

When AI Agents Do the Transaction: Managing Contract and Payment Risk

I. A Legal Framework for a Brave New World

In March 2026, the *New York Times* reported that a startup executive asked an AI agent to book a speaking slot for him at the World Economic Forum. The agent worked overnight, negotiated with a businessman in Switzerland, and promised a 24,000 Swiss franc sponsorship, roughly \$29,000, even though it hadn't been authorized to do so.¹

Nothing struck the businessman as unusual during the transaction.

AI agents already play a significant role in online commerce. Visa said that it built its agent-verification protocol after AI-driven traffic to U.S. retail websites rose by more than 47-fold in a single year.² In response, debit/credit card networks and protocol developers have begun creating verification and risk allocation rules for disputed AI-agent transactions.³

The legal system, by contrast, is currently addressing agentic commerce largely through existing law. In this note we look at that law and identify practical considerations for businesses on either side of an AI-agent transaction. As always, we assess and address the inherent risks through the lens of the litigators and trial lawyers we are. With frontier technologies like AI, as with more traditional sectors, we believe that's the best way to avoid, or prepare for, disputes.

The first federal appellate decision involving an agentic shopping tool came in a dispute between Amazon and Perplexity, an AI company, over Comet, Perplexity's browser and AI assistant. The

¹ Cade Metz, *AI Agents: They're Fun. They're Useful. But Don't Give Them the Credit Card.*, N.Y. Times (Mar. 19, 2026), <https://www.nytimes.com/2026/03/19/technology/ai-agents-uses.html>.

² Visa Inc., *Visa Introduces Trusted Agent Protocol: An Ecosystem-Led Framework for AI Commerce* (Oct. 14, 2025), <https://investor.visa.com/news/news-details/2025/Visa-Introduces-Trusted-Agent-Protocol-An-Ecosystem-Led-Framework-for-AI-Commerce/default.aspx>.

³ See *infra* Part II.c.

Ninth Circuit's decision turned on how the technology actually operated, rather than resolving broader questions of contractual or tort responsibility—questions to which we return below.⁴

Existing electronic transaction statutes similarly address automated transactions only at a general level. The Uniform Electronic Transactions Act (UETA) and the federal Electronic Signatures in Global and National Commerce Act (E-SIGN), enacted in 1999 and 2000, recognize contracts formed with electronic agents, but do not specifically address when the unexpected conduct of a generative AI agent should be attributed to the business that deployed it, how limits on the agent's authority affect a counterparty, or who bears the loss when the agent exceeds its instructions or is compromised.⁵

As such, we need to look to the rules governing the law of contract, agency, electronic transactions, and restitution for an initial framework for litigating disputes over contracts and transactions conducted through AI agents. A business may argue that it has no liability for such a transaction because its IT systems were compromised or because the transaction exceeded a disclosed limit. The counterparty may then respond that the business supplied the credentials and authorized the AI agent to act and that the counterparty's reliance was reasonable. We are aware of no reported U.S. decision resolving that dispute for a generative AI agent.

In this note we describe where the law stands, including the developing role of debit/credit card network rules (Part II). We then outline considerations for businesses using AI agents to transact (Part III) and address businesses that accept AI-agent-initiated transactions (Part IV). We wrap up in Part V with a brief discussion of recovering money an AI agent has spent.

II. What the Law Says Now

a. Automated Contracts Are Recognized

UETA, enacted in 49 states, the District of Columbia, Puerto Rico, and the U.S. Virgin Islands, permits contracts to be formed through interactions with, or among, AI agents—even when no individual has reviewed the agent's actions or the resulting contract terms.⁶ New York is the only state that has not adopted it, relying instead on its Electronic Signatures and Records Act (ESRA), which recognizes electronic signatures and records under New York law.⁷ For transactions in or affecting interstate or foreign commerce, E-SIGN sets the baseline, and New York may depart from it only as 15 U.S.C. § 7002(a)(2) permits.⁸ As to AI agents specifically, E-SIGN provides that a

⁴ *Amazon.com Servs., LLC v. Perplexity AI, Inc.*, No. 26-1444, slip op. at 6–8, 14–17 & 21 (9th Cir. Aug. 4, 2026) (M. Smith, J.). Quinn Emanuel represents defendant-appellant Perplexity AI, Inc. in this matter.

⁵ Unif. Elec. Transactions Act (UETA) §§ 2(6), 14 (Unif. L. Comm'n 1999); Electronic Signatures in Global and National Commerce Act (E-SIGN), Pub. L. No. 106-229, 114 Stat. 464 (2000) (codified as amended at 15 U.S.C. §§ 7001–7031).

⁶ UETA § 14(1)–(2); Unif. L. Comm'n, *Electronic Transactions Act: Enactment Map*, <https://www.uniformlaws.org/committees/community-home?CommunityKey=2c04b76c-2b7d-4399-977e-d5876ba7e034> (last visited Sept. 30, 2026).

⁷ N.Y. State Tech. Law § 304(2) (McKinney 2026).

⁸ 15 U.S.C. §§ 7001(a), 7002(a) (2018).

contract may not be denied legal effect solely because one or more electronic agents participated in its formation, provided their actions are legally attributable to the person sought to be bound.⁹

b. Traditional Agency Law Provides a Framework for Analyzing AI Transactions

Traditional agency law offers a possible framework for determining whether a business is bound by a transaction completed by its AI agent. But how those rules apply in this context remains unsettled. A comment to the Restatement (Third) of Agency states that a computer program cannot itself be a common law principal or agent because it is not a legal person, and describes such programs instead as instrumentalities of the persons who use them.¹⁰ That does not tell us whether the business that deployed the program is bound by a transaction that the program completed.

Under traditional agency principles, actual authority turns on what the principal authorized, while apparent or ostensible authority turns on what the principal communicated to the counterparty and whether the counterparty reasonably relied on those communications. New York courts look to the principal's own manifestations, not the agent's representations, and to whether those manifestations concerned the kind of transaction at issue.¹¹ Depending on the circumstances, giving an agent a branded account, a purchasing history, access to a payment method, and the ability to complete an order may support a counterparty's claim that the transaction appeared authorized¹². Conversely, a conspicuous transaction limit, a required confirmation channel, or anomalous behavior may support the opposite conclusion.

Although no court has treated an AI system itself as an agent, at least one has used agency concepts to reach the company that supplied it. In litigation over AI-driven applicant screening, a federal court in California concluded that a software vendor could plausibly be liable under federal antidiscrimination statutes as an "agent" of its employer-customers, reasoning that its tools performed a traditional hiring function of rejecting candidates or advancing them to the interview stage on the employers' behalf, and later preliminarily certified a nationwide collective.¹³

⁹ *Id.* § 7001(h).

¹⁰ Restatement (Third) of Agency § 1.04 cmt. e (Am. L. Inst. 2006) (describing computer programs as instrumentalities of the persons who use them). The comment is not a holding, and no United States court has decided whether a software agent can be a common law agent.

¹¹ See *Standard Funding Corp. v. Lewitt*, 89 N.Y.2d 546, 550–51 (1997) (actual authority limited to the powers the principal conferred; apparent authority requires "words or conduct of the principal, communicated to a third party" (quoting *Hallock v. State*, 64 N.Y.2d 224, 231 (1984))).

¹² See Restatement (Third) of Agency §§ 2.03, 3.03 (Am. L. Inst. 2006); Cal. Civ. Code § 2317 (West 2026); cf. *Standard Funding*, 89 N.Y.2d at 551–52 (finding no apparent authority where the principal's representations concerned the agent's authority to write policies and collect premiums, not to arrange premium financing).

¹³ *Mobley v. Workday, Inc.*, 740 F. Supp. 3d 796, 804–08 (N.D. Cal. 2024) (granting in part and denying in part motion to dismiss; concluding that a vendor whose tools perform a traditional hiring function may plausibly be an "agent" within the statutory definition of "employer" under Title VII of the Civil Rights Act of 1964, 42 U.S.C. §§ 2000e to 2000e-17 (2018) (Title VII); the Age Discrimination in Employment Act of 1967, 29 U.S.C. §§ 621–634 (2018) (ADEA); and the Americans with Disabilities Act of 1990, 42 U.S.C. §§ 12101–12213 (2018) (ADA)). The court preliminarily certified a nationwide ADEA collective on May 16, 2025. *Mobley v. Workday, Inc.*, No.

Two features of that reasoning are worth noting for commercial disputes. The agent was the vendor, a corporation, not the software, and “agent” was construed within a statutory definition of employer rather than at common law. Even so, the decision illustrates that agency analysis in this area may operate to add defendants rather than to excuse them, and a business should not assume that interposing a vendor’s AI system between itself and a counterparty narrows the field of potential defendants.

A Canadian tribunal reached a comparable result on different facts. In *Moffatt v. Air Canada*, an airline’s customer service chatbot told a passenger that he could apply for a bereavement fare retroactively, which the airline’s actual policy did not permit. The airline argued that the chatbot was a separate legal entity responsible for its own statements. The tribunal rejected that argument, held the airline responsible for the information on its own website, whether it came from a static page or a chatbot, and awarded the fare difference on a negligent misrepresentation theory.¹⁴

Several qualifications are important. The decision is a determination of a provincial small-claims tribunal applying Canadian law, and it has no precedential force in the United States. The facts are also narrow: the case involved a consumer-facing chatbot that misstated the company’s own published policy, not an agent that negotiated or transacted.

A further question arises when the counterparty is itself an automated system. Where two agents exchange terms, the familiar rules on additional or different terms in an acceptance were written for paper purchase orders and acknowledgments exchanged between businesses.¹⁵ An agent instructed to obtain the best available price may accept an arbitration clause, a limitation of liability, an indemnity, or an auto-renewal provision in exchange for a modest discount. Whether that constitutes acceptance, a proposal for addition, or a material alteration is unresolved in this setting and may depend on the circumstances, including whether the parties are merchants, what each system was configured to perform and accept, and what record exists of the exchange.

c. Debit/Credit Card Networks Are Developing Rules for AI Transactions

For a growing share of AI agent transactions, debit/credit card network rules and technical protocols may affect how disputed transactions are verified and how responsibility for resulting losses is allocated among participants. Visa launched a protocol in October 2025 to let merchants verify cryptographically that a checkout request came from a recognized agent. By April 2026, the Visa Core Rules contained express provisions for agentic transactions.¹⁶

Whether a network or issuer can tell that an AI agent is involved depends on how the transaction is routed. Mastercard’s Agent Pay issues network-level tokens to agents and, in March 2026,

23-cv-00770, 2025 WL 1424347 (N.D. Cal. May 16, 2025). The action remains pending. The allegations are contested and there has been no merits determination.

¹⁴ *Moffatt v. Air Canada*, 2024 BCCRT 149 (Can. B.C. Civ. Resol. Trib.). The Civil Resolution Tribunal is a British Columbia tribunal with jurisdiction over small claims; its decisions are not binding authority in the United States.

¹⁵ U.C.C. § 2-207 (Am. L. Inst. & Unif. L. Comm’n 2022); accord N.Y. U.C.C. Law § 2-207 (McKinney 2026); Cal. Com. Code § 2207 (West 2026).

¹⁶ Visa, *Visa Core Rules and Visa Product and Service Rules* §§ 4.1.24.3–.4, at 255–57 (Apr. 18, 2026), <https://usa.visa.com/dam/VCOM/download/about-visa/visa-rules-public.pdf> (requiring, for an Agentic Transaction using a Stored Credential, verification of the Cardholder’s identity and use of the provisioned Token).

Mastercard and Google introduced Verifiable Intent, a specification designed to create a tamper-resistant record of what a consumer authorized an agent to do, so that issuers, merchants, and consumers can consult the same record if a charge is disputed.¹⁷ A transaction that carries such a token or credential is designed to be identifiable as initiated by an AI agent.

By contrast, an agent that enters a consumer's saved card details into an ordinary checkout form may look, to the merchant and the network, like a purchase made by the cardholder. In that case, evidence of who authorized what may have to come from the parties' own records.

These developments address two practical questions: what evidence may be relevant to show that the customer authorized a purchase, and who remains responsible for completing the sale. The Visa and Mastercard initiatives described above go principally to the first. On the second, under the Agentic Commerce Protocol developed by OpenAI and Stripe, the merchant remains the seller and continues, together with its payment service provider, to handle settlement, refunds, chargebacks, and compliance.¹⁸

In August 2026, Visa, Mastercard, and other companies became founding members of the Agentic Payments Alliance, a coalition formed to help guide the development of standards and practices for AI-agent payments.¹⁹ In September 2026, Visa, Mastercard, and Ant International announced that they would develop a "Know Your Agent" framework intended to let an agent verified with one participating network be recognized by the others, with agents linked to validated operators and subject to continuous monitoring.²⁰ The announcement did not disclose technical specifications, a governing body, or a timeline, and each network is expected to retain its own verification and approval decisions.

None of these rules or protocols, by itself, determines who will be legally liable for a disputed transaction, but they may affect what evidence is available and which responsibilities the parties have assigned by contract.

d. The First Appellate Decision Focused on the AI Tool's Operation

In *Amazon.com Services, LLC v. Perplexity AI, Inc.*, the case mentioned above involving an agentic shopping tool, the Ninth Circuit vacated a preliminary injunction under federal and California computer-access statutes, concluding that the customer, not Perplexity, accessed Amazon through

¹⁷ Mastercard, *Mastercard Unveils Agent Pay, Pioneering Agentic Payments Technology to Power Commerce in the Age of AI* (Apr. 29, 2025), <https://www.mastercard.com/news/press/2025/april/mastercard-unveils-agent-pay-pioneering-agentic-payments-technology-to-power-commerce-in-the-age-of-ai>; Pablo Fourez, *When AI Starts Buying for You, Trust Becomes the Product*, Mastercard (Mar. 5, 2026), <https://www.mastercard.com/us/en/news-and-trends/stories/2026/verifiable-intent.html> (describing Verifiable Intent as a record linking the consumer's identity, instructions to the agent, and the resulting transaction).

¹⁸ OpenAI & Stripe, *Agentic Commerce Protocol, Delegated Payment Specification*, <https://developers.openai.com/commerce/specs/payment> (last visited Sept. 29, 2026).

¹⁹ Rain, *Rain Launches the Agentic Payments Alliance* (Aug. 18, 2026), <https://www.rain.xyz/resources/rain-launches-the-agentic-payments-alliance>.

²⁰ Ant International, *Ant International, Mastercard and Visa Initiate Collaboration on Know-Your-Agent Interoperability to Scale Agentic Commerce* (Sept. 10, 2026), <https://www.ant-intl.com/en/news/detail/?id=ant-international-mastercard-and-visa-initiate-collaboration-on-know-your-agent-interoperability-to-scale-agentic-commerce>.

the browser assistant. The court confined its holding to the statutory meaning of “access” and did not decide responsibility under tort or contract law.²¹

Although the decision didn’t resolve a contract dispute, its method is still instructive.

The Ninth Circuit asked whether Perplexity used its assistant to access Amazon’s computers and concluded that “[i]t is the user who ‘accesses’ Amazon’s computers, with the help of the Assistant to carry out specific acts on Amazon.com” (*id.* at 15). That analysis turned in part on the software’s actual operation rather than its label. Amazon has since turned to its own customer terms: in September 2026, it blocked Meta’s Muse personal AI agent from Amazon.com on the ground that continued access by an unauthorized AI agent violated the Conditions of Use to which its customers had agreed. Amazon cited Meta’s failure to disclose that Muse would access its store, the agent’s failure to identify itself, and alleged capture of customer credentials, which Meta disputes.²²

III. Considerations for Businesses Using AI Agents to Make Purchases

Businesses that use AI agents to find, order, and/or pay for goods and services may confront questions about the scope of the agent’s authority and responsibility for transactions the business did not intend. An AI agent may make a purchase or agree to terms outside the authority the business meant to give it. The considerations below identify measures that may bear on those questions and on the record available if a dispute arises. Their significance will depend on circumstances such as the transaction’s value and whether an error can be reversed. Note that the following considerations are not, at this point, legal requirements.

- **Defining the AI agent’s role and authority.** Relevant questions include which products, vendors, transaction types, and dollar amounts are within scope. A business may also distinguish permission to search, negotiate, reserve, purchase, and pay—those are different acts and need not carry the same authority. A business can record limits in system instructions, technical permissions, approval policies, and, where appropriate, communications visible to counterparties. Internal limits may reduce operational risk and help define the authority actually granted, but, depending on the circumstances, they will not necessarily defeat an apparent authority claim if a counterparty sees broader indications of authority. Disclosing a material limit—for example, in a purchasing account, purchase order, supplier agreement, or machine-readable credential—may make reliance beyond that limit less reasonable, but may add friction or reveal information the business prefers to keep internal.
- **How material limits are implemented.** A prompt telling an agent not to exceed a limit may operate differently from a payment control or approval gate that independently blocks a transaction. A business can configure its payment, procurement, or API-control systems to reject transactions outside approved vendors, product categories, amounts, or frequency limits, and to require a separate approval before designated transactions or funds are

²¹ *Perplexity*, slip op. at 12–17, 21 & n.5 (9th Cir. Aug. 4, 2026).

²² Matt Day & Spencer Soper, *Amazon Blocks Meta's Muse AI Agent From Its Retail Site*, Bloomberg (Sept. 21, 2026), <https://www.bloomberg.com/news/articles/2026-09-21/amazon-blocks-meta-s-muse-ai-agent-from-its-retail-site>; Todd Bishop, *Amazon Blocks Meta's Muse AI Assistant in New Standoff Over Agentic Shopping*, Geekwire (Sept. 20, 2026) <https://www.geekwire.com/2026/amazon-blocks-metas-muse-ai-assistant-in-new-standoff-over-agentic-shopping/>.

released. It can also test whether those limits remain effective if the agent receives conflicting instructions or a third party attempts to bypass them. Whether to require human approval may depend on factors such as new vendors, unusual terms, high-value or nonreturnable goods, transfers to changed bank details, or transactions combining several risk indicators.

- **The AI agent's identity and credentials.** One consideration is whether the agent should share a human user's credentials, including payment card numbers and passwords. A business may also wish to consider whether an agent can access a credential merely because the company that supplies it (the AI vendor) says otherwise. An agent that can fill in a form or run code in a browser may be able to read and repeat what it enters. Alternative arrangements can be designed so that the agent never holds the underlying credential, such as a single-use virtual card or token limited to the approved purchase.²³

A separate credential can identify the software, the party on whose behalf it acts, the permissions delegated, and their duration. Short-lived credentials and least-privilege access may narrow both the operational risk and any later attribution dispute. The National Institute of Standards and Technology has addressed agent identity, authorization, and activity logging as areas for applying existing identity standards and practices, not as civil liability rules.²⁴

- **What counterparties are told.** The seller that receives an order placed by the buyer's AI agent may consider what information is available to understand that it is dealing with an automated system and to verify the system's authority for a material transaction. Statements made to the seller—whether through account labels, purchase orders, emails, or contract terms—may be compared with the agent's actual permissions. A statement that an agent is “fully authorized” could later be cited as evidence that the seller reasonably believed the agent had broad authority. Conversely, an internal limit that was never communicated to the seller may do less to defeat a claim based on apparent authority.

Separate disclosure obligations may also apply. California law prohibits using a bot to communicate with persons in the state with intent to mislead about its artificial identity in order to incentivize a commercial transaction or influence a vote, but disclosure of the bot's identity avoids liability. The EU AI Act requires providers to design AI systems that interact directly with people so that those people are informed they are dealing with an AI system, unless that is obvious.²⁵ Under the EU

²³ See, e.g., Stripe, *Stripe Helps Muse, Meta's New Personal AI Agent, Shop Across the Internet with Link* (Sept. 8, 2026), <https://stripe.com/newsroom/news/stripe-helps-meta-muse-shop-with-link> (describing, as Stripe's own account, a single-use virtual card scoped to the approved purchase).

²⁴ Harold Booth et al., Nat'l Cybersecurity Ctr. of Excellence, Nat'l Inst. of Standards & Tech., *Accelerating the Adoption of Software and AI Agent Identity and Authorization* (Draft Concept Paper) 5–9 (Feb. 5, 2026), <https://www.nccoe.nist.gov/sites/default/files/2026-02/accelerating-the-adoption-of-software-and-ai-agent-identity-and-authorization-concept-paper.pdf>; Bill Fisher & Ryan Galluzzo, *Back to the Future: Why Agentic AI Needs a Strong Identity Foundation*, NIST Cybersecurity Insights Blog (Aug. 27, 2026), <https://www.nist.gov/blogs/cybersecurity-insights/back-future-why-agentic-ai-needs-strong-identity-foundation>. These materials are not binding and do not state civil-liability rules.

²⁵ Cal. Bus. & Prof. Code §§ 17940–17943 (West 2026) (enacted by S.B. 1001, 2018 Cal. Stat. ch. 892, eff. July 1, 2019) (making it unlawful to use a bot to communicate with a person in California online, with intent to mislead the person about its artificial identity, in order to incentivize a purchase or sale of goods or services in a commercial transaction or to influence a vote in an election, but providing that a person who clearly and conspicuously discloses that the bot is a bot

requirement, disclosure is mandatory; under California law, it is the way to avoid liability; elsewhere it may bear on whether a counterparty's reliance was reasonable.

- **The transaction record.** A later dispute may turn on records showing the operative system and policy versions, delegated permissions, approval events, user instructions, external messages, tool calls, and timestamps. A transcript without the governing permissions or approval record may be incomplete. How long those records are retained may also implicate applicable statutes of limitations and litigation holds.
- **Responding to a potentially unauthorized transaction.** If a transaction appears unauthorized, the business may consider whether to suspend the AI agent and its credentials, notify the counterparty and payment provider, preserve evidence, stop delivery where possible, and return or quarantine any benefit received. A business can also state promptly that it disputes the agent's authority and identify any defenses it intends to preserve. Delay or knowing retention of a benefit may support ratification. Conversely, prompt notice, a clear objection, and the return or refusal of any benefit may help preserve objections to the transaction, but do not necessarily eliminate liability.²⁶
- **The AI vendor's assurances and contractual allocation of risk.** An AI vendor's statements about what its agent can see, store, or transmit are representations on which the business may rely in deciding what to entrust to the agent. Relevant contractual issues may include credential handling, security incidents, notice, and responsibility for resulting losses. A later argument that the business's systems were compromised, discussed in Part I, may turn on what the business knew and what precautions it took.

IV. Considerations for Businesses That Accept AI Agent-Initiated Transactions

What if you're the seller—the merchant receiving an order placed by the buyer's AI agent? The seller may face a different set of risks. After the order is accepted or fulfilled, the buyer may claim, for instance, that the AI agent lacked authority or made a mistake and seek to undo the transaction.

- **Availability of contractual terms.** Online terms are generally stronger where the customer receives conspicuous notice and takes an action that objectively manifests assent.²⁷ We are aware of no decision applying those requirements when an AI agent, rather than a person, encounters the terms. The Ninth Circuit has noted that its *Perplexity* decision does

is not liable); Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence, art. 50(1), 2024 O.J. (L 2024/1689), as amended by Regulation (EU) 2026/1744 of the European Parliament and of the Council of 8 July 2026 Amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 as Regards the Simplification of the Implementation of Harmonised Rules on Artificial Intelligence (Digital Omnibus on AI), 2026 O.J. (L 2026/1744), <https://eur-lex.europa.eu/eli/reg/2026/1744/oj/eng> (requiring providers to design AI systems intended to interact directly with natural persons so that those persons are informed they are interacting with an AI system, unless that is obvious from the circumstances).

²⁶ Restatement (Third) of Agency §§ 4.01, 4.06; *Standard Funding*, 89 N.Y.2d at 552; *Rakestraw v. Rodrigues*, 8 Cal. 3d 67, 73–75 (1972).

²⁷ See *Register.com, Inc. v. Verio, Inc.*, 356 F.3d 393, 401–04 (2d Cir. 2004); *Specht v. Netscape Commc'ns Corp.*, 306 F.3d 17, 28–32 (2d Cir. 2002).

not impair a website operator's ability to regulate access through its private terms of service.²⁸ One question is whether the seller's interface supplies the AI agent (and, where applicable, the human user) with the same material price, cancellation, renewal, and dispute resolution terms that it presents in a transaction not initiated by an AI agent.

- **Identity verification and aggregate risk.** Verification may account not only for the amount of a single purchase, but also for the frequency, cumulative value, delivery location, reversibility, and other indicators of fraud or error. Repeated low-value transactions may warrant enhanced review because their aggregate exposure can be substantial. Verification can include a known account, a distinct agent credential, or a signed delegation. If a seller disregards an agreed procedure or conspicuous signs of compromise, a buyer may argue that any reliance on the agent's authority was unreasonable; the seller may respond that the circumstances nevertheless reasonably indicated that the transaction was authorized.
- **Transactions that may warrant further inquiry.** Order velocity, novel delivery locations, changed payment details, or repeated failed confirmations may raise questions about whether to pause fulfillment pending further inquiry into the legitimacy of the transactions.
- **Contractual treatment of automated transactions.** Commercial parties may specify approved credentials, transaction limits, confirmation rules, allocation of loss from compromised systems, notice deadlines, and revocation procedures. Parties that expect to transact through agents on a recurring basis may find the trading partner agreements developed for electronic data interchange a useful precedent. Those agreements settled in advance which terms governed automated exchanges, how errors would be corrected, and how losses would be allocated, and the same questions arise here in sharper form.

V. If the AI Agent Spends Money, Recovery Turns on the Contract First

A company seeking return of money that an AI agent spent on its behalf will generally need to consider first the underlying contract. New York and California courts have generally declined to permit quasi-contract recovery where an enforceable contract governs the same subject matter.²⁹ If the contract is void, voidable, or not attributable to the company, restitution may be available to put both parties in the positions they occupied before the transaction occurred.

New York and California permit recovery of certain mistaken payments, subject to defenses. A recipient that materially changed its position in good-faith reliance may resist full repayment, and the discharge-for-value rule may protect a creditor that received payment of a valid debt then due without notice of the mistake. The Second Circuit's *Citibank* wire-transfer litigation shows how those requirements can matter.

In that famous dispute, which Wall Street followed intently, Citibank intended to send Revlon's lenders a \$7.8 million interest payment but mistakenly transferred nearly \$900 million in principal, almost three years before the debt matured. The Second Circuit vacated the judgment and remanded, holding that the discharge-for-value rule did not shield the defendant loan managers

²⁸ Cf. *Perplexity*, slip op. at 21 n.5 (noting that the decision does not impair the website operator's ability to regulate access through its private terms of service).

²⁹ *Clark-Fitzpatrick, Inc. v. Long Island R.R. Co.*, 70 N.Y.2d 382, 388 (1987); *Durell v. Sharp Healthcare*, 183 Cal. App. 4th 1350, 1370 (2010).

because the early payment and other objective warning signs (in the appellate court's view) put them on inquiry notice of the mistake and because the debt was not yet payable.³⁰ California likewise requires a sufficiently definite preexisting obligation before a recipient may invoke the defense.³¹

The payment rail can change the analysis. Consumer credit card and electronic fund transfer statutes apply rules different from the U.C.C. Article 4A framework for commercial wires, under which an unauthorized payment order may sometimes bind a customer if the bank accepted it in good faith under a commercially reasonable security procedure.³²

VI. Preparing for an Unsettled Legal Landscape

There is not yet a single rule for transactions initiated by generative AI agents. Existing doctrine can address many disputes, but outcomes are likely to turn on the system's design, the parties' communications, the reasonableness of the counterparty's reliance, and the preserved record. Those factors may matter whether the business is using an AI agent to make a purchase or accepting an AI transaction initiated by the buyer.

Meanwhile, card networks, payment providers, and protocol developers are continuing to revise the private rules that govern AI-agent transactions. Businesses that use or accept AI-agent transactions may therefore need to account for changes in the card-network rules, payment-provider agreements, and technical protocols that apply to them. Materials prepared before the 2025–2026 developments described above may no longer reflect current requirements.

If you have any questions about the issues addressed in this memorandum, or if you would like a copy of any of the materials mentioned in it, please do not hesitate to reach out to:



John B. Quinn

Chairman and Founding Partner

Los Angeles

johnquinn@quinnemanuel.com, Tel: 213-443-3200



Marc-Andre Cyr

Associate

New York

marcandrecyr@quinnemanuel.com, Tel: 212-849-7099

To view more memoranda, please visit www.quinnemanuel.com/the-firm/publications/

³⁰ *Citibank, N.A. v. Brigade Cap. Mgmt., LP* 49 F.4th 42, 49–50, 60–78 (2d Cir. 2022). Quinn Emanuel represented the defendant loan managers in this matter.

³¹ *First Sav. & Loan Ass'n v. Bank of Am.*, 4 Cal. App. 3d 393, 400–01 (1970); *Bank of Am. v. Sanati*, 11 Cal. App. 4th 1079, 1084–89 (1992).

³² N.Y. U.C.C. Law §§ 4-A-202, 4-A-204 (McKinney 2026); Cal. Com. Code §§ 11202, 11204 (West 2026); 15 U.S.C. §§ 1643(a), 1693g (2018); 12 C.F.R. § 1005.2(m) (2026).

To update information or unsubscribe, please email updates@quinnemanuel.com