



Apr 20, 2026

Categories:

[Agentic Commerce](#)

[Series](#)

[Publications](#)

Authors:

[Nilesh \(Neal\) Patel](#)

[J. Dylan Grafe](#)

[Katelyn E. Moody](#)

When the AI Purchase Goes Wrong, Who Pays, and Who Can Prove It?

By the time a disputed transaction reaches the loss-allocation stage in agentic commerce, the system has already failed somewhere upstream.

If no one in the transaction ecosystem can say with confidence who acted, what the agent was authorized to do, or whether the resulting transaction amounted to assent to the transaction terms, the fight over who bears the loss is not a separate doctrinal curiosity, but rather the commercial consequence of an unstable transaction stack.

Even where identity, authority, and assent are reasonably specified, someone still has to absorb the consequences when the transaction goes wrong. The merchant may have fulfilled an order it thought it was entitled to trust. The payment intermediary may have authorized a charge through an apparently valid channel. The platform may say the agent acted within configured parameters in accordance with principal instruction. The principal may insist that the result was mistaken, excessive, or outside the scope of any authority they meant to grant.

At that point, the question is direct: who bears the loss, and who can prove why?

That is the real loss-allocation problem in agentic commerce. This article is not an attempt to build a full theory of every downstream remedy. Rather, it seeks to explore where many contested agent transactions will first get sorted in practice: payment disputes, merchant-reliance fights, and evidentiary contests over what the parties in the transaction ecosystem can actually prove. In many of the disputes that will matter most, the financial answer and the evidentiary answer will be inseparable.

Once the Transaction Happens, Someone Is Holding the Paper

Articles 2 through 4 addressed the threshold layers of the problem.

[Article 2](#) argued that identity is the first instance of instability in agentic commerce. If the parties operating in the system cannot say who is acting in a legally meaningful way, subsequent questions about attribution, reliance, and fraud posture are weak from the start.

[Article 3](#) addressed authority. Even if the system can identify the actor, it still has to explain whether the principal actually authorized the transaction the agent produced, and determine to what extent the scope of any authority given may have been exceeded.

[Article 4](#) then narrowed the contract-formation problem to its real role. Once identity and authority are reasonably specified, the remaining formation questions concern what terms governed the transaction, how assent was manifested, when the contract formed, and how discretion affected the enforceability story.

Loss allocation comes after all of that, but it is where the commercial pressure becomes unavoidable.

Once the order is placed, payment is initiated, goods ship, services are delivered, or a subscription renews, the system is no longer operating in abstraction. Someone has paid, shipped, relied, or extended value. Someone is now holding the paper.

That is why this part of the series matters. A transaction stack can tolerate uncertainty for only so long. Once value moves, the market must decide who the loss of a contested or unauthorized transaction should sit with: the principal, the merchant, the platform, the payment intermediary, or some combination of these parties. And in many cases, it will decide that question by asking a more operational one first: whose version of events can be proved and which parts of the transaction can be substantiated?

The Hard Cases Are Not Pure Fraud Cases

The easiest disputes to understand are the traditional ones.

If an attacker hijacks credentials, steals delegated tokens, spoofs an agent identity, or otherwise injects a fraudulent transaction into the system, familiar intuitions return quickly. The fight may still be

messy, but at least the commercial story is recognizable: this was not a genuine delegated transaction at all.

Nuanced, more complex disputes arise outside of traditional fact patterns.

The agent may be genuine. The account may be linked. The device may be trusted. The checkout path may be legitimate. The payment credential may be valid. The merchant may have fulfilled the order in good faith. The platform may insist that the system behaved as designed.

And still the principal may dispute the result.

That can happen in several ways. The agent chose the wrong merchant, accepted a bad substitute, bought at the wrong time, exceeded the principal's instructions (e.g., purchased more than authorized), bound the principal to a subscription or automatic renewal transaction the principal did not agree to, or acted within a technical permission envelope that was broader than the principal's real understanding.

Those disputes matter because they do not look like system failure in the usual sense. They look like contested delegated transactions. That is exactly why loss allocation becomes so important. The market will have to decide whether these cases should be treated more like fraud, more like user error, more like merchant-side reliance problems, more like product-design failures, or as a new hybrid category that pulls from several of those frames at once. However, while the market generally tries to decide who should bear the risk of loss, in the meantime it may be beneficial for the parties to allocate loss by contract, which could strengthen trust in emerging agent systems, reassure cautious principals, and avoid a regime in which risk is often defined only after harm occurs through litigation or rulemaking.

The Real Problem Is Misauthorized Commerce

A useful way to think about these disputes is to identify agentic transactions across a few categories, while recognizing that the boundaries between them are often the very thing being contested.

First, there is the classic unauthorized transaction. The system may show that something happened, but the principal never actually delegated the relevant authority, and the transaction is best understood as a fraud event, credential-compromise event, or outright impostor problem.

Second, there is the plainly authorized transaction. The user delegated clearly, the agent stayed within bounds, the merchant relied reasonably, the payment cleared, and the outcome is commercially and legally ordinary even if the principal later regrets it.



Figure 1: The Misauthorized Commerce Spectrum: most contested agent transactions fall in the middle zone, where the category boundaries are themselves what the parties dispute

The third, and perhaps most challenging category to consider, are the misauthorized cases. The principal did delegate something. The agent did act through a genuine system. The merchant may have relied on legitimate signals. However, the result still feels unauthorized in the way that matters once the transaction is challenged. The difficulty is that whether a given transaction falls into this third category (vs. one of the first two) is precisely what the parties will dispute. The category lines are not self-executing.

These misauthorized cases may become one of the most important dispute zones in agentic commerce. They include situations where the system acted within configured rules but outside the principal's real expectations, cases where discretion drifted into contested judgment, and cases where the commercial stack allowed a transaction to look valid even though the legal or practical basis for reliance was thinner than participants assumed.

Many future disputes will not turn on whether the transaction was "authorized" in some binary sense. They will turn on whether the wrong party is being forced to absorb the consequences of a transaction that was technically valid, legally contestable, and still commercially completed. Those disputes will often rise or fall on the quality of the record each participant can produce once the transaction is challenged.

The Four-Party Stack Makes the Dispute Harder, Not Easier

A typical contested transaction may involve at least three relevant actors: the principal or user, the merchant, and the payment intermediary. In agentic commerce, a new AI platform, such as a third-party tool with agentic commerce features, is added to the stack. When a dispute arises and an agentic AI tool is in the mix, each party may have a plausible account of what happened, and each may also have only part of the record.

Agentic Commerce: Liability & Transactional Flows




Figure 2: The four-party stack, the contested zone, and the evidence gap that runs through every disputed agent transaction.

The principal may say, with some force, that they never approved this purchase, this merchant, this price, this substitute, or this renewal.

The merchant may say it received a valid order through an approved channel, saw no visible red flags, delivered value, and should not be made to finance ambiguity inside someone else's delegation model.

The payment intermediary may say the credential was valid, the channel was recognized, the authorization request fit the available signals, and the rails performed as designed.

The platform, enabling agentic commerce tools, may respond that the user's settings, prompts, purchase history, or delegated permissions made the transaction permissible within the system's logic.

Each position may be plausible. That is exactly the problem. Agentic-commerce disputes are not difficult because no one has an argument. They are difficult because everyone does, and because the relevant proof available to substantiate transactions is fragmented across a transaction stack that was not designed primarily to explain itself after the fact. A transaction can be technically valid, legally disputed, and operationally hard to reverse all at once.

Merchant Reliance Will Be Tested Early

Merchants will want clean rules.

From the merchant's perspective, the instinct is straightforward. If the order came through a recognized channel, from a trusted platform or authenticated environment, with a valid payment method and no visible warning signs, the merchant will argue that it should be entitled to rely and fulfill.

That position is not unreasonable. Commerce does not work if every counterparty must independently reconstruct the full delegation history behind every agent-mediated order.

But the merchant case for reliance is not as clean as it may initially appear.

A merchant may be relying on a stack of signals that do different kinds of work: an authenticated session, a platform credential, a tokenized permission set, a payment authorization, a history of prior transactions, a merchant-approved integration, or the surrounding platform's representation that the agent is entitled to transact.

Those signals may justify operational trust. They do not necessarily justify putting the whole loss on the principal if the transaction later turns out to fall outside the principal's intended authority or understanding.

That tension will show up quickly. The merchant will say it received a valid order, fulfilled it, and should not finance the ambiguity inside someone else's delegation system. The principal will say the fact that the system could process the order does not mean the principal should bear the loss for an agent action they did not really authorize. The platform will often say the transaction fit the configured rules. The payment intermediary may say the payment credentials were valid and the rails performed as designed.

Each of those positions may be plausible. The outcome may therefore depend less on abstract instinct than on what the merchant knew, what the platform represented, what the payment layer recorded, and whether the record can show anything meaningful about the principal's actual delegation.

Payment Systems Will Not Want to Become the Universal Backstop

Payment intermediaries are unlikely to welcome this ambiguity.

The more agentic commerce grows, the more pressure there will be to decide whether disputed agent-mediated purchases should be treated as unauthorized payments, authorized but regretted payments, merchant disputes, or something in between.

That matters because payment systems are built around operational categories. They want to know whether a charge should be approved, reversed, disputed, or pushed into an existing loss-

allocation framework. They do not want every agent-mediated dispute to become a bespoke inquiry into human intention.

That pressure will shape the market.

If payment intermediaries conclude that too many agent-mediated purchases look operationally valid but commercially contestable, they will push for clearer allocation rules upstream. They may demand better confirmation logic, tighter delegated-authority structures, stronger merchant controls, clearer transaction metadata, or more legible evidence trails.

This does not mean payment intermediaries bear permanent legal responsibility for delegation failures that originate with platforms or users. But they do absorb the first operational hit in chargeback and dispute scenarios. And that first-hit exposure gives them significant practical incentive to reshape the architecture upstream. They may narrow their willingness to treat certain disputes as unauthorized payments if the real problem looks more like product design or delegation ambiguity, effectively forcing the cost back onto the parties best positioned to prevent it.

This is where existing consumer-payment doctrine starts to matter in a specific and consequential way.

Reg E, the Electronic Fund Transfer Act's implementing regulation, governs unauthorized electronic fund transfers and establishes the liability framework when a consumer disputes a payment as unauthorized. Reg Z governs open-end credit accounts and provides similar protections for disputed credit card charges. Both frameworks were designed for a world where the consumer either authorized a specific transaction or did not. In agentic commerce, that binary breaks down.

The open question is whether an agent-initiated payment is "unauthorized" within the meaning of those regulations when the consumer enabled the agent generally but did not specifically approve the individual transaction in dispute. This is of particular concern given the language of Reg E, which requires actual authority to initiate a transfer. A consumer who delegates broad purchasing authority to an agent and then disputes a specific purchase may argue that the individual charge was never authorized in any

meaningful sense, even where valid credentials were used.

Even if a principal asserts the charge was unauthorized, other parties in the transaction ecosystem (such as, payment intermediaries and financial institutions) may argue that liability is governed by traditional payment-law doctrines and the exceptions existing within those frameworks. For example, financial institutions may point to Reg E's treatment of transactions effectuated through an access device, under which the use of valid credentials can narrow or foreclose a principal's claim that a transaction was unauthorized, resulting in potentially conditioned or limited statutory protections tied to notice and timing.

This tension illustrates how, even under settled law, agent-mediated transactions can force a contested allocation of loss that turns on authorization, attribution, and proof rather than clear doctrinal labels. That mismatch of historical concepts with agentic tools has practical consequences. Financial institutions processing agent-mediated payments will need to develop internal frameworks for categorizing these disputes before regulators or courts resolve the question formally. The financial institutions best positioned to do that are the ones that understand both the technical structure of delegated agent credentials and the doctrinal contours of the existing regulatory regime. For most financial institutions, that gap between technical and legal understanding is where the first operational failures will occur.

Chargebacks and Reversals Will Become a Proxy Battlefield

One reason this issue matters commercially is that many of the first serious fights may not arrive as appellate doctrine. They may arrive as ordinary payment disputes.

A principal or cardholder may dispute a transaction that an agent placed. A merchant may insist that it relied reasonably and delivered value. A platform may say the order came through an approved integration path. A payment intermediary may have to decide whether the dispute fits an existing reversal or chargeback framework.

In that setting, the formal legal question and the operational dispute question may diverge.

The legal system may ask whether the principal authorized the agent, whether the merchant relied reasonably, and whether the terms of the transaction support enforcement. The payment system may instead ask a narrower question: is this transaction best coded as unauthorized, defective, disputed merchandise, user error, or something else the current rules can process?

That mismatch matters. It means that in the early stages of agentic commerce, loss allocation may be shaped less by elegant doctrinal reasoning than by the nearest available operational category.

That is not a bug in the system. It is how markets often absorb novel transaction problems before a cleaner legal framework emerges. It also means that whoever can fit the dispute into an existing evidentiary bucket first may enjoy a large practical advantage.

A further consequence is that principals will inevitably seek guidance from their financial institutions, pulling banks into the center of these disputes. As one of the most consumer-facing participants in the transaction system, financial institutions may absorb increased support burdens, heightened dispute costs, and reputational risk, effectively becoming casualties of a proxy battle over loss allocation they did not design.

These Will Often Be Evidence Disputes Disguised as Payment Disputes

This is the point at which the logging question stops being a compliance afterthought and becomes legal infrastructure.

In many contested transactions, everyone in the stack will have some data, but no one will have the whole story in a form a court, payment intermediary, regulator, or counterparty can easily use. The user may have a rough memory and an account interface. The platform may have prompts, settings, execution traces, and internal rules. The merchant may have order data, checkout records, and fulfillment records. The payment intermediary may have authorization data, routing data, and dispute codes. None of those

records, standing alone, may answer the decisive question.

What will matter is whether the record can reconstruct the transaction in a way that is legally meaningful. Five specific gaps will be decisive in most contested agent transactions:

- **User instruction:** What did the user actually direct the agent to do, and how broadly or narrowly was that instruction stated? If the system cannot reproduce the original prompt or delegation in a form that connects to the specific transaction, the principal's account of what they intended becomes difficult to contest or confirm.
- **Permission state at execution:** What were the operative settings, thresholds, merchant restrictions, substitution logic, renewal parameters, and escalation rules at the exact moment the agent acted? Permissions configured days or weeks earlier may not reflect the principal's intent at the time of the disputed transaction.
- **Agent decision path:** What did the agent actually do, and why did it choose that path rather than an available alternative? The ability to reconstruct agent reasoning, not just agent action, will be decisive in misauthorized-commerce disputes where the contested question is whether the system exercised judgment beyond its delegated scope.
- **Merchant terms at execution:** What terms did the merchant present, when were they presented, through what interface or mechanism, and to whom – the agent, the platform, or the user? A generic reference to what the merchant's website usually says will not answer the question if terms were dynamic, versioned, or changed between the user's delegation and the agent's execution.
- **Payment authorization record:** What payment authorization occurred, on what basis, with what transaction metadata, and through which credential or token? The authorization record is often the most complete piece of the evidentiary puzzle, and also the most likely to be invoked first in a payment dispute framed through existing chargeback frameworks.

If the answer to those questions is missing, incomplete, or locked in a format no counterparty can use, the dispute is not just about law. It is about missing proof. And in many cases, the party with the better record will not merely win – it will control how the dispute is framed, which doctrinal category it lands in, and whether the fight even gets to a court.

Product Design Is Quietly Allocating Risk Already

Companies are effectively making loss-allocation decisions whether they say so or not. They make them through product design.

A platform that allows silent purchasing below a threshold is making a risk-allocation choice. A merchant that accepts agent-generated orders through one channel but not another is making a risk-allocation choice. A payment intermediary that treats certain transactions as presumptively authorized is making a risk-allocation choice. A system that requires renewed confirmation for new merchants, unusual categories, material substitutions, or recurring commitments is making a risk-allocation choice.

These choices may not be framed that way internally. They may appear as user-experience decisions, conversion decisions, or workflow-optimization decisions. In practice, they determine who is most likely to bear the cost when an agent-mediated transaction later goes bad. They also determine what record will exist when someone has to explain the failure.

That is why loss allocation cannot be treated as a purely downstream litigation problem. It is being decided at the design stage.

The system that optimizes hardest for speed and low-friction completion may also be choosing to make principals absorb more risk unless it creates compensating controls elsewhere. The merchant that insists on clearer checkout visibility or narrower acceptance conditions may be trying to avoid becoming the insurer of platform-side ambiguity. The payment intermediary that tightens review standards may be signaling that it will not serve as the universal absorber of delegation failures. The platform that declines to retain intelligible execution records may be making a quieter but

equally consequential decision about who will lose when a dispute becomes a proof contest.

Where the Burden Lands: The Factors That Actually Predict It

The burden will not land where the theory is cleanest. It will land where the market decides it belongs, and that decision will be driven by a recognizable set of practical factors, not abstract doctrine.

Three factors tend to predict where loss settles in novel transaction disputes, and all three apply directly to agentic commerce.

The first is control over the delegation model. The party that designed the permission architecture, set the thresholds, structured the agent's decision envelope, and controlled the logic that converted a broad user instruction into a specific purchase may be more or less likely to bear loss when that architecture produces a contested result. This factor cuts both ways. Platforms that designed well, with clear user disclosures, granular permission controls, and legible escalation logic, are in a stronger position than those that encouraged broad delegation while disclaiming responsibility for its outputs. Some platforms may also affirmatively choose to absorb a capped level of agent-error liability as a market-development strategy, effectively standing behind their agent in exchange for building user trust and driving adoption. The design choices, and what was disclosed about them, are what courts and intermediaries will examine.

The second is the quality of the record. In the absence of clear contractual allocation or settled doctrine, courts and payment intermediaries will rely heavily on who can show what happened. The party with the most complete, legible, and legally usable transaction record will be best positioned to frame the dispute, resist chargeback, and survive litigation. The party that cannot reconstruct the authority chain, the agent's decision path, or the operative terms at execution will bear more risk than its contractual position might suggest.

The third is reasonable reliance. Merchants that accepted transactions through channels that obscured delegation ambiguity, or that relied on platform representations without understanding their limits, face a harder reliance argument than merchants that required clearer agent identification, imposed transaction-level controls, or preserved their own evidence of what signals they received. The quality of reliance (not just the fact of it) will matter.

Early market practice will likely be shaped by a mix of contract design, network rules, platform terms, merchant acceptance decisions, payment categories, and the practical ability of each party to explain what happened. But those factors are not random. They reward the parties that built the right architecture from the start.

The Cases That Matter Most Will Feel Ordinary

The cases that matter most are not likely to be the most futuristic ones. They will often be mundane and commercially familiar.

A household agent reorders groceries from the wrong merchant at a much higher price. A business procurement agent buys from a vendor that was outside policy but within the system's general category rules. A travel agent books a nonrefundable fare through a valid platform flow even though the user expected more flexibility. An embedded shopping assistant accepts a recurring-delivery structure the principal never agreed to. A replenishment agent buys a substitute that seems commercially reasonable to the system but materially wrong to the buyer.

In each case, the transaction may be real, the system may be genuine, and the loss may still be very much in dispute. That is why the legal and commercial problem is not well captured by asking whether AI purchases are simply authorized or unauthorized. The better question is who should bear the cost when a real delegated system produces a result the market no longer agrees how to characterize – and what record exists to decide the point.

Courts Will Apply Familiar Doctrine to Unfamiliar Facts

Courts are unlikely to invent a wholly new law of loss allocation just because AI agents are involved. The more likely path is that disputes arrive through familiar categories: unauthorized transactions, agency disputes, contractual risk-allocation clauses, restitution claims, consumer-protection theories, merchant-chargeback fights, and arguments about reasonable reliance.

That means businesses should resist the temptation to imagine that novelty alone will drive the result. In misauthorized-commerce disputes specifically, three questions are likely to be most decisive:

Who controlled the delegation architecture, and what did that control include? Courts will look hard at the party that designed the permission model, set the agent's decision envelope, and controlled the logic that turned a broad instruction into a specific purchase. That party will carry more of the legal and commercial burden when the architecture produces a contested result than any platform disclaimer is likely to shift.

What record exists, and who built it? In the absence of clear contractual allocation or settled doctrine, the evidentiary record will often be dispositive. The party that cannot reconstruct the authority chain, the operative terms at execution, or the agent's decision path will be in a structurally weaker position regardless of its formal legal arguments.

What did each party actually know, and what could they reasonably have known? Foreseeability and reasonable reliance will carry significant weight. Merchants that relied on thin signals, platforms that encouraged reliance without disclosing the limits of their delegation model, and intermediaries that processed transactions without adequate metadata will face harder questions than parties that built legible, appropriately cautious systems.

The law is not starting from zero. But the operational discipline required to make those doctrines work in the agentic context is higher than many current implementations are designed to support.

What Companies and Financial Institutions Should Be Doing Now

For Companies Building, Enabling, or Confronting Agentic Commerce

Five concrete steps follow from the analysis above:

- **Classify your failure modes explicitly and design separate controls for each that are consistent with applicable network rules and consumer-protection regimes.** Fraud, compromised credentials, bad substitutions, wrong-merchant decisions, policy breaches, overbroad permissions, and contested renewals are different problems that warrant different responses. A single “authorization” framework that treats all of these as equivalent will produce both over-blocking and under-protection. Map your failure modes first, then build controls calibrated to each. Where payment networks or consumer-protection laws mandate specific classifications, disclosures, and remedies, product controls should account for those requirements.
- **Audit your product design for embedded risk-allocation decisions you have not made consciously.** Threshold rules, escalation logic, merchant whitelists, renewal defaults, substitution permissions, and silent-purchase settings are all loss-allocation choices in disguise. Conduct a deliberate audit of where those choices sit in your architecture, who bears the loss under each scenario, and whether that allocation reflects a decision you have actually made or a default you inherited.
- **Be explicit with merchants about what your platform’s trust signals actually cover.** If your platform wants merchants to rely on agent-generated orders, define in writing what sits behind that reliance (i.e., what credentials, permissions, and controls are in place) and what happens when a transaction is later disputed. Those representations should be aligned with applicable network rules and should not imply protections or liability shifts that conflict with statutory consumer rights or network dispute frameworks. Platforms that encourage merchant reliance without disclosing the limits of their delegation model are setting up loss-allocation fights they will be poorly positioned to resolve.

- **Build the transaction record to answer the five decisive questions now (not after the first dispute).** The user's instruction, the permission state at execution, the agent's decision path, the merchant's terms at execution, and the payment authorization record are the five evidentiary points that will determine most misauthorized-commerce disputes. Design your logging architecture to capture all five in a form that is legible, retrievable, and legally usable. Operational logs are not sufficient. Evidence-quality records are.
- **Manage and understand how allocation of liability is contemplated in agreements with entities in your stack, recognizing that certain allocation rules are fixed by statute or network rules rather than private contract.** Identify the initial loss-bearing party for each failure category and whether loss-allocation is dictated by applicable law or network rules, ensuring that allocation is reflected in your platform agreements, merchant terms, and API contracts, and verify that your insurance and indemnification structure is consistent with the allocation you intend. If the answer is "I do not know," the market will decide for you later, and under worse conditions.

For Financial Institutions

Financial institutions sit at a different point in the agentic commerce stack and face obligations that cannot be delegated away. As agent-mediated transactions scale, banks should focus on three immediate priorities.

- **Treat agent disputes as operational-risk events.** Agent-initiated disputes will surface first through Reg E error claims, Reg Z disputes, and chargebacks. Banks should ensure their dispute operations, reason-code selection, and provisional credit processes are calibrated to handle contested delegation scenarios without forcing them prematurely into ill-fitting categories such as fraud or user error. This requires advance coordination between legal, compliance, and operations teams before the first wave of disputes arrives.

- **Audit how agent-originated transactions are identified, coded, and recorded.** Banks should assess whether their systems can distinguish agent-initiated transactions from traditional card or credential-directed activity, and whether they are retaining evidence sufficient to satisfy regulatory error-resolution requirements and network rules. Where identification is not possible today, banks should anticipate greater ambiguity and cost, and proactively determine how to avoid becoming the residual risk bearer for agent-initiated disputes they did not create or control (to the extent possible through contract, reserve, and indemnification arrangements).
- **Align partner and platform agreements with non-delegable obligations.** Banks cannot contract around Reg E, Reg Z, or network rules; however, they can contract around reimbursement. Financial institutions should review sponsorship, BIN, and program-manager agreements to ensure that loss reallocation, indemnities, reserves, and insurance align with the reality that the bank may be the initial loss-bearing party in disputed agent transactions. Agreements that are silent on agent-initiated disputes leave the bank in the weakest possible position when the first wave of chargebacks arrives.

Taken together, these steps allow banks to navigate early agentic-commerce disputes proactively without becoming the default backstop for risks they neither designed nor controlled.

Bottom Line

Loss allocation in agentic commerce is where the transaction's unresolved instability becomes financially real.

Identity, authority, and assent remain foundational. But once money moves, goods ship, or services are delivered, the dispute can no longer be managed as a theoretical problem about how the system works. The market has to decide who bears the cost when the system worked well enough to transact and not well enough to command agreement afterward.

The hardest fights will not always be about obvious fraud. They will often be about real delegated systems producing contested results (misauthorized commerce) and about merchants, platforms, payment intermediaries, and principals all offering plausible but incomplete reasons why someone else should bear the loss.

In many of those fights, the deciding issue will not be a grand doctrinal insight. It will be whether anyone built a transaction record capable of showing what happened, what authority existed, what terms governed, and why the system acted as it did.

The parties that build that record now, before the disputes arrive, will be in a fundamentally different position than those that reconstruct it after the fact.

This article was prepared with the assistance of generative AI tools. The analysis, conclusions, and legal positions are the authors' own.

Agentic Commerce Series

- [Part 1 — AI Agents Are Starting to Act Inside the Transaction, and Commerce Law Is Not Ready](#)
- [Part 2 — The Identity Problem: Authentication, Fraud, and Who's Actually Buying](#)
- [Part 3 — The Authority Problem: When Does an Authorized Agent Become an Unauthorized Buyer?](#)
- [Part 4 — Contracting by Agent: When the Agent Clicks, Who Assents?](#)
- [Part 5 — When the AI Purchase Goes Wrong, Who Pays, and Who Can Prove It?](#)
- [Part 6 — The Gatekeepers of Agentic Commerce](#)
- [Part 6 \(A\) — The Payment Infrastructure Layer: How Network Rules Already Govern Agentic Commerce](#)
- [Part 7 — Agentic Commerce, Price Control, and the New Antitrust Questions](#)
- [Part 8 — What Epic Reveals About Agentic Commerce](#)

© 2026 FBT Gibbons LLP. May be considered attorney advertising in some jurisdictions.