Proskauer

Contract Law in the Age of Agentic AI: Who's Really Clicking "Accept"?

New Media and Technology Law Blog on **April 9, 2025**

In May 2024, we released Part I of this series, in which we discussed agentic AI as an emerging technology enabling a new generation of AI-based hardware devices and software tools that can take actions on behalf of users. It turned out we were early – very early – to the discussion, with several months elapsing before agentic AI became as widely known and discussed as it is today. In this Part II, we return to the topic to explore legal issues concerning user liability for agentic AI-assisted transactions and open questions about existing legal frameworks' applicability to the new generation of AI-assisted transactions.

Background: Snapshot of the Current State of "Agents"[1]

"Intelligent" electronic assistants are not new—the original generation, such as Amazon's Alexa, have been offering narrow capabilities for specific tasks for more than a decade. However, as OpenAI's CEO Sam Altman commented in May 2024, an advanced AI assistant or "super-competent colleague" could be the killer app of the future. Later, Altman noted during a Reddit AMA session: "We will have better and better models. But I think the thing that will feel like the next giant breakthrough will be agents." A McKinsey report on AI agents echoes this sentiment: "The technology is moving from thought to action." Agentic AI represents not only a technological evolution, but also a potential means to further spread (and monetize) AI technology beyond its current uses by consumers and businesses. Major AI developers and others have already embraced this shift, announcing initiatives in the agentic AI space. For example:

- **Anthropic** announced an updated frontier AI model in public beta capable of interacting with and using computers like human users;
- **Google** unveiled Gemini 2.0, its new AI model for the agentic era, alongside Project Mariner, a prototype leveraging Gemini 2.0 to perform tasks via an experimental Chrome browser extension (while keeping a "human in the loop");
- **OpenAI** launched a "research preview" of Operator, an AI tool that can interface with computers on users' behalf, and launched beta feature "Tasks" in ChatGPT to

facilitate ongoing or future task management beyond merely responding to real time prompts;

- **LexisNexis** announced the availability of "Protégé," a personalized AI assistant with agentic AI capabilities;
- **Perplexity** recently rolled out "Shop Like a Pro," an AI-powered shopping recommendation and buying feature that allows Perplexity Pro users to research products and, for those merchants whose sites are integrated with the tool, purchase items directly on Perplexity; and
- **Amazon** announced Alexa+, a new generation of Alexa that has agentic capabilities, including enabling Alexa to navigate the internet and execute tasks, as well as Amazon Nova Act, an AI model designed to perform actions within a web browser.

Beyond these examples, other startups and established tech companies are also developing AI "agents" in this country and overseas (including the invite-only release of Manus AI by Butterfly Effect, an AI developer in China). As a recent Microsoft piece speculates, the generative AI future may involve a "new ecosystem or marketplace of agents," akin to the current smartphone app ecosystem. Although early agentic AI device releases have received mixed reviews and seem to still have much unrealized potential, they demonstrate the capability of such devices to execute multistep actions in response to natural language instructions.

Like prior technological revolutions—personal computers in the 1980s, e-commerce in the 1990s and smartphones in the 2000s—the emergence of agentic AI technology challenges existing legal frameworks. Let's take a look at some of those issues – starting with basic questions about contract law.

Note: This discussion addresses general legal issues with respect to hypothetical agentic AI devices or software tools/apps that have significant autonomy. The examples provided are illustrative and do not reflect any specific AI tool's capabilities.

Automated Transactions and Electronic Agents

Electronic Signatures Statutory Law Overview

A foundational legal question is whether transactions initiated and executed by an AI tool on behalf of a user are enforceable. Despite the newness of agentic AI, the legal underpinnings of electronic transactions are well-established. The Uniform Electronic Transactions Act ("UETA"), which has been adopted by every state and the District of Columbia (except New York, as noted below), the federal E-SIGN Act, and the Uniform Commercial Code ("UCC"), serve as the legal framework for the use of electronic signatures and records, ensuring their validity and enforceability in interstate commerce. The fundamental provisions of UETA are Sections 7(a)-(b), which provide: "(a) A record or signature may not be denied legal effect or enforceability solely because it is in electronic form; (b) A contract may not be denied legal effect or enforceability solely because an electronic record was used in its formation."

UETA is technology-neutral and "applies only to transactions between parties each of which has agreed to conduct transactions by electronic means" (allowing the parties to choose the technology they desire). In the typical e-commerce transaction, a human user selects products or services for purchase and proceeds to checkout, which culminates in the user clicking "I Agree" or "Purchase." This click—while not a "signature" in the traditional sense of the word—may be effective as an electronic signature, affirming the user's agreement to the transaction and to any accompanying terms, assuming the requisite contractual principles of notice and assent have been met.

At the federal level, the E-SIGN Act (15 U.S.C. §§ 7001-7031) ("E-SIGN") establishes the same basic tenets regarding electronic signatures in interstate commerce and contains a reverse preemption provision, generally allowing states that have passed UETA to have UETA take precedence over E-SIGN. If a state does not adopt UETA but enacts another law regarding electronic signatures, its alternative law will preempt E-SIGN only if the alternative law specifies procedures or requirements consistent with E-SIGN, among other things.

However, while UETA has been adopted by 49 states and the District of Columbia, it has not been enacted in New York. Instead, New York has its own electronic signature law, the Electronic Signature Records Act ("ESRA") (N.Y. State Tech. Law § 301 et seq.). ESRA generally provides that "An electronic record shall have the same force and effect as those records not produced by electronic means." According to New York's Office of Information Technology Services, which oversees ESRA, "the definition of 'electronic signature' in ESRA § 302(3) conforms to the definition found in the E-SIGN Act." Thus, as one New York state appellate court stated, "E-SIGN's requirement that an electronically memorialized and subscribed contract be given the same legal effect as a contract memorialized and subscribed on paper…is part of New York law, whether or not the transaction at issue is a matter 'in or affecting interstate or foreign commerce.'"[2]

Given US states' wide adoption of UETA model statute, with minor variations, this post will principally rely on its provisions in analyzing certain contractual questions with respect to AI agents, particularly given that E-SIGN and UETA work toward similar aims in establishing the legal validity of electronic signatures and records and because E-SIGN expressly permits states to supersede the federal act by enacting UETA. As for New York's ESRA, courts have already noted that the New York legislature incorporated the substantive terms of E-SIGN into New York law, thus suggesting that ESRA is generally harmonious with the other laws' purpose to ensure that electronic signatures and records have the same force and effect as traditional signatures.

Electronic "Agents" under the Law

Beyond affirming the enforceability of electronic signatures and transactions where the parties have agreed to transact with one another electronically, Section 2(2) of UETA also contemplates "automated transactions," defined as those "conducted or performed, in whole or in part, by electronic means or electronic records, in which the acts or records of one or both parties are not reviewed by an individual." Central to such a transaction is an "electronic agent," which Section 2(6) of UETA defines as "a computer program or an electronic or other automated means used independently to initiate an action or respond to electronic records or performances in whole or in part, without review or action by an individual." Under UETA, in an automated transaction, a contract may be formed by the interaction of "electronic agents" of the parties or by an "electronic agent" and an individual. E-SIGN similarly contemplates "electronic agents," and states: "A contract or other record relating to a transaction in or affecting interstate or foreign commerce may not be denied legal effect, validity, or enforceability solely because its formation, creation, or delivery involved the action of one or more electronic agents so long as the action of any such electronic agent is legally attributable to the person to be bound."[3] Under both of these definitions, agentic AI tools—which are increasingly able to initiate actions and respond to records and performances on behalf of users—arguably qualify as "electronic agents" and thus can form enforceable contracts under existing law.[4]

AI Tools and E-Commerce Transactions

Given this existing body of statutory law enabling electronic signatures, from a practical perspective this may be the end of the analysis for most e-commerce transactions. If I tell an AI tool to buy me a certain product and it does so, then the product's vendor, the tool's provider and I might assume—with the support of UETA, E-SIGN, the UCC, and New York's ESRA—that the vendor and I (via the tool) have formed a binding agreement for the sale and purchase of the good, and that will be the end of it unless a dispute arises about the good or the payment (e.g., the product is damaged or defective, or my credit card is declined), in which case the AI tool isn't really relevant.

But what if the transaction does not go as planned for reasons related to the AI tool? Consider the following scenarios:

- **Misunderstood Prompts**: The tool misinterprets a prompt that would be clear to a human but is confusing to its model (e.g., the user's prompt states, "Buy two boxes of 101 Dalmatians Premium dog food," and the AI tool orders 101 two-packs

of dog food marketed for Dalmatians).

- **AI Hallucinations**: The user asks for something the tool cannot provide or does not understand, triggering a hallucination in the model with unintended consequences (e.g., the user asks the model to buy stock in a company that is not public, so the model hallucinates a ticker symbol and buys stock in whatever real company that symbol corresponds to).
- **Violation of Limits**: The tool exceeds a pre-determined budget or financial parameter set by the user (e.g., the user's prompt states, "Buy a pair of running shoes under $100" and the AI tool purchases shoes from the UK for £250, exceeding the user's limit).
- **Misinterpretation of User Preference**: The tool misinterprets a prompt due to lack of context or misunderstanding of user preferences (e.g., the user's prompt states, "Book a hotel room in New York City for my conference," intending to stay near the event location in lower Manhattan, and the AI tool books a room in Queens because it prioritizes price over proximity without clarifying the user's preference).

Disputes like these begin with a conflict between the user and a vendor—the AI tool may have been effective to create a contract between the user and the vendor, and the user may then have legal responsibility for that contract. But the user may then seek indemnity or similar rights against the developer of the AI tool.

Of course, most developers will try to avoid these situations by requiring user approvals before purchases are finalized (i.e., "human in the loop"). But as desire for efficiency and speed increases (and AI tools become more autonomous and familiar with their users), these inbuilt protections could start to wither away, and users that grow accustomed to their tool might find themselves approving transactions without vetting them carefully. This could lead to scenarios like the above, where the user might seek to void a transaction or, if that fails, even try to avoid liability for it by seeking to shift his or her responsibility to the AI tool's developer.[5] Could this ever work? Who is responsible for unintended liabilities related to transactions completed by an agentic AI tool?

Sources of Law Governing AI Transactions

AI Developer Terms of Service

As stated in UETA's Prefatory Note, the purpose of UETA is "to remove barriers to electronic commerce by validating and effectuating electronic records and signatures." Yet, the Note cautions, "It is NOT a general contracting statute – the substantive rules of contracts remain unaffected by UETA." E-SIGN contains a similar disclaimer in the statute, limiting its reach to statutes that require contracts or other records be written, signed, or in non-electronic form (15 U.S.C. §7001(b)(2)). In short, UETA, E-SIGN, and the similar UCC provisions do not provide contract law rules on how to form an agreement or the enforceability of the terms of any agreement that has been formed.

Thus, in the event of a dispute, terms of service governing agentic AI tools will likely be the primary source to which courts will look to assess how liability might be allocated. As we noted in Part I of this post, early-generation agentic AI hardware devices generally include terms that not only disclaim responsibility for the actions of their products or the accuracy of their outputs, but also seek indemnification against claims arising from their use. Thus, absent any express customer-favorable indemnities, warranties or other contractual provisions, users might generally bear the legal risk, barring specific legal doctrines or consumer protection laws prohibiting disclaimers or restrictions of certain claims.[6]

But what if the terms of service are nonexistent, don't cover the scenario, or—more likely—are unenforceable? Unenforceable terms for online products and services are not uncommon, for reasons ranging from "browsewrap" being too hidden, to specific provisions being unconscionable. What legal doctrines would control during such a scenario?

The Backstop: User Liability under UETA and E-SIGN

Where would the parties stand without the developer's terms? E-SIGN allows for the effectiveness of actions by "electronic agents" "so long as the action of any such electronic agent is legally attributable to the person to be bound." This provision seems to bring the issue back to the terms of service governing a transaction or general principles of contract law. But again, what if the terms of service are nonexistent or don't cover a particular scenario, such as those listed above. As it did with the threshold question of whether AI tools could form contracts in the first place, UETA appears to offer a position here that could be an attractive starting place for a court. Moreover, in the absence of express language under New York's ESRA, a New York court might apply E-SIGN (which contains an "electronic agent" provision) or else find insight as well by looking at UETA and its commentary and body of precedent if the court isn't able to find on-point binding authority, which wouldn't be a surprise, considering that we are talking about technology-driven scenarios that haven't been possible until very recently.

UETA generally attributes responsibility to *users* of "electronic agents", with the prefatory note explicitly stating that the actions of electronic agents "programmed and used by people will bind the user of the machine." Section 14 of UETA (titled "Automated Transaction") reinforces this principle, noting that a contract can be formed through the interaction of "electronic agents" "even if no individual was aware of or reviewed the electronic agents' actions or the resulting terms and agreements." Accordingly, when automated tools such as agentic AI systems facilitate transactions between parties who knowingly consent to conduct business electronically, UETA seems to suggest that responsibility defaults to the users—the persons who most immediately directed or initiated their AI tool's actions. This reasoning treats the AI as a user's *tool*, consistent with the other UETA Comments (e.g., "contracts can be formed by machines functioning as electronic agents for parties to a transaction").

However, different facts or technologies could lead to alternative interpretations, and ambiguities remain. For example, Comment 1 to UETA Section 14 asserts that the lack of human intent at the time of contract formation does not negate enforceability in contracts "formed by machines functioning as electronic agents for parties to a transaction" and that "when machines are involved, the requisite intention flows from the *programming* and *use* of the machine" (emphasis added).

This explanatory text has a couple of issues. First, it is unclear about what constitutes "programming" and seems to presume that the human intention at the programming step (whatever that may be) is more-or-less the same as the human intention at the use step[7], but this may not always be the case with AI tools. For example, it is conceivable that an AI tool could be programmed by its developer to put the developer's interests above the users', for example by making purchases from a particular preferred e-commerce partner even if that vendor's offerings are not the best value for the end user. This concept may not be so far-fetched, as existing GenAI developers have entered into content licensing deals with online publishers to obtain the right for their chatbots to generate outputs or feature licensed content, with links to such sources. Of course, there is a difference between a chatbot offering links to relevant licensed news sources that are accurate (but not displaying appropriate content from other publishers) versus an agentic chatbot entering into unintended transactions or spending the user's funds in unwanted ways. This discrepancy in intention alignment might not be enough to allow the user to shift liability for a transaction from a user to a programmer, but it is not hard to see how larger misalignments might lead to thornier questions, particularly in the event of litigation when a court might scrutinize the enforceability of an AI vendor's terms (under the unconscionability doctrine, for example).

Second, UETA does not contemplate the possibility that the AI tool might have enough autonomy and capability that some of its actions might be properly characterized as the result of its *own* intent. Looking at UETA's definition of "electronic agent," the commentary notes that "As a general rule, the employer of a tool is responsible for the results obtained by the use of that tool since the tool has no independent volition of its own." But as we know, technology has advanced in the last few decades and depending on the tool, an autonomous AI tool might one day have much independent volition (and further UETA commentary admits the possibility of a future with more autonomous electronic agents). Indeed, modern AI researchers have been contemplating this possibility even before rapid technological progress began with ChatGPT.

Still, Section 10 of UETA may be relevant to some of the scenarios from our bulleted selection of AI tool mishaps listed above, including misunderstood prompts or AI hallucinations. UETA Section 10 (titled “Effect of Change or Error”) outlines the possible actions a party may take when discovering human or machine errors or when “a change or error in an electronic record occurs in a transmission between parties to a transaction.” The remedies outlined in UETA depend on the circumstances of the transaction and whether the parties have agreed to certain security procedures to catch errors (e.g., a “human in the loop” confirming an AI-completed transaction) or whether the transaction involves an individual and a machine.[8] In this way, the guardrails integrated into a particular AI tool or by the parties themselves play a role in the liability calculus. The section concludes by stating that if none of UETA’s error provisions apply, then applicable law governs, which might include the terms of the parties’ contract and the law of mistake, unconscionability and good faith and fair dealing.

* * *

Thus, along an uncertain path we circle back to where we started: the terms of the transaction and general contract law principles and protections. However, not all roads lead to contract law. In our next installment in this series, we will explore the next logical source of potential guidance on AI tool liability questions: agency law. Decades of established law may now be challenged by a new sort of “agent” in the form of agentic AI…and a new AI-related lawsuit foreshadows the issues to come.

---

[1] In keeping with common practice in the artificial intelligence industry, this article refers to AI tools that are capable of taking actions on behalf of users as “agents” (in contrast to more traditional AI tools that can produce content but not take actions). However, note that the use of this term is not intended to imply that these tools are “agents” under agency law.

[2] In addition, the UCC has provisions consistent with UETA and E-SIGN providing for the use of electronic records and electronic signatures for transactions subject to the UCC. The UCC does not require the agreement of the parties to use electronic records and electronic signatures, as UETA and E-SIGN do.

[3] Under E-SIGN, “electronic agent” means “a computer program or an electronic or other automated means used independently to initiate an action or respond to electronic records or performances in whole or in part without review or action by an individual at the time of the action or response.”

[4] It should be noted that New York’s ESRA does not expressly provide for the use of “electronic agents,” yet does not prohibit them either. Reading through ESRA and the ESRA regulation, the spirit of the law could be construed as forward-looking and seems to suggest that it supports the use of automated systems and electronic means to create legally binding agreements between willing parties. Looking to New York precedent, one could also argue that E-SIGN, which contains provisions about the use of “electronic agents”, might also be applicable in certain circumstances to fill the “electronic agent” gap in ESRA. For example, the ESRA regulations (9 CRR-NY § 540.1) state: “New technologies are frequently being introduced. The intent of this Part is to be flexible enough to embrace future technologies that comply with ESRA and all other applicable statutes and regulations.” On the other side, one could argue that certain issues surrounding “electronic agents” are perhaps more unsettled in New York. Still, New York courts have found ESRA consistent with E-SIGN.

[5] Since AI tools are not legal persons, they could not be liable themselves (unlike, for example, a rogue human agent could be in some situations). We will explore agency law questions in Part III.

[6] Once agentic AI technology matures, it is possible that certain user-friendly contractual standards might emerge as market participants compete in the space. For example, as we wrote about in a prior post, in 2023 major GenAI providers rolled out indemnifications to protect their users from third-party claims of intellectual property infringement arising from GenAI outputs, subject to certain carve-outs.

[7] The electronic “agents” in place at the time of UETA’s passage might have included basic e-commerce tools or EDI (Electronic Data Interchange), which is used by businesses to exchange standardized documents, such as purchase orders, electronically between trading partners, replacing traditional methods like paper, fax, mail or telephone. Electronic tools are generally designed to explicitly perform according to the user’s intentions (e.g., clicking on an icon will add this item to a website shopping cart or send this invoice to the customer) and UETA, Section 10, contains provisions governing when an inadvertent or electronic error occurs (as opposed to an abrogation of the user’s wishes).

[8] For example, UETA Section 10 states that if a change or error occurs in an electronic record during transmission between parties to a transaction, the party who followed an agreed-upon security procedure to detect such changes can avoid the effect of the error, if the other party who didn’t follow the procedure would have detected the change had they complied with the security measure; this essentially places responsibility on the party who failed to use the agreed-upon security protocol to verify the electronic record’s integrity.

Comments to UETA Section 10 further explain the context of this section: “The section covers both changes and errors. For example, if Buyer sends a message to Seller ordering 100 widgets, but Buyer’s information processing system changes the order to 1000 widgets, a “change” has occurred between what Buyer transmitted and what Seller received. If on the other hand, Buyer typed in 1000 intending to order only 100, but sent the message before noting the mistake, an error would have occurred which would also be covered by this section.” In the situation where a human makes a mistake when dealing with an electronic agent, the commentary explains that “when an individual makes an error while dealing with the electronic agent of the other party, it may not be possible to correct the error before the other party has shipped or taken other action in reliance on the erroneous record.”

View original.

Related Professionals

- **Peter J. Cramer**
  Associate

- **Caroline E. Rimmer**
  Associate