Agentic Commerce Needs Legal Infrastructure, and The Courts Are Coming

Bridget McCormack, CEO, American Arbitration Association
David Hoffman, William A. Schnader Professor at Penn Carey Law, the University of Pennsylvania

April 22, 2026

Monday Morning

A mid-sized industrial manufacturer deploys a procurement agent to handle routine supplier onboarding across its North American operations. A configuration error causes the agent to accept counterparty terms that include a broad indemnification clause running in the supplier's favor; under the clause, the manufacturer agrees to indemnify the supplier against any third-party claim related to the goods, including claims arising from the supplier's own conduct. Over the weekend, the agent completes 12,000 transactions across hundreds of suppliers. On Monday, operations proceed normally. The goods are accepted, shipments move, and invoices are paid.

That manufacturer faces serious legal risks, under current doctrine. Black-letter ratification rules make principals liable for affirming a prior act with notice of terms: whether the manufacturer can escape liability by pointing to the scale that agentic commerce produces, or will find itself subject to an expensive lesson in the bite of the conscious ignorance rule, will almost certainly be impossible to resolve cheaply. Adjacent doctrines, such as apparent authority and negligent control of agents, provide additional avenues to bind our buyer. One might therefore think that agentic commerce will produce workaday problems with solutions that courts and arbitrators can readily deploy. Not so.

The investment in agentic commerce is significant and growing daily. But the investment thesis has a blind spot. Every commercial system in history has required legal infrastructure to function at scale: mechanisms for forming enforceable agreements, resolving disputes without litigation, satisfying regulatory obligations, and producing evidence when things go wrong. Traditional e-commerce built that infrastructure deliberately, over two decades. Agentic commerce is not building it. The protocols being deployed right now are sophisticated about payments, identity, and interoperability. They are largely silent on what happens when a transaction goes sideways.

That silence is why doctrines like ratification and apparent authority will not resolve these cases cheaply. They were built for transactions with humans in the chain—reviewing terms, signing records, flagging irregularities as they arose. Agentic commerce is removing humans from those roles faster than it is building substitutes. This paper argues that the

legal infrastructure absent from agentic commerce is not a gap the market will fill incrementally. It is an active, accumulating liability that courts and regulators will allocate on their own timelines, under doctrines not designed for autonomous AI agents. The doctrines are already in place, and the market is already in flight; the organizations building this market—and the institutions that support them—should act to shape legal governance rather than have it imposed, case by case.

A note on scope. This paper is not about whether AI agents are reliable, whether they will make good commercial decisions, or whether governments should regulate them. Those are AI questions. This paper is about the legal infrastructure that records what was agreed, by whom, and under what authority. AI variability and legal infrastructure are orthogonal concerns. A flawed agent that produces a clear legal record is in a fundamentally different position than a flawed agent that produces none.

Why the E-Commerce Framework Doesn't Transfer

Traditional e-commerce has not produced systemic legal chaos at scale because, over two decades, it built working legal architecture: a consent model courts recognize, where clickwrap and analogous mechanisms can establish manifestation of assent; a private dispute resolution layer of chargebacks, network arbitration, and platform guarantees that resolves hundreds of millions of disputes annually without involving the courts; statutory compliance backed by federal and state consumer protection law and meaningful enforcement; and a verifiable evidentiary record—terms at a URL, a timestamped click, a logged transaction, payment-network records that survive the dispute.

These pieces are interdependent. They also share a foundational premise: a human is present at the point of transaction. Agentic commerce removes that premise, and each of these mechanisms behaves differently when the transacting actor is an autonomous agent rather than a human.

Clickwrap works because a human clicked, and browsewrap works because a human had reasonable notice. An agent's programmatic retrieval of a terms URL is something different, and no court has yet decided whether it constitutes the kind of notice and consent that contract law requires. The failure of the notice model is not just about who is clicking but what is being consented to: terms drafted for human readers, surfaced by a protocol call, and accepted at machine speed.

Chargebacks and network arbitration cover credit and debit card transactions. Agentic commerce is increasingly settling on stablecoin and crypto rails (stablecoin transfer volume reached $33 trillion in 2025) that are irreversible by design. When something goes wrong on those rails, there is no chargeback to file, no network arbitration to invoke, and no platform guarantee to claim. That leaves litigation.

Federal and state consumer protection statutes require disclosures "clearly and

conspicuously" before a transaction completes. Agent-to-agent transactions complete in milliseconds. There is no temporal window for pre-transaction disclosure to a human principal.

Finally, the lack of any evidentiary record of the agreement compounds these problems. Two agents may negotiate terms dynamically, agree to conditions neither human principal reviewed, and execute at machine speed. There is no standardized mechanism for recording what was offered, what was accepted, the agent's authority parameters, or the governing framework that applies.

Insurance may act as a shock absorber, but a costly one. When legal ambiguity is hard to price, premiums reflect it.

The Doctrine

Courts can't wait for new legislation; they decide the questions before them as they come. The doctrines that govern agentic commerce disputes were written for human agents, deterministic systems, and paper contracts, but, by their terms, they apply to AI agents transacting autonomously.

Common law agency doctrine—apparent authority and ratification—will impose contractual obligations on principals whether they specifically intended them. UETA Section 14 and the federal E-SIGN Act provide the statutory foundation for electronic agent contracting, attributing intention from "the programming and use of the machine." Those statutes were drafted for deterministic systems where the principal fully specified the agent's behavior.

An AI agent with broad autonomy and the capacity for emergent behavior stretches the attribution chain in ways the drafters did not contemplate. The programming and use of a deterministic script is a direct expression of intent; the programming and use of a Large Language Model (LLM) is the delegation of intent to a probabilistic process. When the agent agrees to a term the principal never considered, the attribution becomes a legal fiction used to preserve commercial certainty at the expense of traditional agency principles.

The UCC's warranty disclaimer rules under Section 2-316 ask whether the disclaimer was rendered conspicuous in the text, which may protect B2B sellers who post compliant terms. For B2C transactions, where an agent accepts terms on behalf of a consumer who never reviewed them, the answer is less clear, and Magnuson-Moss adds federal disclosure requirements that go beyond posting text. UCC Section 2-207's protections against surprise terms in the battle of the forms assume parties who can review and object, a structural premise that agentic transactions complicate.

The FTC Act Section 5 and state UDAP statutes already reach AI-related commercial

conduct, and many provide per-violation penalties, fee-shifting, and class-action mechanisms. The clickwrap and browsewrap doctrines that have governed e-commerce consent for two decades depend on human manifestation of assent. When automated code and surrounding legal documentation diverge, courts will reach for parol evidence, integration, and equity to interpret the relationship (See Cohney and Hoffman's contract stack framework).

Cross-border transactions are more complicated still. UNCITRAL adopted a Model Law on Automated Contracting in 2024, addressing automated contracting directly, but no jurisdiction has yet enacted it. The EU AI Act is in phased enforcement; its interaction with existing contract and agency doctrines is untested.

Ratification

Each of these doctrines will produce hard cases. Ratification is worth additional focus because its operation in agentic contexts compounds in ways the other doctrines might not.

Generally speaking, a principal who accepts the benefits of her agent's transactions—knowing that it lacks knowledge of their specifics but proceeding anyway—has ratified those transactions. That includes every specific term the agent agreed to, including terms the principal never saw and would never have consciously authorized. Unlike apparent authority, ratification requires no reasonable belief on the counterparty's part. It requires only that the principal accepts the benefit, knowing what it does not know.

When an agent transacts with a large and diverse population of counterparties, each presenting its own terms, ratification operates as continuous automatic incorporation of whatever terms each counterparty publishes. UETA Section 14 attributes intention from the programming and use of the machine. The scholarly literature on algorithmic contracting supports the framework (Scholz, 2017). *Mobley v. Workday* suggests courts are moving toward applying these principles to AI systems. Critics may argue that ratification requires the principal to have actual knowledge of the material facts of the transaction, and certainly, this is a standard statement of the doctrinal rule. In an agentic environment, however, courts are likely to view the intentional deployment of an autonomous agent with authority to bind as a form of assumption of risk or willful blindness. If a principal delegates the authority to contract at machine scale without implementing a monitoring architecture, they may be estopped from claiming lack of knowledge. As in *Mobley v. Workday*, the focus shifts from what the principal *knew* to what the system they built *did*.

Courts Have No New Tools

Judges do not need to understand agentic architecture to rule on disputes involving agents. They have contract law, agency doctrine, the UCC, consumer protection statutes, and

negligence. They will take the fact patterns agentic commerce generates and fit them into the categories that govern them.

The litigation won't be straightforward or quick, either, given the imperfect evidentiary record. When these disputes reach courts, there will often be no standardized record of the terms offered, the terms accepted, or the governing framework. That makes litigation longer, more expensive, and less predictable for everyone. Machine-scale transactions produce machine-scale litigants. A single misconfiguration creates thousands of similarly situated claimants in hours.

Jurisdictional fragmentation guarantees inconsistent results. An agent acting for a buyer in California interacts with a seller's agent in Ireland, transacts through a protocol developed in San Francisco, and settles through a network based in Singapore. Without contractual governing-law clauses agreed in advance, the question is answered after the fact by whichever court the case is filed in. The same pattern will be decided differently in East Texas, the Southern District of New York, and the Commercial Court in London.

What the First Cases Will Look Like

The ratification class action. An enterprise deploys an agent with broad procurement authority. The agent accepts counterparty terms across hundreds of vendors. A term buried in one vendor's conditions—an IP assignment clause, a liability waiver, a non-compete provision—surfaces months later. The enterprise has accepted the deliveries; ratification is complete. Multiplied across every vendor whose terms the agent accepted without review, the exposure is substantial.

The consumer protection action. An agent acting on behalf of consumers accepts warranty disclaimers that fail Magnuson-Moss disclosure requirements or agrees to terms that violate state UDAP statutes. These statutes often carry per-violation penalties, fee-shifting provisions, and class-action mechanisms. All three compound at machine scale.

The evidentiary contest. Two agents negotiate dynamically. When a dispute arises, discovery produces the published terms, negotiated modifications from each agent, the agents' logs, and the platform's configuration. The court must determine which document constitutes the integrated agreement, whether the published file serves as an integration clause, and whether the dispute-resolution clause in the published terms extends to the negotiated modifications.

The liability cascade. A single disputed transaction touches the buying enterprise, the supplier, the procurement platform, the agent framework provider, the model provider, tool and API vendors, and payment processors. Without contractual allocation of liability agreed in advance, every participant can face exposure on every available theory. Third-party complaints multiply. The litigation becomes a multi-year exercise in allocating fault across an ecosystem that never agreed in advance on how fault should be allocated.

And Also, Torts…

The liability bubble is not contained by contract. When agentic failures cause systemic harm to third parties—such as a pricing agent triggering a flash crash or a procurement agent monopolizing a local supply chain—the dispute moves from contract to tort. Without a clear evidentiary record of the agent's instructions and safety parameters, these cases, too, will be litigated—under theories of negligence and strict liability for ultra-hazardous activities.

Let's Build the Infrastructure that Fits

The legal infrastructure for agentic commerce will be built. The question is whether it is built deliberately or reactively through litigation and regulatory enforcement. Built deliberately, it can include consent mechanisms that courts will recognize, clear terms for all parties, effective dispute-resolution layers that absorb friction before conflict reaches litigation, standardized evidentiary records, and contractual allocation of liability aligned with governing law. Built reactively, it will be by courts in thousands of jurisdictions working with very imperfect tools, on timelines no participant controls, under doctrines that will likely produce results no one in the ecosystem would have chosen.

Traditional e-commerce took two decades to build its legal architecture. Agentic commerce does not have two decades. The transactions are happening now, and so is the exposure.

Bridget McCormack is President and CEO of the American Arbitration Association–International Centre for Dispute Resolution and former Chief Justice of the Michigan Supreme Court. Dave Hoffman is the William A. Schnader Professor of Law at the University of Pennsylvania Carey Law School and a widely cited scholar whose research and teaching focus on contract law.