Agentic AI Commerce: The Next Wave of Online Shopping and Retailer Risk

James Gatto and Ericka Schulz

Agentic AI commerce is here, and it is predicted to start scaling this year. We are talking about AI agents that can autonomously perform tasks for a user and can adapt over time as the agent learns more about the user, all without direct human intervention. Consumers can engage various agents to search, select, purchase, and pay for a variety of goods and services.

Many retailers offer their own AI agents or shopping bots to help their customers' shopping experiences. Retailers that offer their own agents can control the user experience, the agents' actions and the risks. But an army of third-party AI shopping agents are being deployed, and users are creating their own agents. Retailers do not control these agents and as detailed below, this creates a host of potential legal and business risks for those retailers.

One major ecommerce platform operator, apparently recognizing the risks, recently obtained a preliminary injunction to prohibit agents from accessing the ecommerce platforms' electronic systems without permission from the platform owner because the agents access the platform *with the user's permission* but access the user's password-protected account, *without authorization of the platform operator*.

Many retailers are not yet ready for the onslaught of AI agents hitting their sites and transacting on behalf of users. But now is the time for retailers to manage the legal risks that will arise with these AI agents.

Why AI Agents Are Different Than Existing Automation Tools

Online shopping already offers automated tools such as recurring purchases, AI-generated review summaries, automated recommendations and optimization engines that adapt based on customer behavior and purchase history. Typically, a person will use these tools to help them search for and buy products of interest to them—or retailers will implement them in the background. These tools are typically provided by the merchant as part of the ecommerce platform.

AI agents are much more autonomous. With agentic commerce, a person gives the AI agent a high level goal or task (e.g., buy me X) with certain constraints (for no more than Y dollars) and the agent will search many sites for a suitable product or service based on that information and what it already knows about the person (e.g., prior transactions, purchase preferences and more). AI agents can discover products, assemble a cart, and execute checkout and payment on the customer's behalf—often without the customer visiting a retailer's site and often without approving the specific transaction (although this is an option).

AI agents differ from automated tools in a number of ways. Most automated tools are rules-based: the user creates the rules (buy product X every 2 weeks from a particular site), and the rules remain fixed unless and until the user changes them. They are predictable. The tools execute exactly what the user specifies. By contrast, AI agents are more autonomous. Users give them goals and constraints, and the agent figures out what to do with some autonomy. And unlike fixed rules, how the agent accomplishes the goals can change over time because, as the agent learns more about the user from prior outcomes and user preferences, the agent adapts its behavior accordingly. For example, the agent may search every 2 weeks for whichever site is offering the best deal on product X and buy it there using a form of payment associated with the user.

AI agents developed by third parties and users, rather than the ecommerce platform operator, are fraught with potential legal and business risks for the ecommerce platform operator.

Various Efforts Are Underway to Understand and Mitigate AI Agent Risks

Various efforts are underway to understand and mitigate these risks. One aspect of the risk relates to agentic-based payments for consumer transactions.

The Consumer Bankers Association (CBA) met in the fall of 2025 to scope out consumer protection issues with agentic payment tools. Some of the key takeaways from this were:

- Agentic payment tools have massive potential to disrupt the existing consumer payments landscape and revolutionize commerce
- Existing consumer protection rules encouraged digital payments to flourish but have an uncertain application to agentic payments; the general rule in the Electronic Fund Transfer Act (EFTA) that limits consumers' liability for unauthorized transactions may not apply when agents are involved
- Some agents may use payments outside the current card networks (e.g., using stable coins) and this may not afford consumers traditional cardholder protections (e.g., no liability for unauthorized transactions; additional chargeback rights, etc.)
- Immediate statutory and regulatory changes are unlikely and may be unnecessary; immediate government action appears unlikely in the short term as the Trump Administration has signaled that policymakers should allow AI and agentic tools to flourish
- Fraud and scams – new fraud schemes with AI agent commerce may cause retailers increased losses from authorized transactions that later proved fraudulent, great customer service burdens, and reputational damage, among other things

Due to the gravity of these issues, the National Institute of Standards and Technology (NIST) has launched (through its Center for AI Standards and Innovation (CAISI)), an AI Agent Standards Initiative to support the development of interoperable and secure AI agent systems. CAISI identified the need to do so based on heightened compliance, governance and cybersecurity risks.

Various AI protocols being developed will relate to AI shopping agents. Some of these include AI-to-system and AI-to-AI interactions, such as: (i) the Model Context Protocol (MCP)[1]; (ii) the Agent-to-Agent (A2A)[2] communication frameworks; (iii) the Agentic Commerce Protocol (ACP)[3]; and the Agentic Payment Protocol (AP2)[4], among others.

A2A frameworks enable AI agents to communicate directly with each other, coordinating complex tasks that require multiple specialized capabilities. In a retail context, this might involve a shopping agent collaborating with a logistics agent to optimize delivery timing, or a price comparison agent working with an inventory agent to provide real-time availability updates. A2A communication allows for sophisticated workflows where multiple AI systems need to work together seamlessly.

ACP establishes standardized methods for AI agents to discover, evaluate, and transact with retailers, defining how shopping agents request product information and compare offerings across merchants. AP2 complements this by standardizing the payment and

---

1 MCP (Model Context Protocol) is an open-source standard for connecting AI applications to external systems.

2 The Agent2Agent (A2A) protocol is a communication protocol allowing interoperability between AI agents from varied providers or those built using different AI agent frameworks.

3 Agentic Commerce Protocol (ACP) is an open standard that serves as the connective layer between merchants and ChatGPT users. It enables ChatGPT to ingest structured catalog data, understand merchant inventory, and surface relevant products in context.

4 The Agent Payments Protocol (AP2) is an open protocol developed with leading payments and technology companies to securely initiate and transact agent-led payments across platforms. It can be used as an extension of the A2A Protocol and MCP Protocol. It purports to rely on industry rules and standards to establish a payment-agnostic framework for users, merchants, and payments providers to transact with confidence across all types of payment methods.

checkout process for AI-mediated transactions, addressing secure credential management and transaction authorization for autonomous purchasing.

MCP is the "language" that allows AI shopping agents to understand and interact with retail systems in a consistent, predictable manner. Just as HTTP standardized web communications and enabled the internet revolution, MCP aims to standardize how AI agents access and interpret business data across all phases of the customer journey, from browsing through checkout to fulfillment and beyond.

The strategic advantage becomes clear: by investing in MCP infrastructure first, retailers build a flexible foundation that can accommodate specialized protocols like ACP and AP2 as they mature, while immediately establishing the technical capabilities needed to serve the emerging AI agent ecosystem. This positions MCP not as a replacement for domain-specific protocols, but as the essential baseline that makes all subsequent AI commerce innovations possible.

Risk to Retailers

We have previously covered some risks with AI shopping agents. See, *When AI Clicks "Pay": The Emerging Compliance Risks of Agentic Commerce*. It is beyond the scope of this article to provide a detailed explanation of all of the potential risks retailers may face when AI agents interact with their platforms. However, some of the issues for retailers to consider include the following:

- *Identification and Authentication of Agents* – retailers must identify and authenticate agents to prevent fraud and other potential harms. It will be important to recognize when website traffic is an agent rather than a person. And it will be necessary to authenticate the agent as legitimate. Many third-party or user-created agents may intentionally or unintentionally operate improperly, leading to fraud or other losses for the retailer. Some bad actors may try to mask the fact that it is an agent rather than a person accessing the site.

- *Authorization* – even if an agent is legitimate, retailers must confirm that the agent is authorized to: (1) act on behalf of the user; and (2) confirm the scope of the authorization. If an agent engages in an unauthorized transaction or exceeds the scope of its authorization, losses may result.

- *Contract Formation* – retailers must manage contract formation issues when: (1) an AI agent attempts to transact on behalf of a person who is not a current customer of the retailer and has not agreed to the retailer's Terms of Service (ToS); or (2) the AI agent is transacting for a current customer under a ToS that has not been updated to address agentic AI usage and to which the user has affirmatively assented. In either case, the user may not have an enforceable ToS for the transaction. Lack of an enforceable ToS may have many legal ramifications. One of the ramifications of this is that some risk mitigations that retailers use (e.g., mandatory dispute arbitration, waiver of class actions/mass arbitration) may be unavailable.

- *Will Courts Bind Agents' Actions* – AI agents are tools not people. Despite the term "agent" which has certain legal connotations, it is not yet clear whether courts will deem AI agents to be legal agents of the user. In many cases, an agent must be a person. If an agent makes an errant purchase, it is not clear whether the user or the platform operator will bear the risk of loss.

- *New Privacy Issues*? – If a third-party agent causes the platform operator to share user data with the third party, this may implicate privacy issues.

- *Limit or Require User Verification of an AI Agent Transaction* – platform operators may want to employ technical guardrails to limit or prevent (at least some) AI agent transactions. This need not be all or nothing. For example, a platform operator may allow low-value transactions but may prevent or require user authorization for higher-value transactions. Other criteria can be used.

- *Deceptive Practices* – Bad actors will always look to profit from early use of technology before the industry develops standard safeguards. There is no reason to believe this will not happen with AI agents. One likely issue may relate to referral fees for AI agent developers. For example, if instead of finding the best deal for a consumer, a third party AI agent is designed to favor transactions where the developer can collect a referral fee through an affiliate program or other arrangement with the platform

operator, without notice to the user, it is predictable that litigation or FTC enforcements for deceptive or unfair trade practices may ensue. While the developer here is clearly the bad actor, the platform operator may get caught up with the legal issues.

What Should Retailers Do Now?

Retailers who have not already done so need to learn about these issues and take steps to mitigate these and other risks. This includes a combination of technical measures, such as those described above, and legal measures. They should investigate the standards and protocols being developed and implement technical safeguards on their platforms to prevent against issues with AI shopping agents. Retailers should also consider revising their ToS, affiliate programs and privacy policies to address the potential issues with AI agents.

The law in this area will continue to evolve. It will be important for retailers to stay abreast of the changes and work with knowledgeable counsel to mitigate risks while leveraging the potential opportunities.

Agentic commerce is not merely a new feature—it is a shift in who (or what) is “shopping,” at what speed, and under what authority. The promise is real: higher conversion, lower friction, and a new channel that could reshape how products are discovered and purchased. But the risks are equally real, and they arrive in familiar forms—fraud, disputes, privacy exposure, consumer protection scrutiny, and litigation—only faster, more scalable, and more difficult to untangle.

Online merchants should treat AC as a near-term readiness issue, not a distant innovation. Those who start now—by tightening authorization concepts, clarifying liability allocation, increasing transparency around incentives, hardening privacy practices, retooling fraud controls for agent-driven traffic, and validating accessibility and bias safeguards—will be better positioned to capture the upside without becoming the default “deep pocket” when the first major agentic failures hit.

Contacts

James Gatto

Partner | AI, Robotics and Quantum Team Co-Leader

+1.202.747.1945

jgatto@sheppard.com

Bio

Ericka Schulz

Partner

+1.858.720.7485

eschulz@sheppard.com

Bio