impact applications, and identifying the most important outcomes and metrics. Proponents argued that such actions would provide greater nimbleness as technologies and business practices continue to evolve. Although speakers reported some changes in regulators’ attitudes about the use of AI in specific contexts, such as combatting financial crimes and fraud, they emphasized the need for truly risk-based approaches that stay focused on the primary policy concerns.

“The technology is going to continue to evolve,” warned **Teresa Heitsenrether**. “What is it that we’re trying to achieve? Like fair outcomes to credit, stable capital markets, the preservation of personal data—those objectives don’t change regardless of the technology. It’s just a question of how do you ensure that you can test for those outcomes.”

Indeed, these same shifts can potentially change the nature of regulation itself. Indeed, speakers noted that these same shifts can potentially change the nature of regulation and supervision themselves by giving regulators more efficient and effective tools to perform oversight functions.

Conclusion



As AI-driven change continues to accelerate across the financial services ecosystem, the future is not being written by any singular financial entity, any one big technology company, or any single regulator. It is being written collectively by a broad range of stakeholders learning and working together. “We’re not going to get there by standing in our own corners,” said **Melissa Koide, CEO and Director, FinRegLab**. “We need to be participating, exploring together, and building this future of intelligent finance through evidence, experimentation, and dialogue.”

The Symposium identified many different and sometimes contradictory points of view, as participants debated the opportunities and limitations of what AI could offer society. Despite polling that increasingly shows that public fears about AI’s negative impacts on society—including loneliness, isolation, extremism, and negative mental health effects—may outweigh optimism about its beneficial effects, participants at the Symposium believed that we could solve “real problems for real people,” as one speaker put it. Against this backdrop, 2026 will prove more important than ever in defining and delivering good governance for AI in financial services—and positive outcomes for people across the country.

“The models are the worst they’re ever going to be today, and they’re only getting better. I think we are having an internet-like or even much greater transformational moment.”

NADIM HOMSANY
SVP, Head of AI
Strategy & Innovation
BECU





© 2026 FinRegLab. All rights reserved.