



x



CONSUMER
BANKERS
ASSOCIATION

Agentic AI Payments: Navigating Consumer Protection, Innovation, and Regulatory Frameworks

January 2026



Agentic AI Payments: Navigating Consumer Protection, Innovation, and Regulatory Frameworks

JANUARY 2026

Table of Contents

1. Executive Summary	1
2. CBA’s Efforts to Understand Agentic AI, Consumer Opportunities, and Potential Risks	4
3. Overview of Agentic Payments: Technology, Capabilities, and Market Context	8
4. Risks, Challenges, and Misalignments: Potential Harms to Consumers	22
5. Existing Legal and Regulatory Frameworks: Coverage and Gaps	36
6. Role of Regulation in the Context of AI	47
7. Potential Solutions Identified During Symposium	48
8. Conclusion	52

Agentic AI Payments: Navigating Consumer Protection, Innovation, and Regulatory Frameworks

1. Executive Summary

The Consumer Bankers Association (CBA)¹ held a Symposium in the fall of 2025 to address consumer protection issues related to the emergence of agentic payment tools. As artificial intelligence (AI) has developed over the last several years, agentic tools that can perform tasks for people without direct human instruction have emerged as the next wave in AI. The expectation is that consumers will engage any number of agentic payments tools to shop, search, purchase, and pay for a range of goods and services or to make other payments. Agentic payment tools have massive potential to disrupt the existing consumer payments landscape and revolutionize commerce. The Symposium addressed both the promise of agentic payment tools, the potential risks they pose to consumers and other market participants, and it identified gaps in existing regulatory structures.

Based on its discussion with industry participants (including CBA member banks, technology companies, merchants, and payment networks), policymakers, and consumer advocates, CBA identified several key takeaways from this Symposium:

- 1. Existing Consumer Protection Rules Encouraged Digital Payments to Flourish but Have an Uncertain Application to Agentic Payments.** Congress enacted the Electronic Fund Transfer Act (EFTA) in 1978, when electronic payments were as new as agentic commerce is today. EFTA established clear rules of the road for consumers and banks, including limitations on consumers' liability for many types of unauthorized transactions.

This regulatory framework provided consumers with the confidence necessary to transition large swaths of their payments from cash and checks to debit cards and ACH over the subsequent decades. More recently, however, new types of fraud emerging from P2P payments and other sources, have identified some gaps in EFTA's structure. As agentic commerce leads to a potential explosion of agentic payments, additional gaps may emerge. Specifically, the general rule in

¹ The Consumer Bankers Association is a member-driven trade association, and the only national financial trade group focused exclusively on retail banking—banking services geared toward consumers and small businesses. As the recognized voice on retail banking issues, CBA provides leadership, education, research, and federal representation for its members. CBA members operate in all 50 states. They include the nation's largest bank holding companies as well as regional and super-community banks. Eighty-three percent of CBA's members are financial institutions holding more than \$10 billion in assets.

EFTA that limits consumers' liability for unauthorized transactions may not apply when agents are involved. When a consumer gives an access device, such as a payment card or payment credentials, to another person and that person initiates transactions the consumer did not intend for that person to initiate, the consumer may be liable for those transactions—not the financial institution that issued the card or credential. Put another way, consumers may be liable for mistakes their agents make and these mistakes could be costly. These exceptions could have significant impacts on the evolution of consumer protection in connection with consumers' use of agentic payment tools.

- 2. New Payments Rails May Emerge for Agentic Payments, Perhaps with Fewer Consumer Protections.** Today, the vast majority of consumer payments occur over the major payment card networks through which consumers initiate debit and credit card payments. These payment card networks have all voluntarily added protections for cardholders that exceed those required by federal law discussed above, including that cardholders have no liability for unauthorized transactions. These card networks also provide consumers with additional chargeback rights. To pay for these expenses, the card networks impose interchange fees on merchants for accepting their cards.

As agentic payment tools develop alongside other payment technologies, such as stablecoins, new payment rails are beginning to emerge. Some crypto platforms and AI technology companies are beginning to develop new payment rails that allow and encourage consumers to pay outside of the traditional card networks. Some anticipate these rails may have lower costs than the interchange fees imposed by the payment card networks. It is still early for these payment rails, but there is no guarantee they will evolve with the same protections as the traditional card networks. Agentic payment tools, seeking lower prices or other benefits, may route payments on these new rails. Consumers may benefit in the short run but could be harmed in the long run if their transactions on these rails are not protected in the same way, particularly if the federal laws discussed above do not apply to some or all crypto transactions.

- 3. Banks Will Play a Key Role in Agentic Commerce.** As agentic commerce evolves, banks will have a major role in that evolution. Banks will be involved in agentic commerce because they are integral to most consumer payment transactions. Banks issue virtually all payment cards, transact most ACH payments, and hold most consumer asset accounts. When consumers use these tools to enable agentic commerce and something goes wrong, many will seek support and compensation from their financial institutions. In addition to protections for unauthorized transactions and other errors, banks have served as a clearinghouse for consumers' financial challenges. If agentic commerce

causes the number of consumer transactions to multiply exponentially, the customer inquiries about these transactions will also increase by the same proportion.

- 4. Immediate Statutory and Regulatory Changes Are Unlikely and May Be Unnecessary.** Despite the issues identified above and during the Symposium, immediate government action appears unlikely in the short term. The Trump Administration has given clear signals that policymakers should allow AI and agentic tools to flourish. This current approach to federal regulation, however, should not cause developers to ignore consumer protections. It is incumbent on industry participants to consider how best to balance access to these exciting new tools with consumer protections present with other types of payments. This may mean industry should consider development of private network rules like those that, in addition to federal consumer finance law, have allowed electronic and card-based payments to flourish. Private network rules applicable to all participants in the agentic payment tool ecosystem—including developers, merchants, marketplaces, banks, and others—could allow for coherent development of a robust system that allows consumers and others to benefit from agentic payment tools while minimizing their potential downsides.

2. CBA's Efforts to Understand Agentic AI, Consumer Opportunities, and Potential Risks

CBA is the only member-driven trade association focused exclusively on retail banking whose mission is to partner with the nation's leading retail banks to promote sound policy, prepare the next generation of diverse bankers to lead the industry, and finance the dreams of consumers and small businesses.

As part of its work, CBA scopes potential policy risks for its members and their customers. It conducts policy scoping work as part of its advocacy and research function, and identifies emerging issues that may require industry collaboration, regulatory guidance, or policy innovation. One of these issues is consumers' use of agentic payment services. CBA initially determined agentic payments warranted a closer look for three primary reasons:

1. **Prior Experience with Fraud and Scams in Consumer Payments:** CBA's member institutions have significant experience with payments fraud and scams. That experience—including losses from authorized transactions that later proved fraudulent, customer service burdens, and reputational damage—raised concerns about how agentic AI and agentic payment tools could amplify existing fraud and scam concerns or create new ones at scale.
2. **Lack of federal regulatory focus:** To date, federal and most state regulators have not issued specific guidance addressing agentic AI and agentic payment tools in consumer payments nor have they indicated an intention to do so in the near future. Given how quickly AI and agentic payments are developing, CBA determined it was important to engage in a public study and analysis of these issues to guide agentic payment tool developers and policymakers.
3. **Rapid technological development:** CBA recognized that agentic AI in payments is advancing quickly, with major technology providers and payment networks announcing new protocols and capabilities regularly. Early industry engagement and policy discussion seemed warranted to help shape this trajectory.

CBA's 2025 Agentic AI Initiative

To better understand the issues presented by agentic payments, CBA launched a multi-phase initiative in 2025:

Webinar Series (Summer 2025): CBA convened two foundational webinars hosted by Kelvin Chen, CBA's Head of Policy, that it made available to the general public. The first, which CBA held in partnership with Visa, included Kelvin and Reed Bouchelle, Visa's VP, North America Product. They examined the technical and operational aspects of agentic payments, including use cases, architecture, and current implementation challenges.² This webinar examined the benefits to consumers of agentic commerce, including time savings, how banks will need to adapt to their customers' desire to engage with agentic payment services, and the need to secure consumer data. The second webinar, conducted with Eric Goldberg, current partner at Davis Wright Tremaine LLP, examined the liability, authorization, disclosure, and regulatory implications of agentic payments under existing regulatory frameworks.³ Kelvin and Eric discussed how Congress, when it enacted EFTA in the late 1970s, established a framework that would later allow electronic payments to flourish. The webinar reviewed this and related laws, discussed in detail below, to establish a framework for further regulatory discussions around agentic payments.

Agentic Payments Symposium (Fall 2025): Building on the webinar series, CBA hosted a two-day symposium on agentic AI in payments in Washington, D.C. The Symposium convened senior representatives from a range of stakeholders in agentic payments, including representatives from: federal consumer protection and banking regulators; large financial institutions who are CBA members; payment processors and card networks; fintech companies; AI developers; merchants; consumer advocates and

² Consumer Bankers Association, *CBA Webinar Spotlight: Agentic AI and the Future of Payments – What Bankers Need to Know* (July 7, 2025), <https://vimeo.com/1099390495/a91bb37409?share=copy&fl=cl&fe=ci>.

³ Consumer Bankers Association, *Agentic AI and Consumer Payments: Legal & Regulatory Implications* (July 23, 2025), <https://vimeo.com/1103770222?share=copy&fl=cl&fe=ci>.

academic experts; payments infrastructure providers; and legal and regulatory experts, including but not limited to:

Akerman LLP	FinRegLab	McGlinchey Stafford PLLC
Alliance for Innovative Regulation	Fintech Takes	National Consumer Law Center, Inc.
Ally Financial Inc.	First Financial Bank	Office of the Comptroller of the Currency (OCC)
Aspen Institute Financial Security Program	FIS Global	Orrick, Herrington & Sutcliffe LLP
Atlantic Union Bank	Fiserv, Inc.	Paul Hastings, LLP
Bank of America, N.A.	Frost Bank	PayPal
Bank of Hawaii	George Mason University	Plaid Inc.
Banner Bank	Georgetown Law	The PNC Financial Services Group, Inc.
Conference of State Bank Supervisors	Google LLC	Relay Network, LLC
Consumer Federation of America	Jenius Bank	Stripe, LLC
Experian Information Solutions, Inc.	JPMorganChase	Synchrony Bank
Federal Deposit Insurance Corporation (FDIC)	KeyBank National Association	TD Bank, N.A.
Federal Trade Commission (FTC)	Mastercard International Incorporated	Zeta Global Holdings Corp.

The first day and a half of the Symposium was intended to educate participants about agent payments, potential use cases, regulatory risks, and potential regulatory solutions. During the second day, participants broke into two working groups where they addressed these issues in more detail in wide ranging brainstorming sessions. The complete Symposium agenda is attached as Appendix A.

Except where noted in this White Paper, the Symposium followed Chatham House rules. Under this structure, CBA invited participants to share their thoughts freely and noted that their comments would not be attributed to them or their employers. This was to allow participants to share information and insights without fear of attribution. This format enabled candid discussion of concerns, uncertainties, and potential solutions.

CBA's expectation that the Symposium was just the first step in a broad discussion of these issues over the next several years.

Important Disclaimer: *Participation in the Symposium by a person or entity should not be construed as endorsement of anything in this White Paper. This White Paper reflects CBA's summary and analysis of the Symposium and is not intended to represent any individual or group consensus.*

3. Overview of Agentic Payments: Technology, Capabilities, and Market Context

For purposes of this White Paper, “agentic AI” means artificial intelligence systems that can act autonomously on behalf of a user, making decisions, and taking actions. Agentic tools can direct transactions without direct human intervention at each decision point. This definition is deliberately narrow and excludes:

- **Generative AI systems** that respond to prompts and produce outputs (text, images, code) but do not autonomously execute transactions on a user’s behalf.
- **Existing automation and smart tools** such as autopay systems, market-making algorithms, or conditional purchase orders that operate based on predetermined rules or triggers but that are not actually autonomous.
- **AI tools that assist human decision-making** but do not execute transactions autonomously (*e.g.*, dashboards that synthesize information or chatbots that answer customer service questions).

Agentic AI systems can:

1. Receive instructions or mandates from users (expressed through natural language, structured parameters, or pre-configured preferences).
2. Autonomously evaluate options and make decisions aligned with those instructions.
3. Execute transactions or take actions.
4. Learn from outcomes and adapt behavior accordingly.
5. Interact with multiple systems, APIs, and counterparties to accomplish tasks.

Agentic systems can complete all of these tasks without further user input.

Table I. An Emerging Consensus on What Agentic AI Is

Feature	Generative AI	Standard Automation (e.g., Autopay)	Agentic AI (Agentic Payment Tools)
Primary Function	Creates content (text, images) based on prompts.	Executes pre-set rules (e.g., “pay bill on 1st”).	Acts autonomously to make decisions and take actions based on broad inputs.
Decision Making	Low; responds to user input.	None; follows rigid triggers.	High; evaluates options and aligns with broad goals.
User Role	Active: User prompts, reviews, and uses output.	Setup: User sets the rule once; passive execution.	Passive: User sets a broad goal; Agent executes transaction.
Example	“Write an email to my bank.”	“Pay my electric bill on the 15th.”	“Find and buy the best laptop under \$500 for a 13-year old.”

Agentic payments are types of agentic AI systems that can initiate, manage, and execute financial transactions at the point of checkout or payment authorization. In an agentic payment scenario, a consumer provides an AI agent with a goal or set of instructions (*e.g.*, “find the best deal on flights to Washington, D.C., in the next 30 days” or “pay my bills as they arrive in the order that is most beneficial to my personal financial situation”). The AI agent autonomously searches for options, evaluates alternatives against the consumer’s criteria. The AI agent then executes a purchase or payment transaction when it determines that the options meet the consumer’s parameters. The agentic payment tool may notify the consumer of the transaction after it has been executed.

Agentic payments differ from existing electronic payments in a material way: the consumer does not actively participate in the transaction at the moment of authorization. Instead, authorization is prospective and conditional, based on parameters that the consumer established in advance by the consumer working in the construct of the agent. We caution that use cases for agentic payments are constantly expanding and that some tools that claim to be agentic payments or involve agentic AI may not be so. The terminology around agentic AI remains unsettled, and some have applied the “agentic” label to products and capabilities that do not meet the above definition. For this White Paper, we use the more restrictive definition outlined above but caution that, as with any new technology, use of the terminology in other contexts may vary.

Current State of Agentic Technology: Emerging Capabilities and Market Reality

Agentic AI capabilities in payments and agentic payments specifically are nascent but developing rapidly. As we approached the 2025 holiday shopping season, it seemed as if fintechs, merchants, and AI companies announced new agentic payment tools every day.⁴ It is too early to say whether consumers used these tools and, if so, whether they worked as intended. Nevertheless, at the time of the Symposium in fall 2025, attendees largely agreed that while there had been limited deployment of agentic payment tools, society is on the doorstep of a rapid expansion.

Speakers noted that a growth in agentic payment tools was likely due to the immense benefits such tools could offer consumers. Symposium attendees largely agreed that agentic tools are expected to reduce consumers' cognitive load. A tool might understand a consumer's travel parameters and select flights and hotels that satisfy those parameters. The consumer would not have to spend hours poring through search results. AI tools might also make it easier for consumers to understand and synthesize thousands of product reviews for basic commodity products. For example, an agentic tool might know details about a consumer's family and health and identify the best dental floss based not only on the consumer's parameters, but the AI tool's ability to synthesize hundreds or thousands of reviews across various websites and merchants. Put another way, an AI tool may narrow down the number of options returned in a search based on its understanding of a consumer's preferences and requirements. And then the agentic payment tool will purchase that item without further consumer input.

The pace of AI advancement means that capabilities that seemed speculative just a few months ago may quickly become practical. Major technology companies (including Google, with its announcement of the Agent Payments Protocol) and payments networks are publicly signaling significant investment in agentic payment infrastructure. Since the Symposium ended, many major merchants and AI companies have announced various AI tools related to consumer activity, including agentic payments tools.⁵ While it is beyond the scope of this White Paper to review those

⁴ Press Release, Mastercard, Mastercard unveils new tools and collaborations to power smarter, safer agentic commerce (Sept. 10, 2025), <https://www.mastercard.com/us/en/news-and-trends/press/2025/september/mastercard-unveils-new-tools-and-collaborations-to-power-smarter-safer-agentic-commerce.html>; Beth Duckett, *Visa and Mastercard both launch new agentic AI payments tools*, DIGIT. COM. 360 (Oct. 16, 2025), <https://www.digitalcommerce360.com/2025/10/16/visa-mastercard-both-launch-agentic-ai-payments-tools/>; Kevin Miller, *Introducing our agentic commerce solutions*, STRIPE (Oct. 7, 2025), <https://stripe.com/blog/introducing-our-agentic-commerce-solutions>; *Powering AI commerce with the new Agent Payments Protocol (AP2)*, GOOGLE CLOUD (Sept. 16, 2025), <https://cloud.google.com/blog/products/ai-machine-learning/announcing-agents-to-payments-ap2-protocol>.

⁵ *Amazon's new AI-powered shopping feature 'Help Me Decide' makes it easy to quickly pick the right product*, AMAZON (Oct. 23, 2025), <https://www.aboutamazon.com/news/retail/amazon-things-to-buy-help-me-decide-gen-ai>; *OpenAI and PayPal Team Up to Power Instant Checkout and Agentic Commerce in ChatGPT*, PAYPAL (Oct. 28, 2025), <https://newsroom.paypal-corp.com/2025-10-28-OpenAI-and-PayPal-Team-Up-to-Power-Instant-Checkout-and-Agentic-Commerce-in-ChatGPT>.

projects in detail, it is important to note that there is significant momentum behind developing consumer tools in this space.

Symposium participants discussed the current state of agentic payments and the AI marketplace. Many agreed that consumers are seeking to use AI tools in all facets of their lives. Data shows use of AI has exploded since 2023⁶ as has investment of industry participants into developing new AI tools.⁷ One study found that 81 percent of consumers expect to leverage AI in their shopping and 42 percent of consumers would allow AI to shop entirely on their behalf, at least in one product category.⁸ If this is borne out, \$1.3 trillion of online commerce will be impacted by agentic commerce tools.⁹

Participants in the CBA Symposium had additional thoughts on agentic payment growth:

- Consumers will continue to seek out AI and agentic payments tools.
- Consumers will favor agentic payments tools they trust and where methods of performing tasks on behalf of the consumer are transparent.
- While some consumers may be brand agnostic as to the AI payment tools they use, others may not be. For example, a consumer might instruct an agentic tool to purchase the best mutual funds for the consumer's situation without consideration for the provider of the mutual fund. Of course, some consumers may instruct AI tools to frequent preferred brands or providers due to prior experience, preferred payment networks, rewards programs, purchase protection, or simply familiarity.
- Merchants' means of connecting to and partnering with AI providers could impose significant technology and infrastructure costs on merchants. Large merchants with sophisticated payment infrastructure may be able to adapt relatively quickly. Smaller merchants may face significant technical and operational hurdles. This could influence the agentic payment tools merchants may make available to their customers.¹⁰

⁶ Bernard Marr, *20 Mind-Blowing AI Statistics Everyone Must Know About Now*, FORBES (June 3, 2025), <https://www.forbes.com/sites/bernardmarr/2025/06/03/mind-blowing-ai-statistics-everyone-must-know-about-now-in-2025/>.

⁷ *The 2025 AI Index Report*, Stanford University, Human-Centered Artificial Intelligence, <https://hai.stanford.edu/ai-index/2025-ai-index-report> (last visited Dec. 5, 2025).

⁸ BCG, AGENTIC COMMERCE, SHOPPING AND PAYMENTS RE-(AI)MAGINED 2 (2025), <https://media-publications.bcg.com/Agentic-Commerce.pdf>.

⁹ *Id.*

¹⁰ Graham Barlow, *Amazon blocks ChatGPT's new shopping agent – what the fallout could mean for you*, TECHRADAR (Nov. 27, 2025), <https://www.techradar.com/ai-platforms-assistants/chatgpt/amazon-blocks-chatgpts-new-shopping-agent-what-the-fallout-could-mean-for-you>.

Anticipated Use Cases and Benefits

Given that agentic payments are in their early days, speakers at the Symposium and participants also hypothesized about potential near- and medium-term consumer use cases for agentic payment tools:

1. **Intelligent Buying of Goods and Services.** Consumers could use agentic payments tools to select and purchase goods and services for them. Scenarios ranged from luxury goods to commodities. A few examples:
 - a. A consumer might instruct an agentic tool to buy the best light bulbs for the consumer so that the consumer does not have to wade through many similar options.
 - b. A consumer might instruct an agent to “buy Taylor Swift tickets once the price falls below \$500 in the next 30 days.” The agent would monitor pricing, evaluate secondary markets, compare options, and execute the purchase when criteria are met. This use case could extend across categories: concert tickets, travel, consumer goods, services, etc. The benefit to consumers is clear: the agent removes the burden of ongoing monitoring and can execute immediately when conditions are met.
2. **Bill Payment and Payment Management.** AI agents could manage subscription services, utility bills, insurance premiums, and other recurring payments. One participant noted this could eliminate autopay. Rather than authorizing recurring payments to specific merchants (which is what autopay does today), a consumer could authorize an agent to monitor bills and pay them intelligently. If a bill seems unusually high, the agent could flag it for human review before paying. If an alternative provider offers the same service at lower cost, the agent could evaluate and execute a switch. This represents a more intelligent and consumer-protective form of automated payment. The agent could also monitor invoices as they arrive, verify them against receipts or user agreements and execute payments it views as in the best interest of the consumer. Symposium participants noted that this use case is already emerging at the B2B level, with HVAC companies, logistics providers, and other service-heavy businesses deploying back-office AI to manage vendor payments.
3. **Financial Management and Optimization.** Agents could manage savings, investment rebalancing, borrowing decisions, and insurance optimization. These use cases may involve more complex decision-making. An agent must evaluate the consumer’s risk tolerance, time horizon, financial goals, and existing portfolio, then execute transactions (purchases or sale of securities,

purchasing insurance, transfers between accounts to maximize yield, etc.) when market conditions or other triggers align with the consumer's objectives.¹¹

4. **Insurance Claims Resolution.** An agent could manage insurance claims processes, coordinating with claims adjusters, providing necessary documentation, and negotiating settlements. This would reduce friction and time required for consumers to resolve claims.
5. **Multi-Agent Scenarios.** As agentic AI becomes more prevalent, scenarios could emerge where a consumer's agent coordinates with agents representing other parties. For example, a consumer's agent might negotiate with a merchant's agent over price, terms, or warranty conditions before executing a purchase. These multi-agent scenarios introduce additional complexity and risk.

Benefits to Consumers and the Marketplace

Symposium participants emphasized several potential benefits of these sorts of agentic payment tools:

1. **Reduced Cognitive Load:** For routine financial decisions, agentic AI can eliminate the need for repeated human decision-making, freeing consumers to focus on higher-value activities.
2. **Better Financial Outcomes:** Agents with access to comprehensive financial data could help consumers make more informed decisions—comparing products, understanding trade-offs, and identifying options—that align with stated goals, freeing consumers from having to calculate those options on their own.
3. **Faster Decision-Making and Transactions:** Rather than consumers spending time researching products, reading reviews, and manually executing purchases, agents can synthesize information and execute transactions rapidly.
4. **Better Prices and Higher Quality:** Agents can continuously monitor prices, compare quality metrics, and identify opportunities, including new merchants or products, that consumers might overlook.
5. **New Entrant Opportunities:** By removing human involvement in repetitive procurement tasks, agentic payments could open opportunities for small merchants or upstart brands with superior products to reach consumers, even if those merchants lack the brand recognition or advertising budgets of larger competitors and incumbents.

¹¹ This is distinct from non-intelligent investment tools that execute limit orders, stop-loss orders, stop-limit orders, etc. Those tools execute rules but do not make intelligent decisions for the consumer.

Agentic Payments Participants

Given the potential breadth of agentic payments, there are many potential participants in the agentic payments ecosystem, including:

- Consumers
- Merchants
- Merchants of Record, and Payment Facilitators
- AI providers and developers
- Banks
- Data aggregators (*e.g.*, Plaid and Finicity)
- Payment processors, including acquirers and payment facilitators
- Non-bank fintechs (*e.g.*, non-bank lenders, money services businesses, neobanks, etc.)
- Crypto companies including stablecoin issuers, crypto platforms, and companies that facilitate crypto movement and self-hosted wallets
- Payment networks and related self-governing bodies (*e.g.*, Visa, Mastercard, American Express, Discover, PCI, NACHA, and SWIFT)
- Digital wallet providers and related intermediaries

As we have seen in the growth of AI partnerships in 2025, entities from one or more of these categories may work together to develop AI tools for consumers' use. While the extent of oversight may vary, many regulators have oversight over providers in some of these categories. For example, large banks are subject to oversight from prudential regulators, the CFPB, the Department of Justice, and others. AI providers and merchants may be regulated by the Federal Trade Commission and state attorneys general but are not subject to regular examinations.¹²

¹² Beyond the scope of this White Paper, there exists a robust debate about the extent of federal and state regulation of AI. We do note that the outcome of that debate might have a significant impact on the issues we have identified.


A crucial part of assessing the involvement of these participants is how they expect to be compensated. Symposium participants suggested that many existing means of payment revenue may be adapted for agentic payments:

- **Referral Fees.** Merchants or others may compensate agentic payment tools for referring business to them.
- **Interchange Fees.** Card networks charge interchange fees to merchants every time they conduct a transaction. ACH transactions also incur fees, albeit much lower than card network fees.
- **Gas Fees.** Stablecoin gas fees are transaction costs paid on blockchains like Ethereum to process transfers of stablecoins.
- **Direct to Consumer Revenue.** Consumers may pay one time or subscription fees to access and use agentic payment tools.
- **Advertising.** Merchants or others may pay AI providers to promote their goods or services to consumers.
- **Data Monetization.** Agentic payment tools may intake substantial data about consumers and their preferences. AI providers may be incentivized to sell this data—aggregated or otherwise—to third parties.
- **Other.** Participants hypothesized that agentic payments may lead to the development of new forms of compensation as well.

Agentic Payments May Use New or Different Payment Rails

Today, consumers primarily use several well-defined payment rails to execute payments, but that may quickly change given the growth of stablecoins and other alternatives.

According to the Federal Reserve, consumers primarily make credit and debit card payments over the major payment card networks, use ACH, or cash:



Source: Federal Reserve 2025 Diary of Consumer Choice¹³

The two largest non-cash rails – card payments and ACH each have their own set of private network rules that govern all aspects of their respective payment ecosystems, including disputes to address. However, growth of other payment rails seems likely and may be accelerated by agentic payment tools. Additionally, some participants expressed concern that the existing rails and related mechanisms may be insufficient to handle the burden some expect to be imposed by agentic payments.

Some participants noted Coinbase’s efforts to take advantage of the http x402 status code to develop the x402 payment protocol for use with agentic payments. According to Coinbase, x402 is an internet-native payment protocol for fast, cheap, AI-friendly payments over HTTP, used for pay-per-use API models and AI services. Specifically, x402 activates the long-reserved “402 Payment Required” status code in HTTP to enable automatic, on-chain stablecoin payments directly within standard website flows. In a typical interaction, a client (often an AI agent or app) requests a resource, the server replies with an HTTP 402 containing price and payment parameters, the client (or a facilitator) submits a signed stablecoin payment, and upon verification the server responds with the requested data. The protocol is directly aimed at agentic commerce. With it, agentic payment tools may be able to autonomously pay for goods with near-instant settlement and low fees compared to card or banking rails that were built for human users transacting in fiat currencies. Importantly, x402 payments settle instantly and do not today contain mechanisms for reversals or chargebacks.¹⁴

¹³ 2025 Diary of Consumer Payment Choice, FED. RSRV., <https://www.frbervices.org/news/research/2025-findings-from-the-diary-of-consumer-payment-choice> (last visited Dec. 5, 2025).

¹⁴ ERIK REPEL, ET AL., X402: AN OPEN STANDARD FOR INTERNET-NATIVE PAYMENTS 8 (2025), <https://www.x402.org/x402-whitepaper.pdf>.

This is just one example. More broadly, the growth of crypto payments, including those making on-chain payments and executing smart contracts serve to change the traditional payment mechanisms from the existing rails, set forth above. Stablecoins are digital tokens typically pegged to the U.S. dollar or other fiat currencies and are backed by reserves such as cash or short-term Treasuries. Stablecoin payments are typically tracked on a blockchain. Smart contracts, embedded in the blockchain, are self-executing blockchain programs that automatically enforce predefined terms upon conditions being met. On-chain payments involve direct, peer-to-peer stablecoin transfers recorded immutably on the blockchain for near-instant settlement without intermediaries. The growth of stablecoin payments and their automated nature makes them a natural venue for agentic payments.

Specific to AI, several Symposium attendees also referenced the Model Context Protocol (MCP). MCP is an open standard that allows AI systems to securely connect to external data sources and tools, like a universal adapter that simplifies access without needing custom setups for each one.¹⁵ In payments, MCP can enable AI agents to interact with services like Stripe or PayPal for real-time tasks such as checking balances, creating invoices, or processing transactions, making automated financial operations faster and more reliable. While MCP is in its early days, Symposium attendees spoke about it quite favorably as a means of speeding up the implementation of agentic and AI-based processes.

Outside of crypto, other payment types continue to evolve including non-card payments, such as buy now pay later (BNPL) and other forms of embedded credit, that are increasing the volume of payments through alternative means.¹⁶ Consumers are also increasingly enticed to use “pay by bank” mechanisms that transmit payments over ACH rather than card network rails. Dispute resolution requirements for these new payment mechanisms vary due to the lack of applicable law or industry protocol.

The emergence of these alternatives caused some attendees to fear a “race to the bottom.” As noted, these emerging protocols may lack the same protections provided by existing payment rails. And, as is discussed further below in the discussion of risks, the application of existing rails to some of these payment methods is, at best, unclear. One participant noted that these nascent private networks may want to limit participants’ liability to consumers and others in order to incentivize AI providers, merchants, and other innovators to participate in their programs. Additionally, given that protections are put in place to protect consumers, merchant agents could make payment rails act in the merchant’s favor or encourage agents to trade those protections for things such as lower prices. Consumers could be unaware of such tradeoffs. Another participant

¹⁵ *Introducing the Model Context Protocol*, ANTHROPIC (Nov. 25, 2024), <https://www.anthropic.com/news/model-context-protocol>.

¹⁶ Patrick Cooley, *BNPL spending storms on*, PAYMENTS DIVE (Dec. 3, 2025), <https://www.paymentsdive.com/news/bnpl-spending-storms-on/806866/>.

observed that consumer distrust of these new rails may occur if consumers have negative experiences where they learn that these new rails are unprotected.¹⁷

Table II. The Payment Rails Trade-Off

Payment Rail	Settlement Speed	Chargeback Rights	Liability Clarity	Reversibility
Card Networks	Delayed	Yes	High	Yes
ACH	Delayed	Limited	Medium	Limited
BNPL / Embedded Credit	Varies	Fragmented	Medium-Low	Varies
Stablecoins / x402	Instant	None	Low	No

Agentic Payment Tools Collection and Use of Consumer Data and Open Banking Issues

The Symposium also discussed the collection and use of data in the development of agentic payments tools. If an AI tool is engaged to find what a consumer wants and can afford, attendees agreed it will need data about the consumer. More broadly, for some applications, data underpins the entire AI technology.

Participants noted that different levels of functionality require different amounts of data. An agent assigned broad tasks may need all of a consumer’s data— financial and non-financial—to entirely accomplish those tasks. When consumers engage tools directly, they may elect to share data with the tool to allow the tool to access payment mechanisms, banking data, and other non-financial data such as their email. AI tools may also obtain data from third parties, such as merchants, banks, and employers.¹⁸ This data could allow consumers to use agentic payment tools to bypass traditional retail structures, which some have termed “open commerce.”

Participants suggested that health and financial data are the most important types of personal data that could be used in agentic tools. One participant suggested, and many agreed, that the most successful AI tools will be those that have access to the most data.

¹⁷ Some argue similar concerns slowed the growth of cryptocurrency use as a replacement for fiat currencies in consumers’ daily transactions. See, e.g., Tonantzin Carmona, *Protecting the American public from crypto risks and harms*, BROOKINGS INST. (June 2, 2025), <https://www.brookings.edu/articles/protecting-the-american-public-from-crypto-risks-and-harms/>.

¹⁸ Required Rulemaking on Personal Financial Data Rights, 89 Fed. Reg. 90838 (Nov. 18, 2024).

However, as discussed below, there was some consensus that data-related issues were not the most significant risks of agentic payment tools.

Participants also addressed the role of open banking in agentic payments. While not a core part of the Symposium, the ongoing debate about open banking and the CFPB's open banking rulemaking was mentioned throughout the two-day event. Without engaging in the debate about the contours of safe access to and sharing of consumer data, participants agreed that open banking will remain part of the payments and financial transaction landscape. Banks understand that consumers should be able to access their information. However, given their role, banks noted they are concerned about the safety of the system and their customers' data. Several participants noted that a key part of open banking is the provision of safe APIs for use by consumers. Some of these are implemented by core processors, data aggregators, and others.

The Role of Banks in Agentic Payments

Symposium participants generally agreed that banks play, and will continue to play, a major role in the development of the agentic payments ecosystem. Banks maintain accounts for consumers and the existing payment rails discussed above all rely on funds moving in and out of these accounts. Banks are the only entities, in most cases, which can issue network branded payment cards, and they are, with some exceptions, the only parties that can initiate and receive ACH payments and send wire transfers. Other than spending cash or perhaps cryptocurrency, virtually every consumer payment transaction conducted today involves a bank. Given their central role, banks have been instrumental in developing a safe and secure ecosystem for their customers' payments. For example, banks were integral in the development of card network and NACHA rules, which allowed electronic payments to flourish.

As a result, banks sit at the center of the agentic payments ecosystem. Bank participants in the Symposium noted that because banks have the ability to leverage their experience and place in the market to benefit customers, they could lead in the development of a safe and secure agentic payments system. One participant noted that existing bank models for managing risk may be adapted to agentic payments. Today, banks rely on a vast number of financial and risk models to engineer core risk and business operations. Banks use these models to both inform decision making, measure risk, and estimate asset values. This has led to the development of Model Risk Management (MRM), which includes strict validation and compliance procedures, which causes banks to operate at higher standards as compared to other industries with less regulatory oversight. Participants emphasized that banks use MRM to identify, assess, and mitigate the potential for adverse outcomes arising from the use of models in their decision-making. In essence, MRM is a quality control system for the models that financial institutions rely on to measure risk, detect fraud, and ensure compliance.

Relatedly, banks maintain particularly tight oversight of their vendors. Regulators require banks to employ models and engage with vendors in a proper way, through sufficient process oversight, due diligence, and testing. Regulators then examine banks to confirm they have complied with these obligations. Given this, banks have substantial experience selecting and monitoring third-party vendors through implementation of guidance from their prudential regulators and the CFPB. Banking is unique in the sense that non-banks largely do not face similar oversight into their vendor management practices. Participants suggested this is a key distinction between banks and other participants in the agentic payments ecosystem.

Participants representing banks agreed that banks cannot be slow in understanding how their customers will be using AI and agentic payments tools because banks should expect customers to reach out for help when agentic transactions go wrong. Today, banks are the primary touch point for customers when there is a dispute or issue involving a payment they have made. While consumers are encouraged to try to work things out with merchants, consumers are instructed to contact their bank either when the merchant is unhelpful or the consumer suspects unauthorized transactions or fraud involving their account. These instructions are not misplaced. The primary consumer protection rules in the Truth in Lending Act (TILA) and EFTA – see Section 5.B below – also disclose to consumers that they must contact their banks about these issues.

More broadly, several participants expressed concern that customers will contact banks with general customer service inquiries and that these inquiries could have several negative impacts on banks and their customers. Several participants posited that banks would receive an outsized number of customer service questions regarding their customers' use of agentic payment tools because AI and technology companies have historically provided limited customer service to consumers.¹⁹ Customers also expect banks to make them whole for failed improper payment transactions. If that persists, and perhaps even if there is more robust support from these providers, customers will turn first to their banks for help. Bank representatives at the Symposium agreed with this based on their experience with scams and frauds from P2P payment services over the last several years. If banks fail to address customers' issues—even if those issues are not the banks' fault—banks fear a loss of goodwill. Conversely, some asserted that banks are natural stewards of trust from customers and others and that this trust allows banks to play a potential role in integrating agentic technologies into customer payments. Potentially, banks can become a trusted digital advocate for their customers. Ultimately, bank participants expressed that they have a shared responsibility with the AI industry to protect customers.

Finally, some bank participants suggested different approaches to regulation generally and specifically to AI. They noted that banks typically maintain close relationships with

¹⁹ Banks, unlike technology companies, are required by federal law to provide customer service. Banks also have historically relied on strong connections and customer service relationships with their customers.

their regulators and that this helps to ensure consistency in regulations and their application to real-life situations. Consumer advocate participants agreed with this principles-based approach that remains technology agnostic.

4. Risks, Challenges, and Misalignments: Potential Harms to Consumers

A central part of the Symposium was the identification of risks and harms to consumers, merchants, and the broader payments ecosystem. Participants catalogued potential harms. As detailed below, the potential harms can be divided into several broad categories. Attendees also identified potential harms to merchants, banks, and other participants in the agentic payments ecosystem.

Most of the discussion focused on potential harms to consumers from agentic payment tools that make mistakes. The potential for mistakes—intentional or otherwise—is significant and could occur due to a range of reasons including technological shortcomings, misaligned incentives, or consumer misunderstandings.

4.A. Agents Not Operating as Intended.

One of the most fundamental risks agentic payments pose is that agentic payment tools may not act in the consumer's best interest, either because they are designed not to, because they fail operationally, or because conflicts of interest distort their decision-making. Below we catalog some specific ways these risks may manifest.

4.A.i. AGENT OPTIMIZES FOR INTERESTS OTHER THAN THE CONSUMER'S.

Symposium participants discussed several ways in which an agentic payments agent may optimize for a person or entity's interests other than those of the consumer. In some ways, this means an agent is not actually an agent. The classic definition of an agent is a person (the principal) who manifests assent that another (the agent) shall act on the principal's behalf and subject to the principal's control.²⁰ Under agency law (discussed below in Section 5.A), if an agent acts for someone else's interest, the agent violates its duty to the principal.

Commenters suggested agents may deviate from their duty to the principal in various ways, such as:

- An AI developer might configure an agent to prefer payment rails or merchants that generate higher revenue for the developer (through referral fees, data monetization, etc.), even though the consumer might benefit from or prefer a different rail or merchant. For example, an agentic tool might choose a means for completing its task that result in higher prices, more difficult delivery or return processes, or fewer legal protections because of the lack of applicable law or private network

²⁰ RESTATEMENT (THIRD) OF AGENCY § 1.01.

rules. Some of these problems may not be apparent to a consumer instructing the agentic tool so the consumer would not know to prioritize this at the outset.

- A merchant could offer incentives (direct payments to the AI developer, data-sharing agreements, affiliate commissions) that cause an agent to favor that merchant's products or services over competitors' offerings—even when competitors offer better prices or quality.
- An agent might be designed to pursue macro-level objectives (*e.g.*, supporting a particular payment ecosystem's growth or adoption) that are favorable to the agent's creator in the long run but that may disadvantage the consumer in the short term.
- A financial services company could configure an agent to favor the company's own products (*e.g.*, credit cards, brokerage accounts) over competitors' products, even when consumers would be better served by the competitors' offerings.
- An agent configured to optimize for a single goal might inadvertently work against the consumer's broader interests. For example, an agent configured to find the lowest-cost insurance premium might not take into account the consumer's need for specific coverage levels, leading to underinsurance. An agent configured to maximize investment returns without adequate understanding of the consumer's risk tolerance might drive losses that the consumer cannot absorb. Or, an agent configured to lower credit card debt might not account for the consumer's need to maintain a certain credit utilization ratio for credit-building purposes.

This risk is especially acute because the consumer may not be aware that the agent is optimizing for interests other than their own. Attendees noted that even where disclosure occurs, many consumers lack the sophistication to evaluate whether the disclosed incentives truly affect their interests. Or, consumers may agree to compromises without understanding how they may be disadvantaged.

4.A.ii. CONSUMER PROVIDES INCORRECT INPUTS TO AN AGENTIC PAYMENT TOOL.

Agentic AI systems are, at their core, probabilistic models trained on historical data. They perform well in scenarios where their training data is representative and where the patterns they learned from remain stable. They perform poorly when consumers provide incomplete or ambiguous

instructions, when the scenarios they encounter differ from their training data, or when the consumer's expressed preferences do not fully specify or account for their true preferences.

This could be the combined fault of the consumer or the agent. While the agent might neglect or poorly explain a key detail, it is also possible the agent asked the consumer for the wrong information. A few examples of this risk are as follows:

- An agentic payment tool struggles to understand what is best for the consumer so the agent buys the wrong thing. This could be, for example, because the consumer failed to provide the agentic payment tool with a key detail. A consumer could instruct an agent to buy groceries but neglect to tell the agent how much freezer space the consumer has, and the agent buys more frozen food than the consumer can store.
- Prompts can be too complex for agents to understand. Relatedly, agents' limited understanding of context could prevent them from asking for additional information/prompts. For example:
 - A consumer tells an agent to “buy the best laptop for my needs,” but the agent has incomplete information about the consumer's use case and recommends a laptop that performs poorly for that use case.
 - A consumer instructs an agent to “book me a hotel room,” but does not tell the agent that the consumer needs a hotel near the conference center, so the agent reserves a room across town.
 - A consumer gives an agent context-dependent instructions (*e.g.*, “only use my credit card rewards points on purchases from my favorite brands”), but the agent lacks the context or capability to evaluate what brands the consumer favors.

As Symposium participants noted, consumers may expect that agents will ask clarifying questions when their instructions are ambiguous. But current agentic AI has limited capability to ask for clarification or to recognize when it lacks critical information. This misalignment between consumer expectations and agent capabilities could generate frustration and harm.

4.A.iii. HIGH TRANSACTION VOLUMES MAY INCREASE PRICES OR REDUCE AVAILABILITY.

If large numbers of agentic agents converge on the same product, price, or merchant at the same time, several adverse effects could occur. For example,

prices could increase due to sudden demand surges, and consumers could end up paying more than they would have if they were making the purchase manually from perhaps a range of different merchants. Availability of the product could be exhausted, leaving some consumers unable to complete purchases. Merchants could also experience sudden traffic spikes from bot behavior that could overwhelm their systems or lead them to block agent-initiated traffic. Timely availability of goods and services could be artificially limited.

4.A.iv. EXISTING PROBABILISTIC MODELS MAY BE MISMATCHED WITH CONSUMER EXPECTATIONS.

Today, consumers expect payments to work instantly and seamlessly and for fraud to be addressed behind the scenes. To manage payment volumes and money flows, providers use probabilistic models based on consumers' past behavior and other data. If agentic tools operate differently, existing models may fail, and consumer payments may not work as consumers want them to. For example, a bank's fraud tool may flag transactions initiated far from the consumer's home. But if agentic tools are operating from cloud-based servers in distinct locales, these fraud tools may not work.

4.B. Consumers Misunderstand Risks and Limitations of Agentic Payments.

Even if agents perform as intended, consumer misunderstanding of how agents work, what they are authorized to do, and what risks exist, could generate substantial harm and confusion.

4.B.i. CONSUMERS MAY LACK ADEQUATE UNDERSTANDING THAT THEY ARE USING AI, AN AGENTIC PAYMENT TOOL, OR HOW THEY WORK.

Many agentic payment products may be integrated into consumer apps, wallets, or financial platforms in ways that obscure the fact that an agentic tool is executing transactions to the consumer. Alternatively, consumers may know that an agent is involved but may lack understanding of how the agent works, what information it is using (or not using), what its limitations are, or how to override or control its decisions.

This lack of transparency could:

- Prevent consumers from understanding whether the agent is optimizing for their interests or for the interests of the agent's creator.
- Prevent consumers from identifying errors or problematic behavior early.

- Limit consumers' ability to make informed decisions about whether to use the agent or to modify the agent's instructions.

4.B.ii. INADEQUATE EDUCATION AND DISCLOSURE ABOUT AGENTS AND THE DUTY THEY OWE CONSUMERS.

At the outset, even when a consumer selects an agent to act on its behalf, the consumer may misunderstand the agent's ability to assist the consumer. This could be because the developer failed to explain this, the agent's limitations, or its biases. Some questions about consumer understanding include:

- Is the consumer aware that an AI agent is executing transactions on their behalf?
- Does the consumer understand the instructions they have given the agent?
- Is the agent optimizing for the consumer's interests or for the interests of other parties?
- What data does the agent use, and where does that data come from?
- What happens if the agent makes a mistake?
- How can the consumer monitor or override the agent's actions?
- What are the consumer's rights if the agent's actions cause harm?

Additionally, consumers may develop unrealistic expectations about whether and when an agent is acting as a fiduciary.²¹ Consumers may reasonably assume that an agent designed to optimize for the consumer's interests is a fiduciary, meaning that the agent's developer has legal obligations to the consumer. Attendees expressed concern that existing law may not clearly establish when agentic AI developers would be deemed fiduciaries or whether fiduciary duties (such as duties of disclosure, care, and loyalty) would apply.²²

4.C. AI Operational Failures and Security Risks.

Beyond consumer misunderstanding, several categories of operational failures and security risks could cause direct harm to consumers.

²¹ In the law, a fiduciary relationship is one where a party, the fiduciary, owes another party the duty of loyalty, care, and good faith. A fiduciary typically must always act in the best interests of the person to whom the duty is owed. *See* RESTATEMENT (THIRD) OF AGENCY § 8.01.

²² Parties need not identify a relationship as one involving a fiduciary for it to exist. *A. Gay Jensen Farms Co. v. Cargill, Inc.*, 309 N.W.2d 285 (Minn. 1981).

4.C.i. AGENT FAILS TO ACT, ACTS AT THE WRONG TIME, OR ACTS INCORRECTLY.

As has been widely reported, AI systems hallucinate and make other mistakes. In the context of agentic payments, attendees suggested these mistakes could include instances where:

- An agent is instructed to purchase a product when its price falls below a threshold, but the agent fails to monitor the price or executes the purchase late (after the promotion has ended).
- An agent is instructed to manage a consumer's credit card utilization to avoid negative credit score impacts, but because the agent doesn't fully understand how credit scoring works or does not have access to the full Metro 2 rules used by credit score companies, it takes actions that actually harm the consumer's credit.
- An agent is configured to rebalance investment portfolios but buys the wrong security due to a data error or misconfiguration.
- An agent buys too many items or the wrong item entirely. For example, a consumer says, "buy me five apples," and the agent, interpreting "five" as pounds rather than units, purchases 5 pounds of apples (approximately 20 apples). Or a consumer says, "buy me the new Nikes," the agent finds a counterfeit \$10,000 "Nike" shoes on a third-party marketplace and purchases them.
- An agent is configured to handle recurring bill payment but does not recognize when a monthly statement includes an unusual or suspicious charge and pays it anyway.
- Agents re-buy items the consumer does not need yet, or at all.
- Agents may not know how to interact with existing marketing, search, and consumer review mechanisms. Thus, for example, an agent could be easily tricked into believing marketing hype.
- An agent making purchasing decisions might not have adequate data about return policies, customer service availability, shipping terms, warranties, or other non-price factors that materially affect the value of a transaction. If the agent uses incomplete or outdated data, it could lead consumers to purchase from merchants with poor customer service or unfavorable return policies.

It is apparent how these operational failures could create significant consumer harm. This harm could generate a range of disputes involving all of the participants in the transaction. For example, if the consumer receives the wrong number of apples, the consumer may seek recourse from not only the provider of the agentic tool, but also the merchant (particularly if the

merchant and AI provider have partnered to offer the service jointly to the consumer) and demand the merchant take back the extra apples. If those efforts fail (and even if they do not), the consumer will likely contact their bank, which handled the payment, for redress, seeking reimbursement for the unwanted apples.

4.C.ii. AGENT UNINTENTIONALLY EXCEEDS THE SCOPE OF ITS AUTHORITY.

An AI agent might execute transactions that fall outside the scope of the consumer's authorization. Examples include:

- A consumer authorizes an agent to make purchases from a specific type of merchant, but the agent executes transactions on unrelated or malicious websites.
- A consumer authorizes an agent to manage household bill payments, but the agent is compromised and executes unauthorized transactions.
- An agent purchases goods or services from merchants on the dark web, executing transactions that the consumer did not authorize and that may violate law.

4.C.iii. FRAUDSTERS DEPLOY SPOOFED OR MALICIOUS AGENTS.

Criminals could create fake agents that purport to be legitimate agentic payment tools but that actually steal financial information or execute unauthorized transactions. Examples include:

- A fraudster creates a spoofed app that appears to be a legitimate agentic payment tool, but it harvests bank credentials and payment information.
- A fraudster creates a malicious agent that mimics the behavior of a legitimate agent, but it has been configured to route transactions to the fraudster's accounts or to execute transactions on behalf of the fraudster rather than the consumer.
- Fraudsters could use imposter agentic payment tools to buy goods from a fraudulent merchant also operated by the fraudster.

4.C.iv. AI AGENTS MAY TRICK FRAUD DETECTION SYSTEMS.

Agentic AI is sophisticated enough that it could be used to circumvent fraud detection and prevention systems. Today's fraud models are trained on consumer behavior and may be ill-equipped to handle AI-generated transactions and payments. This could manifest in a number of ways, including:

- An agentic system could distribute transactions across multiple accounts, merchants, and time periods in ways designed to evade velocity limits and fraud scoring systems.
- An agent could learn from feedback about which transactions are blocked and adapt its behavior to avoid triggering those blocks.
- An agent could use information about fraud detection rules to craft transactions that minimize detection risk.

4.C.v. PRIVACY AND DATA SECURITY CONCERNS.

For agents to function effectively, they require access to substantial amounts of consumer data, potentially including:

- Complete transaction history;
- Bank account details and balances;
- Investment portfolios and performance;
- Credit history and scores;
- Health and insurance information;
- Preferences, interests, and browsing history;
- Location data; or
- Biometric data (in some scenarios).

Each data point creates a potential vector for data security incidents, privacy breaches or data misuse. Specific risks include:

- Unauthorized access by third parties if the agent's infrastructure is compromised.
- Data monetization: The agent's creator or operator could sell or share consumer data with other parties without adequate consumer understanding or consent.
- Data retention: The agent's creator might retain data long after it is needed for authorized purposes.
- Cross-use of data: Data collected for one purpose (*e.g.*, optimizing payment decisions) might be used for other purposes (*e.g.*, behavioral advertising, credit scoring, insurance underwriting) without adequate disclosure or consent.

Data protection laws in the financial services context, such as the Gramm-Leach-Bliley Act and state privacy laws, provide some baseline protections, but many agentic payment applications may be created by non-financial services companies not covered by these regimes or below certain revenue or customer thresholds, creating regulatory gaps.

4.C.vi. BIAS AND DISCRIMINATION IN AI DECISION-MAKING.

AI systems can reflect and amplify biases present in their training data or embedded in their design. Potential harms include:

- An agent configured to optimize for creditworthiness based on historical lending data might incorporate historical discrimination or redlining practices, leading the agent to steer certain consumers away from credit opportunities.
- An agent might be biased against certain merchants or product categories if its training data over-represents certain demographic preferences.
- An agent might make purchasing decisions that reflect discriminatory patterns, such as recommending lower-quality or higher-priced products to certain consumers based on their demographic attributes.

Existing anti-discrimination law, such as the Fair Housing Act and the Equal Credit Opportunity Act and fair lending guidance, address discrimination by lenders and merchants, but the application of these frameworks to AI agents remains largely untested.

4.C.vii. ACCESSIBILITY AND FAIRNESS CONCERNS.

Agentic payment systems may not be equally accessible or beneficial to all consumers.

- **Language access issues:** AI agents may not provide adequate support for limited-English-proficiency consumers, potentially leading to misunderstandings or errors.
- **Disability access:** Consumers with disabilities (blind consumers unable to access visual interfaces, deaf consumers unable to access audio alerts) may struggle to use, monitor, or control agentic payment systems.
- **Socioeconomic disparities:** Consumers without savings, investments, or credit access may be unable to benefit from agents designed to optimize investment returns or manage credit strategically. Additionally, lower-income consumers using agentic payment systems

might face different pricing than higher-income consumers if algorithms incorporate socioeconomic proxies.

- **Algorithmic bias:** Agent decision-making could reflect and amplify historical discrimination or bias.

4.C.viii. PRODUCT DECLINE.

Much like web developers target search engines, merchants may develop products intended to satisfy agents rather than consumers themselves. This would likely prompt merchants to create sub-par products intended to “trick” agents. By shifting focus from the consumer, merchants also may not keep consumer protections, or the consumers best interests, front of mind given the focus on convincing agents to purchase their goods and services.

4.D. Potential Harms to Banks and Financial Institutions.

Agentic payments also create risks for the financial institutions that support payment systems. Symposium attendees suggested the following risks:

4.D.i. LIABILITY FOR INCREASED NUMBERS OF UNAUTHORIZED TRANSACTIONS.

Under the EFTA (Regulation E) and TILA (Regulation Z), financial institutions are liable for certain categories of unauthorized transactions, with limited exceptions. A key exception, discussed below in Section 5.B, exists in some circumstances when a consumer voluntarily provides an access device to another person. But the applicability of this exception to agentic AI is unclear, at best. Even if some disputes fall under these exemptions, many others may not, and banks may face an increased number of disputes where they must reimburse their customers.

4.D.ii. HIGH VOLUMES OF DISPUTES AND CHARGEBACKS.

If agentic payments scale substantially, the volume of consumer disputes and chargebacks could overwhelm existing systems. This is particularly true if agents make mistakes that generate widespread disputes (*e.g.*, all instances of an agent purchasing wrong products). This could also occur if fraud increases significantly or if merchants or agents use dispute/chargeback systems as a primary dispute resolution mechanism rather than merchant customer service.

Current chargeback and dispute resolution systems were designed for human-initiated transactions and small volumes. They may lack the operational capacity to handle orders of magnitude increasing in volume such as those we will likely see with agentic payment tools.

4.D.iii. LIABILITY FOR GOODS NOT DELIVERED OR DELIVERED DAMAGED.

For credit card transactions, Regulation Z protections apply when a consumer purchases goods or services that are not delivered or are delivered damaged. Under these protections, consumers can dispute the charge and force the merchant or issuer to resolve the issue. With agentic payments, disputes over goods not delivered or damaged may increase substantially if:

- Agents make purchasing decisions without human review, leading to more transactions susceptible to merchant fraud or simply merchants unable to fulfill large influxes of unexpected orders.
- Transaction volumes increase, increasing the absolute number of failed transactions.
- Merchant accountability systems break down if merchants cannot identify individual consumers behind agent-initiated transactions (for example, if thousands of agents initiated transactions through a proxy wallet).

4.D.iv. CONSUMER CONFUSION MAY INCREASE CUSTOMER SERVICE BURDENS ON BANKS AND INCREASE REPUTATIONAL RISKS.

Consumers who are confused about why a transaction occurred, whether it was authorized, or what happened when the transaction went wrong, will likely turn to their banks for help. Customer service costs could increase substantially. Additionally, consumers who cannot get adequate help from banks may develop negative perceptions of those banks, creating reputational and attrition risk.

4.D.v. BANKS MAY STRUGGLE TO MAINTAIN ADEQUATE VISIBILITY INTO AGENT BEHAVIOR.

If agents are developed by third parties and operate through merchant APIs or open banking infrastructure, banks may have limited visibility into agent behavior. This could make it difficult for banks to:

- Detect fraudulent agents or compromised legitimate agents;
- Identify patterns of agent failure or abuse;
- Explain to consumers what the agent did or why;
- Resolve disputes effectively; or
- Manage operational and reputational risk.

4.E. Potential Harms to Merchants and the Broader Ecosystem.

Several externalities—including several discussed above—could negatively impact merchants and other market participants. As one Symposium attendee suggested, merchant harms redound to consumer harms. And, merchants could be incentivized to harm the system for the longer term in exchange for the immediate benefit of increased sales. Other potential harms have larger impacts.

4.E.i. LIABILITY SHIFT TO MERCHANTS.

If agentic payments generate disputes, merchants could face liability for fraud by AI agents, especially if merchants are contractually responsible for validating transactions and preventing fraud. Merchants might face chargeback liability even if the consumer authorized the agent to make purchases, because chargebacks are typically liability-neutral for issuers but generate significant costs for merchants.

4.E.ii. HIGH VOLUMES OF BOT TRAFFIC INCREASES MERCHANT COSTS.

Agentic agents could generate large volumes of traffic on merchant websites as they search for products, monitor prices, and execute transactions. This traffic could:

- Overrun merchant infrastructure, causing slowdowns or outages.
- Trigger merchant costs for hosting, processing, and infrastructure even for agents that do not ultimately purchase products.
- Trigger merchant fraud and bot detection systems, potentially blocking legitimate agent traffic. As one Symposium participant noted: “The cost of running these systems is high. If there are a million bots hitting your website, you are paying for that traffic. And, if the bots are not buying anything, that costs you money.”

If agentic payment tools create a surge of business for an unexpecting small business, that business’s inability to fulfill those orders could have significant long-term effects on the business.²³

4.E.iii. DETERIORATION OF REVIEW SYSTEMS AND CONSUMER TRUST MECHANISMS.

If agentic agents bypass the normal consumer shopping experience (reading reviews, comparing products, evaluating merchant feedback), the incentives

²³ For example, in 2022, a condiments business run by a small-scale creator—exploded in popularity after going viral on TikTok, leading to massive order influxes that its limited operation could not handle. This caused inconsistent product batches, labeling errors, shipping delays, and customer service backlogs, eroding trust and sparking complaints. See Miles Klee, *How the ‘Pink Sauce’ Chef Survived Her Downfall*, ROLLING STONE (Jan. 25, 2023), <https://www.rollingstone.com/culture/culture-features/pink-sauce-tiktok-chef-cancelation-walmart-1234667556/>.

for consumers to leave reviews or for merchants to maintain high quality standards could erode. This could reduce transparency and increase the potential for fraud or misrepresentation about products and their attributes.²⁴

4.E.iv. **BROADER MARKET DYNAMICS AND ANTITRUST CONCERNS.**

Several longer-term ecosystem risks also emerged from Symposium discussion:

- **Market manipulation:** If agents begin to do the same or similar things, that could have strange impacts. For example, imagine the impacts of an agentic payment tool that buys everyone the same sneakers or retirement planning tools that make the same investment decisions for everyone.
- **Market consolidation:** Entities with access to large consumer data sets that have relationships with merchants and payment networks or control over agentic AI platforms could use agentic payments to cement their market dominance pushing out or limiting the success of other platforms.
- **Antitrust risk:** If a platform controls both the consumer agent and merchant infrastructure (*e.g.*, a payments network that develops both consumer agents and merchant acceptance systems), that platform could abuse its control to favor certain merchants or payment methods.
- **Smaller merchant displacement:** Agentic agents that prioritize price and efficiency might systematically favor large merchants at the expense of smaller competitors, leading to further market consolidation.
- **Trade balance implications:** If agentic agents become sophisticated at sourcing products and services globally, this could have implications for consumer spending patterns and trade balances.
- **Decline of physical stores and related impacts.** If bots are buying things, consumers may not need to enter stores to test items, try on clothes, etc. furthering the consumer behavior and patterns we are seeing now and the subsequent impacts it is having on small businesses, real estate, etc.

²⁴ See Rachit Kamdar & Siva Viswanathan, *Impact of AI on Reviews and Outcomes* (2025), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5329033.

Table III: Examples of Vulnerabilities by Stakeholder

Risk Category	Consumer Impact	Bank/Issuer Impact	Merchant Impact
Misaligned Incentives	Agent buys goods that benefit developer (kickbacks), not user.	Reputational damage when customers complain.	Loss of trust if perceived as “bribing” agents.
Operational Failure	Agent buys wrong item or quantity (e.g., 500 vs 5 apples).	Increased volume of disputes and chargebacks.	Surge in returns; inventory management issues.
Data & Privacy	Excessive data harvesting; potential breaches.	Difficulty managing Third-Party Risk (TPRM).	Loss of direct relationship with the customer.
Legal / Liability	Financial loss if “unauthorized” liability caps don’t apply.	Operational burden of proving “authorization.” TPRM	Potential liability shift for bot-driven fraud.
Market Stability	Exclusion from markets if user lacks AI tools.	Loss of visibility into transaction behavior.	Website crashes due to high-volume bot traffic.

5. Existing Legal and Regulatory Frameworks: Coverage and Gaps

The next phase of the Symposium identified potentially applicable legal and regulatory regimes that may apply to some of the potential agentic payment harms discussed above and the gaps in those regimes.

5.A. Common Law Principals.

Agency. Given the focus on artificial agents, a key area of law is agency liability. While agency law varies to some extent across the states, most states have adopted similar definitions of principals and agents through adoption of the Restatement (Third) of Agency. During the Symposium, some attendees debated whether traditional agency principles apply to nonhuman agents such as agentic tools. The Restatement defines an agency as:

the fiduciary relationship that arises when one person (a principal) manifests assent to another person (an agent) that the agent shall act on the principal's behalf and subject to the principal's control, and the agent manifests assent or otherwise consents so to act.²⁵

Notably, this definition assumes the agent is an individual or one or more individuals acting through a business association capable of understanding and exercising duties of loyalty, care, and disclosure to the principal. An immediate question then is whether and when an agentic payment tool is an agent under the Restatement. If it is, then the consumer, the principal, can avail themselves of the protections afforded to principals. While the answer to that question is beyond the scope of this White Paper, it could have profound effects on the development of agentic payment tools. For example, if an agentic tool is not bound by the divided loyalty rules described below, agentic tools may not be inclined to obtain a consumer's consent to proceed while conflicted. Other questions exist about the scope of agents' liability under these rules and what prerequisites are necessary to establish a principal/agent relationship under the law. Agents owe a number of duties to their principal, including the duty of loyalty, obedience, disclosure, care, and confidentiality. The scope of these duties often vary based on the agent's assigned tasks.²⁶

²⁵ RESTATEMENT (THIRD) AGENCY § 1.01.

²⁶ For example, in Washington, agents must "(a) act in accordance with the principal's reasonable expectations to the extent actually known by the agent and, otherwise, in the principal's best interest; (b) Act in good faith; and (c) Act only within the scope of authority granted in the power of attorney." Agents must also, with some exceptions, "(a) Act loyally for the principal's benefit; (b) Act so as not to create a conflict of interest that impairs the agent's ability to act

A key issue in agency law is the treatment of agents who have divided loyalties. The Restatement generally prohibits agents from engaging in self-dealing and when there are conflicts of interests.²⁷ Typically, divided loyalties must be disclosed to the principal who then must provide informed and voluntary consent to the agent.²⁸ For example, if a consumer wants to buy a home and retains a real estate agent, if the agent is going to be compensated by someone other than the homebuyer, the agent must disclose that fact to the homebuyer.²⁹ In most jurisdictions, if the homebuyer does not consent, the agent cannot proceed while keeping their divided loyalties.

Product Liability and Tort. Tort law governs civil wrongs, such as intentional acts and negligent acts where one party causes harm to another's person, property, or rights regardless of whether the parties have a contractual relationship between them. Generally, tort law allows an injured party to seek remedies like monetary damages. Product liability, a subset of tort law, holds manufacturers, distributors, and sellers strictly liable for injuries caused by defective products that are unreasonably dangerous, without requiring proof of negligence.

Like any consumer product, Symposium attendees suggested that consumers may be able to pursue claims against the providers of the agentic tool if it causes harm to the consumer. Of course, costs in time and potential legal fees required to vindicate those rights may be substantial. In addition, providers of most online tools use arbitration clauses that restrict consumers' ability to bring class actions on behalf of other potentially injured customers to vindicate those rights en masse. Given that, Symposium attendees were skeptical of consumers' ability to assert smaller consumer tort and product liability claims against providers of agentic tools in significant numbers.

Contract. Contracts between consumers and providers may give consumers certain additional rights and will usually be specified in terms of service. However, consumers typically have limited bargaining power and rarely read these contracts. Attendees suggested consumers would be unable to extract any significant protections via contract.

impartially in the principal's best interest; (c) Act with the care, competence, and diligence ordinarily exercised by agents in similar circumstances...." WASH. REV. CODE § 11.125.140,

²⁷ RESTATEMENT (THIRD) AGENCY § 8.01.

²⁸ RESTATEMENT (THIRD) AGENCY § 8.06.

²⁹ See, e.g., CAL. BUS. & PROF'L CODE § 10176(d).

5.B. Consumer Finance Protection Laws: EFTA, TILA, and Related Frameworks.

Consumers routinely rely on rights afforded to them under the federal consumer protection laws when they experience problems with their bank accounts, credit and debit cards, and other financial transactions. This section reviews the key laws and their potential shortcomings in the context of agentic payments.

EFTA and Regulation E (Electronic Funds Transfers).

EFTA, enacted in 1978, and its implementing regulation (Regulation E), establish substantial protections for consumers who use electronic payment methods. Key protections relevant here include the consumer liability cap, burden of proof, and the access device exception.

A consumer's liability for unauthorized electronic funds transfers is capped at \$50 if the consumer reports the unauthorized transfer within two days of discovering it, and at \$500 if the consumer misses that deadline but reports it within 60 days. If the consumer does not report unauthorized transfers within 60 days, the consumer's liability is uncapped.³⁰ The burden falls on the financial institution to prove that a transfer was authorized; if the institution cannot establish authorization, the consumer is not liable.³¹ An important exception to the general liability cap is that a consumer's liability for unauthorized transfers initiated through a lost or stolen access device may not be limited if the consumer gave their access device to another person and that person exceeded the scope of their authority.³² This exception does not apply if the consumer received a benefit from the transaction, notified the financial institution that transfers are no longer authorized, or if the third-party obtained the consumer's access device through fraud or robbery.³³

Key issue for agentic payments: This “access device” exception raises the core question about authorization in agentic scenarios. If a consumer authorizes an AI agent to access the consumer's bank account and initiate transactions, has the consumer “provided” an access device that triggers the exception? What if the agent exceeds its authorization? What if the agent is compromised and subsequently makes unauthorized transactions? If the bank's liability is limited in these circumstances, that could have substantial impacts on consumers who would incur losses from such transactions instead of their bank. Relatedly, questions exist about the application of Regulation E to payments conducted in

³⁰ See 12 C.F.R. § 1005.7.

³¹ 15 U.S.C. § 1693g(b).

³² 12 C.F.R. § 1005.2(m).

³³ 12 C.F.R. § 1005.2 cmt. 2(m)-3. The exception also does not apply if the third-party uses force to induce the consumer to provide their access device. *Id.* cmt. 2(m)-4.

cryptocurrency. Given the increased potential of alternative payment rails, the likelihood of agentic payment tools using payment methods where Regulation E does not apply is substantial.

TILA and Regulation Z (Truth in Lending).

TILA, and its implementing regulation (Regulation Z), provides similar protections for consumers when they purchase goods or services with their credit cards. For example, Regulation Z limits consumers' liability for unauthorized transactions to the lesser of \$50 or the amount of the unauthorized transaction prior to the consumer's notification to their card issuer.³⁴ This exception does not apply if the transaction is made by someone who has "actual, implied, or apparent authority to use the consumer's credit card."³⁵ Similar to Regulation E, the burden in Regulation Z is on the credit card issuer to prove authorization.³⁶

Key issue for agentic payments: As with Regulation E, it is, at best, unclear whether a consumer's authorization of an AI agent to initiate purchases constitutes "actual, implied, or apparent authority to use" the consumer's credit card. If an agent exceeds its authorization, or if an agent is compromised, how do these protections apply? The resolution of this question when applied to agentic payments is key to evaluating consumers' potential liability for use of these new tools.

Dodd-Frank § 1033 and Open Banking.

The Dodd-Frank Act's Section 1033 consumer data rights provision mandates that the CFPB establish rules allowing consumers to access their financial data and to authorize third parties to access that data. As discussed above, the CFPB was, and still is, considering revisions to those rules, which may be finalized in the near future. Symposium participants acknowledged that data sharing, and related liability, is beneficial when looking at agentic tools and really goes to the core of agentic tools working at their highest functionality. The agentic tools need to be equipped with the correct data and an adequate amount of data to work efficiently and precisely.

Implementing regulations (proposed by the CFPB and still under development in 2026) could enable consumers to authorize agentic AI systems to access and utilize their financial data for decision-making purposes. However, banks fear

³⁴ 12 C.F.R. § 1026.12(b)(1)(ii).

³⁵ 12 C.F.R. § 1026.12(b)(1). See *Drainman v. Am. Express Travel Related Servs. Co.*, 892 F. Supp. 1096, 1098 (N.D. Ill. 1995); *Citibank (South Dakota), N.A. v. Gifelman*, 773 A.2d 993, 997 (Conn. App. Ct. 2001) (cardholder who received a \$25 per month stipend for permitting a third-party to be a secondary user on his account could not claim unauthorized use, even if the unauthorized use may have been committed by someone other than the secondary user, because the cardholder benefited from the use).

³⁶ 15 U.S.C. § 1643(b).

that unfettered access to consumer data could substantially increase consumer risk and banks' liability.³⁷

Key issue for agentic payments: Section 1033 is an important enabler of agentic payments, as agents will require access to consumer financial data to function effectively. However, Section 1033 itself does not address liability, authorization, consent, or consumer protection issues specific to agentic payments. Attendees expressed concern that the unknown future of the CFPB's 1033 rulemaking could impact the expansion and development of agentic payments.

UDAP / UDAAP.

The FTC Act prohibits most non-bank entities from engaging in unfair and deceptive acts and practices.³⁸ The Dodd-Frank Act additionally prohibits banks with over \$10 billion in assets and nonbank providers of consumer financial services from engaging in unfair, deceptive, and abusive acts and practices.³⁹ While a full overview of the conduct prohibited by these statutes is beyond the scope of this White Paper, Symposium attendees repeatedly noted that AI-related conduct that causes consumers harm could lead to a cognizable UDAP / UDAAP claim. In 2023, the CFPB, FTC, Department of Justice, and the Equal Opportunity Employment Commission jointly issued guidance indicating that the use of automated systems—including agentic tools—could violate UDAAP laws.⁴⁰ Then-CFPB Director Rohit Chopra noted that “[c]ompanies must take responsibility for their use of these tools.”⁴¹

Key issue for agentic payments: At the Symposium, attendees frequently suggested that many of the harms discussed could give rise to a UDAP/UDAAP claim. For example, if an agentic payments agent asserts it is in fact the consumer's agent but has undisclosed loyalties to a third-party, that failure to provide disclosure could be viewed as deception. Similarly, if there are undisclosed limitations on an agentic payments tool—for example it only can buy from certain merchants or buys items that cause consumers to incur additional undisclosed fees—that conduct could be viewed as unfair.

³⁷ See Sean Oblack, *Banks Challenge CFPB Rule Jeopardizing Security and Privacy of Consumer Data*, BANK POL'Y INST. (Oct. 22, 2024), <https://bpi.com/banks-challenge-cfpb-rule-jeopardizing-security-and-privacy-of-consumer-financial-data/>.

³⁸ 15 U.S.C. § 45.

³⁹ 12 U.S.C. § 5536.

⁴⁰ FED. TRADE COMM'N ET AL., JOINT STATEMENT ON ENFORCEMENT EFFORTS AGAINST DISCRIMINATION AND BIAS IN AUTOMATED SYSTEMS (2023), https://www.ftc.gov/system/files/ftc_gov/pdf/EEOC-CRT-FTC-CFPB-AI-Joint-Statement%28final%29.pdf.

⁴¹ Director Rohit Chopra, Consumer Fin. Prot. Bureau, Prepared Remarks on the Interagency Enforcement Policy Statement on “Artificial Intelligence” (Apr. 25, 2023), <https://www.consumerfinance.gov/about-us/newsroom/director-chopra-prepared-remarks-on-interagency-enforcement-policy-statement-artificial-intelligence/>.

Model Risk Management Regulatory Requirements.

Symposium attendees, particularly those representing banks, emphasized the role of MRM in managing banks' handling of newly introduced risks from agentic payments. In financial services, MRM addresses risks from predictive models used for credit scoring, fraud detection, and other market risks. MRM also helps institutions avoid penalties and build trust. Prudential regulators emphasize robust development, validation, and oversight to limit adverse outcomes with AI.⁴²

Key issue for agentic payments: As AI models proliferate in fintech, MRM will need to evolve to address newly identified risks. However, not all industry participants must concern themselves with MRM so those companies may not have the same risk protections.

E-SIGN Act and State UETA Laws.

Some Symposium attendees referenced the role of electronic disclosure and consent in a world where agents may be receiving disclosures and providing authorizations and consumers may not themselves be performing these actions. Currently, the E-SIGN Act requires those desiring to provide electronic disclosures to consumers, such as banks providing loan disclosures, to satisfy certain prerequisites before proceeding.⁴³ These prerequisites anticipate a consumer will receive and review the E-SIGN disclosures and any subsequent disclosure. The same is true about the electronic signature of documents.

Most states, except New York, have adopted the Uniform Electronic Transaction Act. Similar in many respects to the E-SIGN Act, these UETA laws have provisions specific to the use of agents. For example, the UETA acknowledges that in an automated transaction, “[a] contract may be formed by the interaction of electronic agents of the parties, even if no individual was aware of or reviewed the electronic agents’ actions or the resulting terms and agreements.”⁴⁴ While this long precedes the development of agentic payment tools, this language may prove useful to those trying to establish contractual relationships without direct consumer involvement.

⁴² See, e.g., Michael Hsu, Acting Comptroller of the Currency, Office of the Comptroller of the Currency, 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>.

⁴³ 15 U.S.C. § 7001.

⁴⁴ UNIFORM ELECTRONIC TRANSACTIONS ACT, NAT’L CONF. OF COMM’RS ON UNIF. STATE LAWS 37 (1999), <https://www.uniformlaws.org/viewdocument/final-act-21?CommunityKey=2c04b76c-2b7d-4399-977e-d5876ba7e034>. An “electronic agent” is defined as “a computer program or an electronic or other automated means used independently to initiate an action or respond to electronic records or performances in whole or in part, without review or action by an individual.” *Id.* at 4.

Key issue for agentic payments: Attendees raised questions about whether and how agentic tools can process and review disclosures for consumers. Some suggested agentic tools may do a better job of analyzing and assessing complex and dense disclosures. Others said that consumers may be at a disadvantage if disclosures are delivered only to agentic tools and not consumers themselves. This raises compliance concerns as well, because many laws require disclosures be delivered to the consumer. These laws do not address whether a provider satisfies its legal obligation by delivering a disclosure only to an agent. Similarly, where federal or state law requires consumer consent, attendees speculated about whether an agent can provide that consent on behalf of a consumer.

Third-Party Risk Management.

Attendees noted that banks and non-bank providers of consumer financial services are obligated to manage risks posed by third parties (*i.e.*, vendors and other partners) with whom they work with to offer services to consumers. Federal regulators have promulgated substantial guidance on the steps these entities must take to assess the risks posed by third parties. In 2023, the federal banking agencies issued final Interagency Guidance on Third-Party Relationships: Risk Management.⁴⁵ This guidance outlines the risk-based framework banks are expected to use to manage third-party risks across the relationship life cycle. It emphasizes that banks retain full responsibility for safe and sound operations and legal compliance despite outsourcing.⁴⁶ Similarly, the CFPB issued guidance in 2012, updated in 2016, regarding how covered entities must exercise oversight of their third-party relationships.⁴⁷

Key issue for agentic payments: Attendees noted that these oversight systems would have to be enhanced to address new issues related to agentic payments. For example, existing contractual relationships may not have been developed to address new risks posed by AI. Banks and covered non-banks will also have to enhance existing due diligence processes to better assess risks posed by vendors' use of agentic payments tools. Some attendees suggested that banks and others should not be liable for AI-related incidents if they were unaware of a third party's use of agentic tools. Others noted that while contracts make vendors

⁴⁵ Interagency Guidance on Third-Party Relationships: Risk Management, 88 Fed. Reg. 37920 (June 9, 2023).

⁴⁶ *Id.* at 37924.

⁴⁷ Press Release, Consumer Fin. Prot. Bureau, CFPB to Hold Financial Institutions and their Service Providers Accountable (Apr. 13, 2012), <https://www.consumerfinance.gov/about-us/newsroom/consumer-financial-protection-bureau-to-hold-financial-institutions-and-their-service-providers-accountable/>; CONSUMER FIN. PROT. BUREAU, CFPB BULLETIN 2012-03, SERVICE PROVIDERS (2012), https://files.consumerfinance.gov/f/documents/201204_cfpb_bulletin_service-providers.pdf; CONSUMER FIN. PROT. BUREAU, COMPLIANCE BULLETIN AND POLICY GUIDANCE; 2016-02, SERVICE PROVIDERS (2016), <https://www.consumerfinance.gov/compliance/supervisory-guidance/compliance-bulletin-and-policy-guidance-2016-02-service-providers/>.

liable for risks they introduce, these vendors may be ill-equipped to manage the increased risk posed by consumers' use of agentic payments tools.

5.C. Private Network Rules.

As is discussed throughout this White Paper, almost all non-cash consumer payments in the United States occur over payment card networks or ACH networks. Each of these networks has its own rules.

The major card networks – American Express, Discover, Mastercard, and Visa – each have a set of rules governing all network participants.⁴⁸ These rules consist of core operating regulations that govern card issuance, transaction processing, authorization, clearing, settlement, and compliance for all network participants, including issuers, acquirers, and merchants. These rules mandate responsibilities for fraud prevention, chargeback and dispute resolution procedures, security standards like data protection, and adherence to local laws. They also detail how violations are punished and how disputes between members are resolved.⁴⁹ Rules are regularly updated to address emerging trends.⁵⁰

These rules are as old as their respective networks. Mastercard established rules in 1966 via the Interbank Card Association, which governed card authorization, settlement, and security details.⁵¹ Early rules for the BankAmericard program mirrored the Interbank Card Association Rules.⁵² These rules became standardized in the 1970s when BankAmericard became Visa. Over the subsequent decades, these rules have been revised and expanded to reflect the growth and change of consumer payments. For example, they adapted to the evolution of debit cards in the 1990s, the Fair Credit Billing Act in 1974 (which amended TILA to add credit card-specific consumer protections), and, more recently, the mobile and touchless payments consumers make with their mobile devices.

⁴⁸ See, e.g., MASTERCARD, MASTERCARD RULES (2025), <https://www.mastercard.com/content/dam/mccom/shared/business/support/rules-pdfs/mastercard-rules.pdf>; VISA, VISA CORE RULES AND VISA PRODUCT AND SERVICE RULES (2025), <https://usa.visa.com/dam/VCOM/download/about-visa/visa-rules-public.pdf>.

⁴⁹ See, e.g., MASTERCARD, MASTERCARD RULES 51-52, 219, (2025), <https://www.mastercard.com/content/dam/mccom/shared/business/support/rules-pdfs/mastercard-rules.pdf>

⁵⁰ E.g., Lynne Marek, *Visa ramps new fraud prevention program*, PAYMENTS DIVE (Sept. 2, 2025), <https://www.paymentsdive.com/news/visa-new-merchant-card-acceptance-fraud-dispute-program/758988/>.

⁵¹ *Brand History*, MASTERCARD, <https://www.mastercard.com/brandcenter/us/en/brand-history.html> (last visited Dec. 9, 2025).

⁵² David L. Stearns, “Think of it as Money”: A History of the VISA Payment System, 1970-1984 (Jan. 2007) (Ph.D. dissertation, University of Edinburgh), <https://era.ed.ac.uk/handle/1842/2672> (“The [original BankAmericard] operating regulations ... are essentially the rules of the game. They stipulate not only card design standards and rules for how the marks can be used in advertising, they also dictate how the various inter-member work should be accomplished and what penalties a member suffers if that member violates the rules. The operating regulations also stipulate the basic rules an acquirer must enforce in their contracts with merchants. They are the central coordination mechanism for the entire system, providing a structure in which the competing members can cooperate.”) (citations omitted).

Today, these rules explain how to resolve a range of potential consumer issues including fraud, lost and stolen cards, processing errors, merchants' failure to deliver goods or services, counterfeit merchandise, cancelled transactions, and others.⁵³ These rules also contain zero-liability policies that protect consumers from unauthorized transactions and fraud.⁵⁴ These protections go beyond the requirements of EFTA and TILA discussed in the prior section. The operators of these networks decided these additional pro-consumer protections were necessary to encourage consumer adoption and use of their card products and to overcome consumers' concerns about potential exposure and risk of such card products.⁵⁵

The NACHA Rules govern ACH payments. Development of these rules in the late 1960s led to the creation of NACHA in 1974.⁵⁶ While NACHA does not operate ACH networks, it does setup and maintain the rules. Like card network rules, the NACHA rules address all aspects of the processing of ACH transactions including consumer authorization rules, returns of transactions and unauthorized transactions, and network disputes. Also like the card network rules, the NACHA rules have evolved to address innovations such as same day ACH.⁵⁷

Key issue for agentic payments: Attendees emphasized that these networks are not designed for the anticipated growth in agentic payments. More fundamentally, many attendees emphasized that new networks may develop, such as X402 or others, which may not include these same consumer protections.

5.D. Potential Concerns with Existing Structures.

During the Symposium, attendees addressed potential problems with the legal structures identified above. Several attendees noted that, at the outset at least, the application of existing laws to new technologies may be unclear. This could

⁵³ See generally VISA, VISA CORE RULES AND VISA PRODUCT AND SERVICE RULES 674-803 (2025), <https://usa.visa.com/dam/VCOM/download/about-visa/visa-rules-public.pdf>; MASTERCARD, CHARGEBACK GUIDE: MERCHANT EDITION (2025), <https://www.mastercard.us/content/dam/public/mastercardcom/na/global-site/documents/chargeback-guide.pdf>.

⁵⁴ *Our Zero Liability Policy lets you shop without worry*, VISA, <https://usa.visa.com/pay-with-visa/visa-chip-technology-consumers/zero-liability-policy.html> (last visited Dec. 9, 2025); *Zero liability protection*, MASTERCARD, <https://www.mastercard.com/us/en/personal/protection-and-security/zero-liability-protection.html> (last visited Dec. 9, 2025); *What is the American Express® Fraud Protection Guarantee?*, AM. EXPRESS, <https://www.americanexpress.com/ca/en/customer-service/faq.what-is-the-american-express-fraud-protection-guarantee.html> (last visited Dec. 9, 2025); *Explore your free benefits and unlimited rewards*, DISCOVER, <https://www.discover.com/credit-cards/member-benefits/> (last visited Dec. 9, 2025).

⁵⁵ Some have argued that the lack of similar protections for crypto transactions has slowed the adoption of crypto for everyday consumer payments. Luke Kowalski et al., *Lackluster Adoption of Cryptocurrencies as a Consumer Payment Method in the United States – Hypothesis: Is This Independent Technology in Need of a Brand, and What Kind?*, 16 J. RISK AND FIN. MGMT. 23, 27 (2022), <https://www.mdpi.com/1911-8074/16/1/23>; Venkata Marella et al., *Understanding the creation of trust in cryptocurrencies: the case of Bitcoin*, 30 ELEC. MKTS. 259, 260 (2020), <https://link.springer.com/article/10.1007/s12525-019-00392-5>.

⁵⁶ *History of Nacha and the ACH Network*, NACHA (Apr. 20, 2019), <https://www.nacha.org/content/history-nacha-and-ach-network>.

⁵⁷ *Id.*

lead to extensive and costly litigation as parties work through these issues. One attendee noted that this could take some time because many cases may be brought with “bad facts,” which, inevitably may lead to bad law.

Attendees discussed whether and how existing law might apply to AI agents. For example, while human agents owe certain duties to their principals, will AI agents owe those same duties? This is particularly relevant if a consumer creates their own consumer agent by themselves, and the AI agent is not provided by a third-party entity that the consumer could pursue a claim against. Or can product liability law support claims against AI bots? Under typical product liability law, a consumer can bring claims against providers of defective products. If a consumer trains an AI tool but then has an issue with how the AI tool acted, is product liability law the right mechanism to resolve those claims? The answer is unclear.

Attendees also noted that in disputes against merchants and other providers, arbitration clauses may block significant relief in tort, contract, or agency law.

Table IV: Regulatory Gap Analysis

Regulatory Domain	Intention of Developers	Framework Today	Agentic AI “Gap”
Model Risk Management (MRM)	Bank-developed predictive models used in regulated financial activities	Strict validation: Banks must validate, test, and monitor models (e.g., credit, fraud) under supervisory guidance and examiner review.	The “non-bank” loophole: Non-bank AI developers and many fintechs are generally not subject to MRM requirements, creating a safety disparity between bank-run and tech-run agents.
Supervisory Examination	Prudentially regulated financial institutions	Proactive oversight: Banks undergo regular, onsite examinations by prudential regulators (i.e., OCC, FDIC, Federal Reserve) to identify risks before consumer harm occurs.	Reactive enforcement: AI developers and merchants are generally not subject to routine examinations and are overseen primarily through after-the-fact enforcement (e.g., FTC/state AGs).
Third-Party Risk Management (TPRM)	Traditional vendor relationships with observable systems and controls	Mandated due diligence: Banks must vet, monitor, and remain liable for vendors under interagency guidance on third-party relationships.	Visibility blind spots: Even diligent banks may lack technical visibility into “black box” AI agents, making it difficult to detect compromise, drift, or conflicts in real time.

Regulatory Domain	Intention of Developers	Framework Today	Agentic AI “Gap”
EFTA (Reg E) & TILA (Reg Z)	Consumer-initiated electronic payments and credit card transactions	Liability caps: Consumer liability is limited for unauthorized transfers, with burden on the financial institution to prove authorization.	The “access device” trap: Sharing credentials or API keys with an AI agent may be deemed furnishing an access device, potentially shifting full liability to the consumer if the AI agent exceeds its authority.
Agency Law	Human agents acting on behalf of principals	Fiduciary duty: Human agents owe duties of loyalty, care, and disclosure, prohibiting undisclosed conflicts or kickbacks.	The “software” exception: It is legally unclear whether AI tools qualify as agents. If not, they may owe no fiduciary duty and can legally prioritize third-party incentives over consumer interests.
Unfair / Deceptive Acts (UDAP / UDAAP)	Human-directed business practices and representations	Prohibits deception: Businesses may not mislead consumers or omit material information about products or services.	Opacity barrier: Proving deception is difficult when AI decision-making is opaque; regulators may lack access to code, training data, or explainability needed to establish unfair or deceptive conduct.
Disclosure Laws (E-SIGN / UETA)	Human receipt, review, and consent to electronic disclosures	Consumer consent: Requires disclosures be delivered to, and consented to by, consumers before electronic contracting.	The “robo-reader” gap: Laws do not clearly address whether an AI agent can receive disclosures or consent on a consumer’s behalf, raising questions about contract validity.

6. Role of Regulation in the Context of AI

Symposium attendees addressed the role of regulations in AI and agentic payments. Debate was robust and some attendees took issue with the question of whether regulation was the right tool to address agentic issues. An attendee explained that, in his view, existing regulations will not prevent consumers from using AI tools. Nor is the uptake of AI impacted by the regulatory landscape. This is because few regulations contain absolute prohibitions.

With respect to the role of regulations in innovation, attendees had differing opinions. Some asserted that regulators must think carefully about this issue and ensure there is a balance between regulating AI tools and allowing innovation to flourish. As an example, one attendee compared regulators' permissible approach to prize linked savings in the United Kingdom with a more restrictive approach in the United States; such programs have floundered here while they are more successful in the U.K. Other attendees noted that regulators should avoid giving definitive answers right now. Such regulations, comments, or guidance in these early days of AI may make for permanent answers on key questions that will stifle innovation.

Finally, attendees speculated that regulations should be outcome-oriented, that regulators should consider small entity carveouts (as opposed to regulatory sandboxes), and that regulators may not be able to induce optimal behavior by shifting liability. Regulators were also encouraged to consider whether, and to what extent, consumers may be able to resolve some issues directly with agents.

7. Potential Solutions Identified During Symposium

The Symposium encouraged and provided an opportunity for attendees to brainstorm potential solutions to the identified risks of agentic payments. Below are some of the primary solutions suggested by the attendees:

7.A. Market-Based Solution, For Now.

Several attendees urged that it is too early in the development of agentic payments for legislators and regulators to adopt new laws or change existing laws and rules to address hypothetical agentic risks. These attendees explained that AI and agentic tools have great promise, and any regulatory efforts could stifle innovation.⁵⁸ Many have urged this approach in connection with other recent technologies, such as cryptocurrency. While there has been back-and-forth with respect to the Securities and Exchange Commission and the Commodity Futures Trading Commission about the proper approach to regulating crypto, the consumer-facing regulators such as the CFPB and FTC have been more hands-off. They have allowed consumers to test these products before imposing significant regulations on them. Of course, the risk here is that the statutory gaps identified above remain open or get wider and, if significant consumer harms develop before regulations or statutes take effect, consumers could have little recourse.

7.B. Licensing.

One attendee suggested state licensing of AI providers in a manner similar to the licensing and regulation of money transmitters and lenders. This would allow state regulators to evaluate business plans, check company leaders' backgrounds, and obtain surety bonds to cover potential future consumer losses. While this may increase oversight of industry participants, particularly at the state level, others expressed concern that licensing can add substantial burdens particularly for startups and new entrants who lack the resources and time it takes to obtain licenses nationwide. For example, obtaining money transmission licenses in the 49 states can take upwards of \$1 million and 2 years.⁵⁹

7.C. FCRA-Style Dispute System.

One attendee suggested development of a dispute system like that used in credit reporting disputes. Under FCRA, consumers can dispute information in their credit file that they believe is inaccurate, incomplete, or unverifiable.⁶⁰ The

⁵⁸ See, e.g., THE WHITE HOUSE, WINNING THE RACE AMERICA'S AI ACTION PLAN (2025), <https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf>.

⁵⁹ Overcoming the challenge of obtaining U.S. money transmitter licenses, TRULIOO BLOG (Nov. 24, 2022), <https://www.trulioo.com/blog/payments/us-money-transmitter-licenses>.

⁶⁰ 15 U.S.C. § 1681i.

consumer reporting agency that received the dispute must conduct a reinvestigation to determine whether there is an inaccuracy and notify the furnisher.⁶¹

The company that reported the information, the furnisher, must also conduct an investigation and report their findings back to the consumer reporting agency.⁶² If the consumer reporting agency and/or the furnisher finds inaccurate information, they must promptly correct, delete, or block such information.⁶³ Applying that to agentic disputes, consumers questioning an AI-directed outcome could seek the data underlying the AI tool and evaluate whether the tool made a decision based on inaccurate or incomplete data. Additionally, such a system could require AI developers to disclose model inputs. Even if consumers do not use this information, regulators could use it to evaluate models, including for potential conflicts of interest.

This framework addresses a critical gap: consumers today have no systematic right to challenge algorithmic decision-making outside the credit context. The FCRA framework is well-established and has proven adaptable to different credit decision contexts. Transparency about data and decision-making could help identify and remedy bias and discrimination.

But FCRA's dispute mechanisms can be slow and cumbersome. If agentic payments operate at a high velocity, FCRA-style processes may be inadequate. Disclosing detailed information about data and algorithms could create intellectual property concerns for developers. FCRA-style rights assume discrete decisions (credit approval/denial), but agentic payments may involve continuous series of transactions; the framework may not map well to the scope of potential agentic commerce transactions. And the administrative burden of responding to disputes could also be substantial for providers without meaningful limits on the types of disputes that could be submitted.

7.D. RESPA-Style Prohibitions on Kickbacks and Self-Dealing.

The Real Estate Settlement Procedures Act (RESPA) prohibits any person from giving or accepting a fee, kickback, or thing of value pursuant to an agreement or understanding for referrals of a real estate settlement service involving federally related mortgage loans.⁶⁴ It also bars unearned fees or fee splitting, though exceptions exist for bona fide payments for services actually performed, such as by attorneys or agents.⁶⁵ A similar statute in the agentic payments space could prohibit companies offering AI tools from entering into agreements with certain

⁶¹ *Id.*

⁶² *Id.*

⁶³ *Id.*

⁶⁴ 12 U.S.C. § 2607(a).

⁶⁵ 12 U.S.C. § 2607(b).

merchants and receiving kickbacks or another thing of value when programming AI tools to only make purchases with those merchants and such AI tools do so. This could address self-dealing and similar risks identified above.

7.E. Increased Role of the FTC.

Some attendees suggested that the FTC, as the primary federal regulator of non-banks and technology companies that offer products and services to the public, should take a greater oversight role in this space. For example, many of the advertising and marketing related conduct that troubled some attendees is already addressed in the FTC actions related to online advertising and marketing. Others suggested that the FTC could use its UDAP authority, in particular, to root out conduct that harms consumers or other industry participants. Already in 2025, the FTC has taken action against a number of companies engaged in AI-related conduct that the FTC determined violated the FTC Act.⁶⁶ Of course, the FTC is a relatively small agency and its approach to “regulation by enforcement” makes it harder for industry participants and others to identify conduct the FTC views as problematic prior to an enforcement action.

7.F. Standard Setting Organization.

A dedicated standard-setting organization could be established by industry or could be a public-private partnership with regulators. Such an organization could establish standards for agent development, disclosure, transparency, and consumer protection. Agents could be “certified” if they meet such standards, and uncertified agents would face restrictions or prohibitions.

This is similar to how the Payment Card Industry Data Security Standard (PCI-DSS) establishes baseline security standards for payment processors. In the agentic payments space, a consortium could also establish baseline standards for agentic AI developers. Compliance with standards could be a condition of participating in major payment networks or integrations. Or, similar to how EMV (chip card) security standards are governed and enforced, the industry could establish certification standards for agentic systems. To participate in commerce

⁶⁶ Press Release, Fed. Trade Comm’n, FTC Order Requires Workado to Back Up Artificial Intelligence Detection Claims (Apr. 28, 2025), <https://www.ftc.gov/news-events/news/press-releases/2025/04/ftc-order-requires-workado-back-artificial-intelligence-detection-claims> (finding Workado’s marketing of its AI Content Detector as “98% accurate” was false, misleading, and non-substantiated as the product was not as accurate as it claimed); Press Release, Fed. Trade Comm’n, FTC Sues to Stop Air AI from Using Deceptive Claims about Business Growth, Earnings Potential, and Refund Guarantees to Bilk Millions from Small Businesses (Aug. 25, 2025), <https://www.ftc.gov/news-events/news/press-releases/2025/08/ftc-sues-stop-air-ai-using-deceptive-claims-about-business-growth-earnings-potential-refund> (alleging Air AI Technologies deceptively marketed and sold materials as consumers often did not receive the promised profits nor were they able to receive the “guaranteed” refund for lost profits); Press Release, Fed. Trade Comm’n, FTC Issues Staff Report on AI Partnerships & Investment Study (Jan. 17, 2025), <https://www.ftc.gov/news-events/news/press-releases/2025/01/ftc-issues-staff-report-ai-partnerships-investments-study> (announcing a staff report of the FTC’s opinions on the structure of CSP and AI developer partnerships and certain rights CSPs gain through such investments, and reminding companies that such partnerships should not undermine open markets, opportunity, or innovation).

(e.g., to integrate with merchant checkout systems or payment networks), an agentic system would need to certify its compliance with standards and would be subject to ongoing monitoring and audit.

There are several benefits of such systems, the most notable being speed. Industry standards can adapt quickly to technological change, faster than statutory or regulatory frameworks. Additionally, certification provides a mechanism for third parties and consumers to assess whether an agentic system meets baseline standards. As noted above with card network rules, such systems can preserve flexibility for innovation while ensuring baseline consumer protections. Finally, precedent exists given decades of success with existing systems in payments.

There are some limitations inherent in this approach, including that standard-setting organizations might reflect interests of those who develop the standards and could discount consumer protection. Second, these systems must also contain mechanisms to punish non-compliance. And lastly, standards may need to be enacted globally so that U.S.-only standards do not inhibit development of agentic tools.

7.G. Contractual.

Some attendees focused on contract formation between the various parties involved in creating agentic tools. In such contracts, industry participants can push for inclusion of additional consumer friendly provisions. Other attendees suggested consumers should or could be guaranteed certain rights that cannot be waived in these contracts or that AI tools use standard contract terms. While this could protect consumers, it was unclear whether and how such provisions could be adopted in the absence of industry standards or applicable law.

7.H. Legislation.

Some attendees suggested the enactment of legislation that would regulate agentic payments tools, although this was the opinion of a distinct minority of attendees. Legislation could establish guardrails or clarify consumer protections. Some attendees suggested that if there was going to be regulation, policymakers should anchor regulations on models or on the people building the models so there is a legal obligation on an actual person or company creating an agentic tool. Others suggested efforts be taken to avoid a patchwork of authorities (a problem particularly acute in financial services), but that it could be difficult for lawmakers to create a uniform set of rules for all actors. One attendee noted that jurisdictional issues could be particularly acute because of the global pace of AI development and because different participants in the United States are regulated by different entities.

8. Conclusion

Attendees agreed that adoption of AI generally and agentic payments tools in particular will only accelerate and that this acceleration is inevitable. In the months since the Symposium, announcements of new agentic payments tools have proliferated. Though it may be too early to assess consumer adoption and consumer harms, if any, resulting from these new tools. Regardless, the growth of agentic payments tools may result in two conflicting ideas:

*“This is going to get radically better,
and it is going to screw things up.”*

Preventing consumer harm must guide those developing these new and exciting tools. Collaboration now is critical before agentic AI scales beyond control. AI introduces complexity in places where consumers can go to rectify agent-created harms. The opacity of agents—including who influences, controls, and owns them—may make accountability more difficult. When these agents implicate consumer payments, the consumers’ banks will inevitably be involved when there are issues. To minimize these issues, all players will need to have appropriate involvement and those in the best position to manage and minimize risks to consumers and others must do so.