


Deloitte.

Together makes progress




No human, no contract:
The legal risks behind
AI commerce

Executive summary

Agentic commerce represents the next evolutionary leap in digital retail. Autonomous artificial intelligence (AI) bots could do all the research and purchasing of the perfect product for you—and at the best deal too—no intervention required. It has the potential for unparalleled convenience and efficiency. The scale is already significant: US retail websites have seen a 4,700% jump in traffic from Generative AI sources.¹ AI agents influenced 20% of all Cyber Week orders globally, totaling \$67 billion in sales.² Morgan Stanley projects that agentic shoppers could reach 10% to 20% of US e-commerce spending by 2030, representing \$190 billion to \$385 billion in transaction volume.³ But does it align with current law?

Like autonomous vehicles, the pace of agentic technology has outpaced established legal frameworks, creating a landscape of risks—both known and unknown.

Current laws governing online transactions are built on the foundational principles of human intent, notice, and consent—concepts that are not easily applied to autonomous AI. For CEOs and business leaders, not addressing this misalignment introduces significant risk. It invites the risk of unenforceable sales contracts, regulatory action from bodies like the Federal Trade Commission (FTC), and significant liability exposure.



The challenge: A legal framework *built for humans*, not bots

The core of the issue lies in a fundamental disconnect. The entire legal structure of e-commerce, from the Uniform Electronic Transactions Act (UETA) to the Electronic Signatures in Global and National Commerce (ESIGN) Act, is predicated on the idea that a “person” must assent to a contract. The law requires a clear action from a human to form a binding agreement. When an AI agent autonomously clicks “I agree” on a terms of service page, who has truly consented? Can a bot, which lacks legal intent, bind its human user to a contract the human user has never seen?

Courts have consistently reinforced this principle. Rulings like *Nguyen v. Barnes & Noble Inc.* have established that for a contract to be valid, the user must have actual or constructive knowledge of the terms. An AI agent’s actions obscure this, making it difficult, if not impossible, to prove the human consumer was aware of your terms of sale. **This creates a critical vulnerability in which transactions facilitated by AI may be deemed legally unenforceable.**

Key risk areas for your business

1. Unenforceable sales and revenue loss

The most immediate risk is to the validity of your sales. If a consumer uses an AI agent to make a purchase, and that transaction is later disputed, you may find that the underlying sales contract is not legally binding. As established in cases like *Berman v. Freedom Financial Network, LLC*, online agreements require two things: reasonably conspicuous notice of the terms and an unambiguous action that manifests assent. It is highly questionable whether a bot's pre-programmed action fulfills this two-part test on behalf of an unaware user, putting your revenue from such transactions at risk.

2. FTC scrutiny and “unfair practices”

The FTC has broad authority to police “unfair or deceptive acts or practices.” The use of bots in commerce is already on its radar, as evidenced by the Better Online Ticket Sales (BOTS) Act, which targets ticket-scalping bots. The FTC’s Unfairness Policy Statement, codified in federal law, allows the agency to act against practices that violate established public policy, even if not explicitly illegal. Allowing your platform to be utilized by autonomous agents without clear guardrails could be interpreted as an unfair practice, inviting regulatory investigation, fines, and damage to your brand’s reputation.

3. Liability for unauthorized access and system misuse

Many websites employ security measures like CAPTCHA to prevent automated abuse. As the court found in *Ticketmaster L.L.C. v. Prestige Entm’t, Inc.*, using a bot to circumvent these measures can constitute a violation of your terms of use and even be considered unauthorized access under statutes like the Computer Fraud and Abuse Act (CFAA). Furthermore, under the precedent set in *Mohon v. Agentra LLC*, your company could even face vicarious liability if your platform is used by third-party AI agents to conduct illegal activities, especially if you knew or should have known about the misuse.

4. A growing patchwork of state-level AI regulations

A complex web of state laws is emerging that mandates disclosure when AI is used in commerce. States like California, Maine, Utah, and Colorado have already enacted statutes requiring that consumers be notified when they are interacting with an AI system. While many of these laws are currently aimed at disclosure by the *seller*, they signal a clear legislative trend toward transparency and accountability in AI-driven commerce. Failing to get ahead of these regulations means constantly reacting to a shifting legal landscape, rather than leading from a position of strength.

A case already in court


The risks outlined above are already being litigated. In November 2025, a major US retailer filed suit against an AI commerce platform over its autonomous shopping agent, among the first cases to directly test the legal boundaries of agent-driven purchasing in court.

The retailer's position rested on three arguments:

1. The agent was impersonating a human consumer;
2. Its terms of service were designed for human visitors, not automated systems; and
3. Its advertising and recommendation revenue was undermined when the agent bypassed those surfaces.

The platform countered that the retailer's core concern was its inability to monetize agent-driven traffic the same way it monetizes human traffic, and that agents operated on behalf of users who had granted them access.

Regardless of outcome, the case signals that the legal questions raised in this paper are no longer theoretical. They are being decided now.



A *proactive strategy* to help mitigate risk and embrace the future

While the legal landscape is fraught with risk, it is not an impasse. Agentic AI commerce is not prohibited, but it is misaligned with current legal doctrines. Companies that take proactive steps to bridge this gap can not only help protect themselves from liability but also build the trust necessary to lead in the next generation of e-commerce.

Here are important steps to consider:

ACTION ITEM	STRATEGIC RATIONALE
Update your terms and conditions (T’s & C’s)	Explicitly address the use of AI purchasing agents. Work with your legal team to evaluate whether your current Terms & Conditions adequately account for bot-mediated transactions and to identify gaps that may require revision.
Implement AI-friendly consent mechanisms	Rethink the “clickwrap” agreement for a new reality. Consider creating a specific, one-time user verification process for authorizing an AI agent to make purchases on their behalf, creating a clear record of “informed consent.”
Remove ambiguous security barriers	If you wish to permit agentic commerce, traditional anti-bot measures like CAPTCHA create a legal contradiction. Removing them for designated AI interaction points could mitigate the risk of a “terms of use” violation for unauthorized access.
Lead with clear disclosure	Get ahead of regulation. Clearly and conspicuously state your policy on AI agents. If you permit them, explain how the process works. This transparency builds consumer trust and demonstrates a good-faith effort to navigate the legal complexities fairly.
Ensure user control at the final step	As emphasized by the FTC and financial regulations (such as Regulation E), authorization is narrowly construed. Implement a final, human-in-the-loop confirmation step for high-value transactions or for the initial setup of an AI agent’s purchasing authority to solidify customer authorization.

Beyond legal risk: Revenue exposure and operational vulnerability

The risks of agentic commerce extend beyond contract enforceability and regulatory exposure. As AI-driven purchasing scales, two additional areas deserve the same level of scrutiny: the impact on marketing-driven revenue and the operational risk of agent-initiated chargebacks.

Revenue exposure: Marketing in an agent-driven world

Amazon generated more than \$68 billion in advertising revenue in 2025,⁴ built on the assumption that human shoppers see ads, feel urgency, and respond to promotions. AI agents do none of these things. They select through structured data and API calls, bypassing the visual and emotional triggers that drive traditional marketing spend. As agent-driven purchases grow, marketing budgets optimized for human browsing may yield diminishing returns. If 20% of Cyber Week orders were already agent-influenced, the question is immediate: How much of your marketing budget is reaching the buyer that actually converted? Whether there is an emerging channel optimized for machine-led discovery is a question worth exploring now.

Operational risk: The chargeback problem

Retail e-commerce chargebacks surged 233% in 2025.⁵ When an AI agent buys on behalf of a consumer, the authorization question becomes acute: If the consumer did not explicitly approve the transaction, the purchase may be disputed as unauthorized under existing chargeback frameworks. The scenarios compound quickly: duplicate orders from competing agents acting on the same intent, purchases that do not match actual preferences, and bad actors using agents to generate fraudulent disputes at scale. Fraud models are built for human purchasing patterns. Whether your systems are prepared for agent-driven volume, velocity, and dispute complexity is an operational question that warrants immediate assessment.

This paper focuses on the legal dimensions of agentic commerce. The revenue and operational implications outlined here are areas for further exploration, not settled conclusions. They are included because the same underlying shift, the removal of the human from the transaction, creates exposure across multiple dimensions simultaneously.

Conclusion

The era of agentic commerce is not on the horizon; it is at our doorstep. While it presents a landscape layered with legal ambiguity and risk, it also offers a pathway to innovation and enhanced consumer experience. The companies that will likely succeed are not those that ignore the risks, but those that confront them head-on. By proactively updating your legal frameworks, embracing transparency, and reengineering consent for an AI-driven world, you can help shield your organization from unnecessary liability and position your business as a trusted leader in the future of digital retail.

Authors



Celia Ludwinski
Specialist Leader, SFL Scientific
Deloitte Consulting LLP
cludwinski@deloitte.com



Kevin Byrne
Principal, US Retail and Consumer
Products AI Leader
Deloitte Consulting LLP
kebyrne@deloitte.com

Special thanks to Emma Guarnieri and Juan Rivas Caban for their contributions.

Endnotes

1. Vivek Pandya, "[Adobe: Generative AI-powered shopping rises with traffic to U.S. retail sites up 4,700%](#)," Adobe, August 21, 2025.
2. Salesforce, "[Salesforce data: AI and agents propel Cyber Week to record \\$336.6B in global spend](#)," press release, December 5, 2025.
3. Morgan Stanley, "[Here come the shopping bots](#)," December 8, 2025.
4. Peter Adams, "[Amazon's annual ad revenue passes \\$68B, boosted by full-funnel strategy](#)," *Marketing Dive*, February 6, 2026.
5. Payscout, "[The 233% surge: Why retail chargebacks soared in 2025 \(and what to do about it\)](#)," accessed May 2026.

SFL Scientific, a Deloitte business, is the artificial intelligence/machine learning (AI/ML) technical arm of Deloitte Consulting LLP.

We pair research-grade science with production engineering to turn frontier methods in agentic large language models (LLMs), mixture of experts, computer vision, and robust evaluation into systems capable of standing up to real-world constraints. Our scientists design the methods and our engineers harden them for cloud, graphics processing unit (GPU), and edge. SFL Scientific also enables client teams through hands-on delivery, platform benchmarking, and applied learning via Deloitte's AI Academy. Together, we help organizations establish AI capabilities that can scale efficiently, withstand change, and create lasting value.



SFL Scientific
a Deloitte business

About Deloitte

As used in this document, "Deloitte" means Deloitte Consulting LLP, a subsidiary of Deloitte LLP. Please see deloitte.com/us/about for a detailed description of our legal structure. Certain services may not be available to attest clients under the rules and regulations of public accounting.

Disclaimer

This paper is for informational purposes only and does not constitute legal advice. Readers should consult qualified legal counsel for guidance specific to their circumstances.

This publication contains general information only and Deloitte is not, by means of this publication, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This publication is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional advisor. Deloitte shall not be responsible for any loss sustained by any person who relies on this publication.