



Apr 14, 2026

Categories:

[Agentic Commerce](#)

[Series](#)

[Publications](#)

Authors:

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

Contracting by Agent: When the Agent Clicks, Who Assents?

By the time lawyers get to contract formation in agentic commerce, much of the hard work has already happened somewhere else.

That is worth stating plainly because this is the point in the series most likely to be overstated. If the system cannot identify who acted in a legally meaningful way, nor explain whether the principal authorized the agent to produce the transaction in question, contract-formation analysis is unstable from the start. In fact, many apparent formation problems in agentic commerce are really upstream problems of [identity](#) and [authority](#), as we posited earlier in this series exploring the [six fault lines of agentic commerce](#).

But that does not make contract formation disappear.

Once those upstream issues are reasonably addressed, the remaining questions are narrower but still important. What terms was the principal agreeing to? By what mechanism were those terms accepted? At what point did acceptance occur? And how much of the final transaction reflected delegated discretion rather than a legally meaningful manifestation of assent?

That is the formation problem in agentic commerce.

The point is not that existing doctrine suddenly has nothing to say. It is that the doctrine was built around recognizable human acts: browsing, clicking, reviewing, and checking out. Agentic systems do not eliminate those acts so much as spread them across time, interfaces, and decision layers. Once that happens, ordinary contract questions become harder to answer cleanly.

The Right Question Is Narrower Than It Appears

Put at a high enough level of abstraction, the whole issue can sound overblown. Courts do not need to decide whether a machine has consciousness, intent, or legal personhood in order to determine whether a transaction binds a human principal. Agency law, automated contracting, and electronic contracting already allow

legal consequences to flow from mechanisms that are not themselves human.

So, the broadest version of the question — whether an AI can form a contract — is not especially useful.

The better question is narrower and more practical: *When does a principal's prior delegation, combined with an agent's later conduct, amount to legally sufficient assent to the terms governing the resulting transaction?*

That question breaks into several concrete ones: What terms govern the transaction? Who presented those terms, and to whom? What act is being treated as acceptance? When, exactly, did the contract form? Did the agent stay within the decision envelope the principal actually delegated?

Those are not philosophical questions. They are ordinary contract questions made harder by agent-mediated transaction design.

Identity and Authority Do Most of the Work — But Not All of It

The previous two articles in this series set up the issue. In the [second article](#), we argued that identity is doing more legal work in agentic commerce than many systems admit — proving that a credentialed device or linked account acted is not the same as proving the human principal intended the specific transaction that resulted. The [third article](#) addressed authority: once the system can identify the actor, it still has to explain what the agent was authorized to do and whether the principal's delegation covers the specific merchant, item, price, substitute, bundle, or timing choice the agent ultimately produced.

In many contested transactions, those two issues will do most of the work. If identity is weak or authority is missing, contract formation may never become the real point of dispute.

But where identity and authority are relatively well specified, formation still matters because it asks a different question: not whether the agent could act in some general sense, but whether the transaction that resulted amounted to assent to a particular set of

terms in a legally recognizable way.

A user may authorize an agent to purchase a replacement filter, book travel within policy, or reorder office supplies below a threshold. A merchant may receive and process the transaction without hesitation. A payment intermediary may authorize the charge. And still a later dispute may turn on formation-layer questions the upstream analysis did not answer: which merchant terms governed, was arbitration incorporated, did warranty limitations become part of the deal, and was the operative acceptance the user's earlier configuration or the agent's later click?

That is why contract formation remains a distinct fault line even if it is not the largest one.

Assent Gets Harder to Locate Once the Human Act Is Distributed

Ordinary e-commerce compresses several legally important steps into a single recognizable event. The user reviews the product, sees the seller, encounters the checkout flow, is presented with terms in some fashion, clicks to complete the purchase, and receives confirmation. That sequence may still raise familiar disputes about clickwrap, browsewrap, incorporation, or conspicuousness, but at least the law knows where to look.

Agentic commerce disrupts that simplicity.

The user may authorize the agent days or weeks earlier. The instruction may be broad rather than transaction-specific. The actual merchant may be selected later by the system. Terms may appear only at execution. Price and delivery conditions may be dynamic. Substitution may occur if the preferred product is unavailable. The user may never see the final checkout page at all.

In that setting, assent does not disappear. It gets harder to locate. The path to contract formation may be spread across the user's initial delegation, the user's standing settings and constraints, the platform's machine-readable permissions, the agent's execution of the transaction, the merchant's presentation of terms, and the

payment system's authorization event.

WHERE DID ASSENT OCCUR?

In traditional e-commerce, one moment carries all the legal weight.
In agentic commerce, assent is distributed across six events — and ambiguous at each one.



The harder it becomes to identify which of those moments carried legal significance, the harder it becomes to say with confidence what contract formed and on what terms.

Four Formation Questions That Survive the Identity and Authority Analysis

The formation problem in agentic commerce should not be inflated into a claim that doctrine collapses the moment an AI agent enters a checkout flow. Most of the familiar apparatus still applies. Courts will still ask whether terms were reasonably presented, whether assent was manifested through the relevant mechanism, whether a principal can be bound through an agent, and whether the record supports the counterparty's account of what happened.

Once identity and authority are accounted for, the residual formation risk tends to cluster around four questions.

1. Assent to What Terms?

In an agent-mediated transaction, several sets of terms may be in play at once: merchant website terms, marketplace terms, platform terms governing the agent's service, embedded-checkout terms, payment-provider terms, API rules, or fulfillment conditions incorporated by reference. In some flows, those layers may align neatly. In others, they may not.

Agentic commerce makes the stacking harder to see. The user may interact primarily with an agent interface while the legal terms that matter most sit downstream at the merchant or payment layer. A platform may present one set of assumptions about what the agent can do while the merchant's checkout terms impose different restrictions or allocations of risk. In a machine-to-machine order flow, key terms may be incorporated through the commercial architecture itself rather than through a familiar human-facing page.

When disputes arise, one of the first questions will be which terms actually govern and whether those terms entered the transaction in a way contract law will recognize.

The terms question also has a second dimension that agentic commerce makes more acute: the procedural unconscionability and contracts of adhesion arguments that have historically struggled in ordinary e-commerce litigation may find more traction here. Courts have generally been reluctant to invalidate standard-form online terms on adhesion or unconscionability grounds, in part because the standard defense holds that the user had a meaningful opportunity to review the terms and chose to proceed. That defense weakens considerably when the agent — not the user — traversed the checkout flow. The user may have had no realistic opportunity to review the operative terms, not because they were hidden but because the user was structurally absent from the interface where those terms appeared.

That is a procedurally stronger unconscionability argument than what exists in ordinary clickwrap cases. Courts will still be reluctant

to void standard terms wholesale on this basis, and the substantive terms themselves have not become more one-sided simply because an agent accepted them. But the procedural posture is meaningfully different when a third system accepted terms the principal never had a realistic opportunity to evaluate. Companies relying on standard checkout terms to govern agent-initiated transactions should treat that vulnerability as a real design consideration, not a theoretical one.

2. Assent by What Mechanism?

In ordinary online contracting, the law usually has a visible candidate: the user clicked the button, checked the box, or continued using the service after receiving notice. The adequacy of that mechanism may still be contested, but at least the act is recognizable.

Agentic systems complicate that analysis because the operative act may be the user's initial instruction, a standing authorization setting, an in-chat confirmation, an automated click executed by the agent, an API call carrying a transaction request, or a payment-authorization event. Not all of those acts prove the same thing. A standing instruction may show broad permission to transact in a category. An automated click may show execution of a downstream purchase event. A payment authorization may show that the rails recognized the transaction. None automatically answers the separate question of whether the principal manifested assent to the full set of governing terms through that mechanism.

That distinction will matter most where the acceptance ritual is no longer obviously human-facing.

3. Assent When?

In agentic commerce, the legally meaningful moment of contract formation may not line up neatly with the moment the user last interacted with the system. The user may delegate authority long before the actual merchant is selected. The agent may monitor prices or inventory conditions and execute only once specified triggers are met. Terms may change between the original instruction

and the final transaction. A merchant may not accept until fulfillment or shipment. A platform may treat the transaction as complete once the agent places the order even if the merchant later rejects or modifies it.

When did acceptance occur? When did the operative terms become fixed? At what point did the principal become bound? If the user delegated at one moment and the agent transacted under materially different downstream conditions later, which moment matters most? These are ordinary contract questions. Agentic flows just make them less tidy.

4. Assent With How Much Discretion?

The fourth issue sits closest to the boundary between authority and formation. The more precisely the user defines the transaction, the easier it is to treat the agent's later conduct as implementation. The more discretion the agent has to choose among merchants, products, delivery options, substitutes, bundles, warranties, or timing choices, the harder it becomes to say that the final act cleanly reflects assent to the exact transaction that occurred.

This is not just an authority question because the same discretionary move that complicates scope may also complicate assent to terms. An agent instructed to book a flight within policy may encounter materially different fare conditions, refund rights, or arbitration clauses depending on carrier and booking channel. A household shopping agent may buy a functionally equivalent replacement product whose terms convert the purchase into a recurring subscription the principal never reviewed.

At that point, the question is not only whether the agent was authorized to choose. It is whether the principal should be treated as having assented to the legal package attached to the chosen result.

Checkout Architecture Is Also Enforceability Architecture

The market tends to describe the difference between in-chat checkout, embedded checkout, redirect flows, and direct merchant-native checkout as a product or conversion issue. In agentic commerce, it is also an enforceability issue.

If the transaction occurs entirely within an agent or platform interface, the merchant may have less direct control over how terms are presented and less visibility into what the user or principal actually saw. If the user or agent is redirected into a merchant-controlled checkout flow, the merchant may have a stronger argument that its terms were presented through a familiar contracting mechanism — even if the human principal never personally reviewed the page. If the transaction occurs through an API or machine-to-machine ordering environment, the system may have strong evidence of execution and weak evidence of anything resembling a traditional assent ritual.

None of those models is automatically unenforceable. The point is narrower: different checkout architectures generate different evidentiary and doctrinal postures. A business deciding where the transaction should occur is therefore deciding, at least in part, what its future formation story will look like.

That connection also runs forward into loss allocation. Formation ambiguity does not just create doctrinal uncertainty in the abstract — it determines who is exposed when the transaction is later contested. A weak formation story is not merely a legal problem. It is a loss allocation problem. The party least able to reconstruct what terms governed, how they entered the transaction, and why the principal should be treated as bound will bear more of the commercial consequence when disputes arise. Companies that treat checkout architecture as a product decision while leaving the formation story unexamined are making a risk allocation decision without knowing it.

The Hard Cases Will Be Ordinary

The contract cases that matter most are unlikely to be the most futuristic ones.

A user authorized the agent to book travel, but the fare selected was nonrefundable and carried terms the user never reviewed.

A household shopping agent bought a functionally equivalent replacement product, but the merchant's terms converted the purchase into a recurring subscription.

An enterprise procurement agent placed a routine order through an API integration, but the supplier's current terms contained changed limitation-of-liability language that no human reviewed during the transaction.

Those are not problems of machine consciousness. They are problems of how conventional contract doctrine interacts with automated choice, layered terms, and distributed execution.

That is also why the "much ado about nothing" critique only goes so far. It is right that identity and authority do much of the work. It is wrong if it assumes that nothing formation-specific remains once those issues are reasonably stable.

Courts Will Apply Familiar Doctrine to Unfamiliar Facts

Courts are unlikely to want a special body of AI contract law merely because agents are mediating the flow. The more likely path is familiar doctrine applied to unfamiliar facts. Judges will still ask ordinary questions about notice, manifestation, incorporation, agency, and reasonableness — then fit the technology into those frameworks as best they can.

That cuts both ways.

The warning is that businesses should not expect courts to rescue poor transaction design by declaring that agent-mediated commerce requires radically new formation rules.

The opportunity is that businesses already know much of what helps ordinary digital contracts hold up: clear presentation of terms, legible records of what was shown and when, evidence of the authority chain, bounded discretion, and decision points that distinguish user instruction from system inference.

The law is not starting from zero. But the operational discipline needed to make those doctrines work in agentic environments is higher than many current implementations seem prepared to support.

What Companies Should Be Doing Now

Companies building, enabling, or confronting agentic commerce should not assume that identity and authority alone solve their contracting risk. Five things follow from that recognition:

- **Map where operative terms are actually presented in the flow and verify that the path is traceable.** If the terms that matter most appear downstream in a merchant or payment flow that neither the principal nor the agent traversed in a documented way, the formation story is weak regardless of how clean the authentication record looks. Fix the gap in the architecture before a dispute requires explaining it.
- **Designate one event as the operative act of assent and design the system around it.** Whether that is an initial user configuration, a transaction-specific approval, an automated click, an API request, or a payment event, the choice should be deliberate and documented. A system that cannot identify which act was supposed to carry legal significance will not be able to defend that choice later.
- **Preserve the actual terms in effect at the moment of execution — not a generic reference to what the website usually says.** If terms are dynamic, versioned, personalized, or changeable by the merchant, the business should snapshot and timestamp the operative terms at the moment the transaction closes. That record is the contract.
- **Set explicit limits on agent discretion and build confirmation back into high-risk decision points.** The more discretion the agent exercises without escalation or renewed confirmation, the harder it becomes to tie the final transaction to something the principal actually assented to. Discretion should be commercially useful and legally bounded — and

that boundary should be visible in the transaction record.

- **Build the transaction log to distinguish what the user instructed, what the system inferred, what the merchant presented, and what the agent ultimately accepted.** If those layers collapse into a single opaque event, the business will struggle to reconstruct its own formation story when a dispute requires one. The log is not just an audit tool — it is the evidentiary foundation of the contract itself.

Bottom Line

Contract formation in agentic commerce is narrower than early commentary may suggest, but it is not trivial.

Identity and authority do much of the work. In many disputes, they will do most of it. But once those questions are reasonably addressed, formation still matters because the law still has to decide what contract formed, on what terms, through what mechanism, and at what point in the agent-mediated flow.

AI agents do not make contract doctrine obsolete. They make ordinary formation questions harder to answer cleanly by distributing human commercial action across settings, permissions, interfaces, and execution events. A weak formation story is not merely a legal problem — it is a loss allocation problem. The businesses that recognize that early, and design their transaction architecture accordingly, will be in a better position than those that assume delegated authority alone completes the legal analysis.

And once the system can tell a coherent story about identity, authority, and assent, the next question becomes unavoidable: when the transaction still goes wrong, who bears the loss?

Please contact the author if you have questions or comments on this article. You can also reach out to any member of the firm's **Data, Digital Assets & Technology** practice group for help navigating AI deployments and governance.

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