The Liability Counter-Pressure Reaches the Banks: AI Agent Risk Is Being Assigned to Consumers

Agent providers disclaim their agents' mistakes. Three major banks restrict access, two have written terms that assign the loss to the customer, and one of those takes effect November 24. The OCC issued a nonbinding examiner resource on November 4, and the federal law that limits what such terms can settle has not yet been tested on this question.

By Basil C. Puglisi, MPA
November 5, 2026

The AI Risk Economy argued that economic pressure favors faster deployment because speed, savings, and profit arrive now while the cost of failure arrives later. Liability and insurance push back, the paper holds, by attaching that cost to the business and setting the terms on which it can be transferred. Events in autumn 2026 test that argument in consumer banking. Two personal AI agents are on the consumer market with terms that assign the consequences of their actions to the user. Four major banks then restricted those agents or wrote contracts that assign the loss to the customer, and one of those contracts takes effect on November 24. The Office of the Comptroller of the Currency published a nonbinding examiner resource on November 4. Beneath all of it sits a federal statute that limits what any of these contracts can settle, and the question these agents raise under it has not yet been put to a court or an agency.

The piece that follows reads the agent providers' terms, the bank contracts, the one proposal that runs the other way, the examiner resource, and the law. The pattern it finds is the one the paper described, with a new address. The first wave of risk assignment reaches the party nearest the agent with the weakest access to the records that would decide a dispute, which is the customer. Liability continues the counter-pressure from above, through a statute, a burden of proof, an examiner's questions, and a competitor's proposal. The records that exist when the first transfer is contested will decide where the loss finally stays.

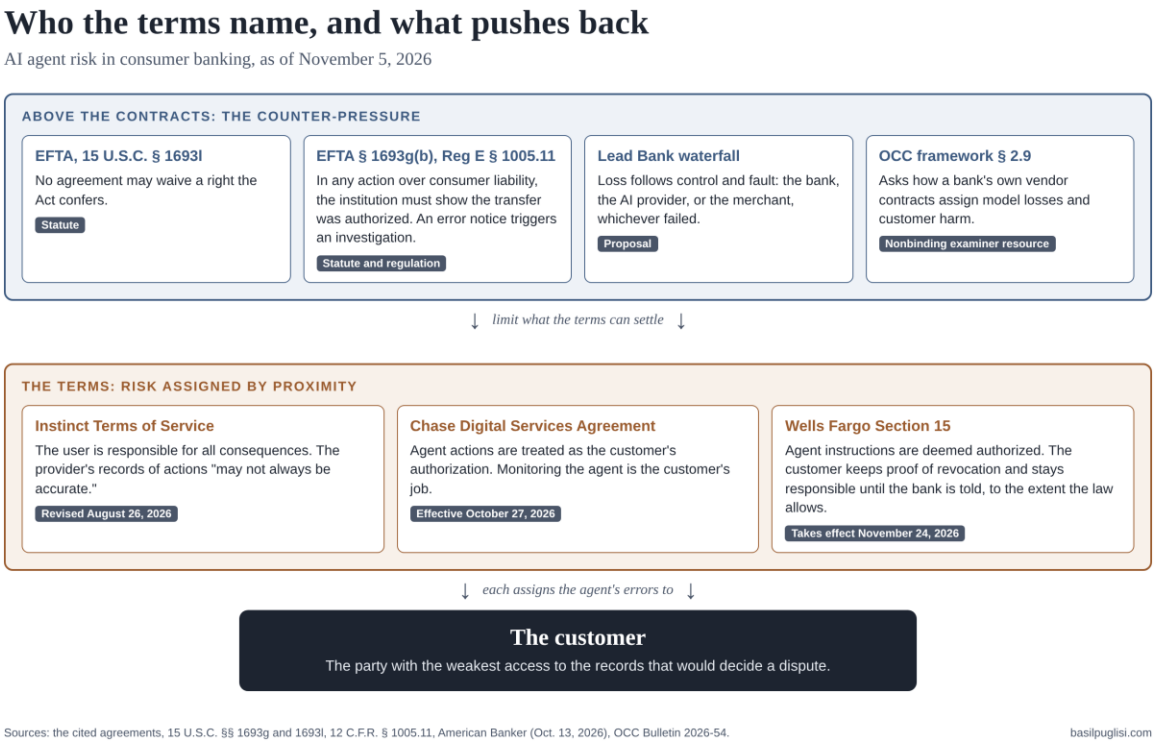


Figure 1. Who the terms name, and what pushes back. An agent provider's terms and two bank agreements assign the agent's errors to the customer. A statute, a regulation, a proposal, and a nonbinding examiner resource sit above them.

The Agent Rush, in Its Own Terms

The mechanism by which corporate incentives systematically prioritize capability advancement over safety validation, which the author calls the Economic Override Pattern, was described in Governing AI in 2025. The AI Risk Economy carried it into the insurance market in 2026. The paper treats it as a working concept, supported by observable evidence and not yet independently validated as formal theory. Its consumer form was on the market by the autumn of 2026, when Meta's Muse and Instinct offered personal agents that can log in, shop, and pay. A client alert from Hunton Andrews Kurth names both products in its account of what the banks did next. It also records Jamie Dimon's statement that JPMorgan Chase will not let outside agents act inside customer accounts, citing the risk of rogue agents.

The provider agreements say who carries the risk. Instinct's Terms of Service are issued by Spear Street Technology, Inc., doing business as Instinct, and were last revised

August 26, 2026. Section 3 states: "You are responsible for your Inputs, Actions, and use of any Output, including all consequences." The user authorizes the service "to take Actions on your behalf as the Services may deem are responsive to your Input." The company is "not responsible for any unintended Actions or any loss, damage, charge, liability." It makes "no representation or warranty that such safeguards will prevent unintended or erroneous Actions." The same section tells the user not to take Actions or rely on Output "where doing so without meaningful human review" would be unreasonable, which places the review duty on the user as well. One sentence matters most for what follows: "Records of Actions available through the Services may not always be accurate." Section 10 closes with binding arbitration and a class action waiver, which a user may opt out of by written notice within 30 days of registering.

For Muse, the record is thinner. Moneywise reported on September 29 that a Meta help page for the agent tells users, "You're responsible for all transactions your Muse makes on your behalf." The same page, by that report, advises them to "keep an eye out for email confirmations, receipts and statements." The Meta page could not be read for this piece, so the sentence stands as reported and not verified. If it reads as reported, the second provider has also made liability for the agent's transactions the consumer's problem, as far as its terms can carry it. A consumer who enters a bank through either product already carries the agent's errors before the bank says a word.

Four Banks, Two Instruments

JPMorgan Chase, Charles Schwab, Citigroup, and Wells Fargo each moved against third-party agents that log in with the customer's own credentials. The Hunton alert tabled the moves on November 3 and dates the three announced actions to October 13 through 20; Wells Fargo's amended agreement surfaced in reporting on November 2 and takes effect November 24. The alert's table shows three institutions blocking and Wells Fargo rewriting its terms, while its prose says four chose a hard block. The table matches the underlying documents, and this piece follows it. Each institution still permits customer-authorized, read-only data sharing through the connections it approves, so none of the four banned AI outright. What they banned, or priced, was an outside agent holding the keys. International Business Times reported on October 23

that Bank of America and Morgan Stanley also restricted agents, and, citing Reuters, that a source said Morgan Stanley had not seen rogue agents accessing its systems.

Chase paired a technical restriction with a contractual allocation. On the bank's third-quarter earnings call on October 13, Chief Financial Officer Jeremy Barnum told analysts that credential sharing is not secure and that the bank blocks it where it can. A customer who hands credentials to an agent that then does something unwanted will find, in his words, that "that's going to be on them" (Volkova, 2026a). The bank's public website Terms of Use, which govern interaction with Chase's systems and are separate from the account agreement, carry a section headed AI Agent or Agentic AI Notices. Any agent that reaches those systems must identify itself as an AI agent at all times and label all agentic traffic. It must also answer truthfully when asked whether traffic comes from a human or a machine. The person who deploys the agent is responsible for actions or errors that result from the agent's incorrect interpretation of instructions, from unintended actions, or from the agent exceeding the authority it was given.

The account contract is the Digital Services Agreement, last updated October 2026, retrieved and read in full on November 5. Effective October 27, 2026, third parties "including External AI, automated tools, or data aggregators" are prohibited from using a customer's login credentials. A customer who permits one anyway is deemed to agree that the bank "will rely on, and treat as evidence of your authorization, any actions taken or instructions given by the third party as if you personally took or gave those actions or instructions." The customer is "responsible for any resulting errors, including where the third party performs unintended actions, makes transactions you did not intend, or exceeds the authority you gave it." It is "solely your responsibility to monitor the third party's activities," and the bank has "no obligation to automatically or subsequently review or screen the third party's activities." The customer assumes "all risks and losses associated with such access or use," to the fullest extent the law permits.

Two further provisions in the same agreement matter here. A separate change, effective January 1, 2027, lets the bank require its mobile app and a trusted device for identity verification, "unless an accommodation is required by law." And the agreement defines a Third-party AI Agent as software "that takes autonomous, semi-autonomous, or non-

autonomous action on behalf of, or at the instruction of, any person (human or entity)." The same three words appear in Wells Fargo's definition.

Schwab's chief executive told Bloomberg on October 16 that the firm blocks unauthorized agents that try to move client money. He put the count at 10 to 20 times a week, as Hunton relays it, a count of blocked attempts rather than of losses. Citigroup said the same day that it does not permit third-party agents to log in or transact inside its digital platforms. Both statements reach this piece through the alert, which links to the original reporting.

Wells Fargo published a different kind of response. Its amended Online Access Agreement, in a version effective November 24, 2026, adds a Section 15, Third-Party Artificial Intelligence Agents or Agentic AI. The section declares no prohibition, though it reserves the bank's authority to restrict any agent, and it allocates the risk of the agents it does not prohibit.

Section 15, Clause by Clause

The upcoming agreement, effective November 24, 2026, is public, and the clauses below are read from its full text as retrieved on November 5.

Section 15(a) defines an AI Agent as any software, application, program, or other technology that takes autonomous, semi-autonomous, or non-autonomous action on behalf of any person or entity. The definition reaches a shopping assistant, a budgeting app that logs in with a password, and a scheduled script that acts on the account alike.

Section 15(b) states that the bank does not provide, control, or endorse any agent and is not responsible for agents or their activities. It does not monitor, review, audit, or certify an agent's security or data handling. Section 15(c) has the customer consent to sharing financial, account, and personal information with the agent, and acknowledge that, in the agreement's words, what the agent does with it sits outside the bank's control.

Section 15(d) carries the allocation, and it opens with the qualifier that governs everything after it: unless otherwise provided by law or regulation. A customer who gives an agent access with the customer's login credentials, or otherwise authorizes an

agent to act, agrees to the following. Any instruction, action, or transaction the agent initiates is deemed initiated and authorized by the customer, without further verification. The bank has no obligation to investigate the validity or scope of the agent's authority, and may rely on the agent's instructions as valid, binding, and as evidence of the customer's authorization. The agreement deems. Whether the law agrees is the question the statute answers below.

Section 15(e) states, in capitals, that use of any agent is at the customer's own risk. The customer is solely responsible for monitoring the agent and keeping it within the scope of authority granted. Unless law provides otherwise, the customer also answers for any actions, errors, or results from the agent's interpretation of instructions or from unintended actions. The same holds for incomplete, inaccurate, or manipulated inputs, a phrase that may include an adversarial prompt injection, and for the agent exceeding its authority. The customer acknowledges that the bank does not monitor, screen, or review agent activity, including any transactions an agent initiates. That sentence concerns the agent's activity; what the bank records on its own channels is a separate matter, taken up below.

Section 15(f) is the clause the trade press noticed. The customer is responsible for properly and timely revoking an agent's authorization, must retain evidence of the revocation, and must provide that evidence to the bank on request. Then comes the sentence that sets the clock: "Until you notify us of such revocation pursuant to Section 17(d), you remain responsible for all the AI Agent's activities, to the extent permitted by law or regulation." Section 17(d), Protecting Your Security, requires the customer to notify the bank immediately of suspected fraudulent activity and names two channels for doing so, a telephone number answered 24 hours a day and a postal address. Under the clause, responsibility is asserted to run until the bank receives that notice, subject to law. The event the contract names is notice to the bank. Notice to the agent is not that event, and a record in the customer's hands is not that event either.

Section 15(g) reserves to the bank the right to change its security procedures for agents, "including requiring that you register your AI Agents with us." The bank may also change, limit, or discontinue any API or other access method, and may limit, prohibit,

suspend, or terminate any agent's access. The section reserves the tools of visibility without offering them; registration is a right the bank keeps, not a channel the customer can use today. Section 15(h) bars anyone under 18 from permitting an agent to act.

Additional provisions elsewhere in the same agreement bear on the agent clauses, and they sort into three groups. The first concerns people the customer appoints. Section 2(b) treats an Authorized Representative's transactions, "including the transactions of their AI Agents," as authorized. That status ends only when the customer tells the bank otherwise and the bank has had "reasonable opportunity to act on your notification." Section 2(d) treats such agents' transactions, including ones the customer did not want or intend, as transactions the customer authorized.

The second concerns the consumer protections the agreement incorporates. Section 5 applies only to electronic fund transfers that debit or credit a consumer account and are subject to Regulation E. It defines an unauthorized transfer as one "initiated by a person other than you without actual authority to initiate the transfer and from which you receive no benefit." It adds that a transaction the customer initiates or authorizes, "even if induced to do so through fraud or deceit," is not an unauthorized transfer. That sentence reaches a customer who is deceived into authorizing a payment, not a transfer initiated by someone who obtained the credentials by fraud. And it states that the section "does not apply to wire transfers that debit or credit a consumer's account" or to remittance transfers, which fall under Section 8.

The third concerns the commercial rules and the baseline. Sections 8(b) and 17(c) state that permitting an agent may adversely impact the commercially reasonable security procedure the customer agreed to under Article 4A of the Uniform Commercial Code, and that the customer accepts that risk. Section 17(c) applies only to transactions subject to Article 4A, "including online transactions from business accounts." Section 19 states the baseline that predates agents: "You are responsible for actions taken by anyone using the Service after signing in with your username and password." Section 20(c) extends the customer's indemnity of the bank to the agents of Authorized Representatives and Delegates, for consumer and business accounts alike. Section 21 is

the dispute resolution program, with binding arbitration and a waiver of class action rights, both listed on the agreement's cover beside the agent responsibilities.

One line on that cover page sets the frame for all of it. The summary lists the customer's responsibilities for artificial intelligence agents "if and when we permit the use of such agents." The terms precede the permission. No general permission for credential-based agents is open in this agreement, and the rules for one are already written.

The Bank's Reason, and What the Agreement Says Elsewhere

Asked what counts as evidence of revocation, a Wells Fargo spokesman told American Banker that customers may authorize or revoke an agent outside the bank. The agreement asks for a record, he said, because the bank "may not otherwise have visibility into when that occurred" (Crosman, 2026b). The sentence deserves a precise reading. It concerns one event, a revocation that happens inside the customer's relationship with the agent, where the bank is not a party. It is not a statement that the bank keeps no records. Section 18 of the same agreement describes "analytics and monitoring tools" and "technical log files," including "interaction monitoring and visual replay for fraud." It also states that the bank "may use automated processes to detect any use of the Service that violates the terms of this Agreement." Those records can reconstruct what happened on the bank's own channels. They cannot show what the customer told an agent hosted elsewhere, or when the customer revoked it there. The bank holds records of the session and the transfer. It lacks the record of the one event it has asked the customer to keep, and it has written that the customer carries the loss until the bank itself is told.

The bank's position has a defense, and it deserves its own paragraph. A credential-sharing agent looks to the bank like the customer. The bank cannot tell, from the login alone, whether the instruction came from the person, from software the person directed, or from software someone else hijacked. It carries fraud, identity, anti-money-laundering, privacy, and operational obligations that do not pause while it finds out. A rule against sharing credentials and a request for proof of revocation answer a real

vulnerability, which is Hunton's reading of the moves and part of Borthwick's. The question this piece asks is not whether the bank has a reason. It is where the loss sits while the reason is being acted on, and who holds the record that decides it.

Two practitioners answered in the same piece. Douglas Borthwick, a former Morgan Stanley and Standard Chartered executive, now runs The Insumer Model, a wallet authentication startup. He called it fair for a bank to want a record it cannot see and unrealistic to expect a customer to produce one. Telling a chatbot to stop leaves nothing anyone else can check, he said, and a screenshot shows that a customer asked, not that the agent stopped. Elaine Duffus of Wolters Kluwer, whose firm is building an agent governance layer, asked a harder question. What happens to the customer who keeps every record the bank asks for and still watches the agent go rogue on an ambiguous instruction? How far down the chain, she asked, would anyone go to find fault?

American Banker's own framing bears repeating. In its reporting, Wells Fargo appears to be the first bank to require customers to prove they turned an agent off, or else carry the loss for what it does. The October 2026 Chase agreement read for this piece deems and disclaims but contains no comparable duty, so the framing holds as far as the documents reviewed reach.

Set the two providers beside the two banks and the shape is plain. Instinct assigns the user all consequences of the agent's actions and says its own record of those actions may be wrong. Wells Fargo deems every agent instruction the customer's and asks the customer for the record of the one event the bank cannot see. Chase deems every agent action the customer's and assigns the customer the monitoring the bank will not do. Each instrument was drafted by the party it protects. Each assigns the loss to the party in the middle, who has the weakest access to the records that would decide a dispute. Whether the assignment holds is the statute's question, and it comes below.

The Allocation That Runs the Other Way

The bank clauses are answers to a question Federal Reserve Governor Christopher J. Waller framed at Sibos in Miami on September 29, describing the issue rather than settling it. "Liability in this context comes down to a pretty simple question," he said.

"Who is on the hook if an agent makes the wrong purchase?" Current e-commerce liability frameworks, including network rules and consumer protection standards, "could potentially be adapted for agentic commerce," he added. There may also be "an opportunity to test and refine new approaches" (Waller, 2026).

Lead Bank in Kansas City, a state-chartered direct bank that also runs a banking-as-a-service platform, has proposed a waterfall. As its chief executive, Jackie Reses, described it to American Banker on October 13, the consumer would not be liable where an agent exceeded delegated authority. The bank would be liable where it controlled the delegation of authority and failed to enforce it. The AI provider would be liable where the failure arose from a model defect, insecure tool use, or inadequate safeguards. The merchant would be liable where it knowingly induced an agent to exceed its authority. Reses argued that the technology to assign and enforce that accountability exists today, through behavioral drift monitoring, attestations, know-your-agent refreshes, and transaction credentials that can be issued and revoked (Crosman, 2026a). The proposal, the infrastructure, and the market are three different things. The bank's own engineering post of September 29 describes per-task account numbers, rail-scoped controls, and revocable access as features of its own platform. It closes with a more measured line: "Today's capabilities are the first step." American Express announced an Agent Purchase Protection for registered agents in April 2026 as something that will apply "in the future" (Business Wire, 2026). The rails that would let a bank see an agent's authority and its revocation are being built, and they are not yet the rails most consumers bank on.

Set the two allocations side by side. Wells Fargo and Chase assign agent overreach to the customer and retain contractual discretion over registration, access, and termination, which could support visibility if the banks build on it. Lead Bank proposes to assign it to whoever controlled the delegation or built the model, and to record every credential issued and revoked. One allocation, written into contracts, follows proximity. The other, proposed, follows control. That is this piece's reading of the two, and nothing above them has yet decided which one holds.

What the Regulators Said, and What They Did Not

The federal banking agencies left agents outside their model risk guidance on purpose, and said where the question should go. SR 26-2 is the revised model risk management guidance the Federal Reserve, the FDIC, and the OCC issued on April 17, 2026. Its third footnote states that generative AI and agentic AI models "are novel and rapidly evolving" and "are not within the scope of this guidance." The footnote continues: "Nonetheless, a banking organization's risk management and governance practices should guide the determination of appropriate governance and controls for any tools, processes, or systems not covered in this document." The agencies excluded these systems from this guidance and left their governance to each institution's broader risk management, within the law that already applies. The institutions answered with the contracts above.

On November 4, the OCC filled part of the space, and was careful to say how little. OCC Bulletin 2026-54 released a supervisory framework for bank use of super intelligence, the executive branch's label under Executive Order 14434 for the technologies the statute defines as artificial intelligence. The bulletin states that the framework does not establish supervisory expectations for banks and that banks are not expected to use it to assess their own SI use. Nothing in it sets enforceable standards or prescriptive requirements. It is a resource for examiners, who are told to tailor their work to the materiality of each use case. The bulletin is addressed to the chief executives of national banks, federal savings associations, and federal branches and agencies, the institutions the OCC supervises; Wells Fargo Bank, N.A. and JPMorgan Chase Bank, N.A. are among them, and Lead Bank, a state-chartered institution, sits outside the OCC's supervision. The framework notes that the April guidance applies to non-generative, non-agentic models, so the new document is more applicable to generative and agentic systems. It records that it was prepared with the assistance of SI and reviewed, edited, and approved by a human.

Four parts of the framework bear on the contracts: its risk tiering, its reviewer standard, its treatment of external agents, and its treatment of vendor contracts.

The first is its overarching principle. Risk classification may be based on potential consequences, transfer of agency, and the degree of human oversight, rather than on the technical sophistication of the system. In the illustrative tiers, which assess a bank's own use cases, a fully autonomous, customer-facing payment agent sits at Tier 1, regardless of how well it performs in testing.

The second is Section 2.3, on human-in-the-loop and human-on-the-loop controls over the bank's own systems. The assigned reviewer may need sufficient expertise, information, and authority to evaluate, reject, modify, pause, or halt the system's output or action. Review triggers, evidence, and procedures may be defined "so that approval is more than a nominal click," and high volumes or bulk approval could reduce review quality. The framework's summary table names the risk as a named human reviewer without practical authority, competency, capacity, or a meaningful review standard.

The third is Section 2.10, on operational resilience, which reaches the agents a bank does not control. Customer, counterparty, fintech, or agentic-commerce systems outside the bank's control could present channel risk, and identity, provenance, delegation, authentication, monitoring, velocity limits, dispute handling, suspension, and revocation may be calibrated to the action requested. The examiner resource lists monitoring and revocation of external agents as controls a bank may calibrate on its own side. One bank's upcoming consumer contract says it does not monitor the agent's activity and that revocation is the customer's job to prove. The framework binds no one, so the two documents do not conflict in law. They describe different oversight postures toward overlapping risk. The examiner resource contemplates bank-side controls over the bank's interaction with outside agents; the contract disclaims surveillance of the agent's own activity and assigns that oversight to the customer.

The fourth is Section 2.9, on third-party and vendor risk. It invites examiners to consider how a bank's contracts with its own vendors "assign responsibility for model related errors or losses," including "customer harm for which the bank may seek clarity, indemnification, caps or loss responsibility." The section speaks to the bank's suppliers, not to an agent a customer brought through the front door, and it expects nothing of anyone. It still supplies the contrast this piece needs. In the one relationship the

examiner resource addresses, the bank is invited to consider recovering customer harm from the party that built the model. In the relationship the consumer contracts address, the bank assigns that harm to the customer who used one. Section 2.14 adds that documentation should let people other than the developers "understand, evaluate, and audit" a bank's system, with records kept for "legal, supervisory, audit, customer-remediation" purposes.

On the liability question itself, the agencies have said nothing binding, and the Comptroller did not resolve it when asked. On Fox Business on October 22, Comptroller Jonathan Gould was asked twice whether customers could get their money back if an unauthorized agent moved it. He "did not specify how such matters would be resolved" and restated his confidence that banks can block the activity (Campbell, 2026). Hunton reports, without naming them, that regulators privately welcomed the restrictions and that bank executives raised agent-initiated transactions with the Federal Reserve. Fed staff, by the same account, have indicated that clarification on authorization and liability for autonomous transactions may follow. Julie Hill, dean of the University of Wyoming College of Law, told American Banker that existing rules generally favor consumers, who are typically not liable for unauthorized transactions. Agents complicate the framework, she said, because the bank can answer a dispute by saying the customer authorized the agent (Volkova, 2026b). Justin Herring of Mayer Brown added that from a technical standpoint it can be nearly impossible to tell whether a request came from a bad actor or the real customer. Both expected rules to arrive after practice, as they usually do.

The Law a Contract Cannot Move

The allocation clauses in Sections 15(d), 15(e), and 15(f) each preserve applicable law with an express qualifier, "unless otherwise provided by law or regulation" or "to the extent permitted by law or regulation." The Chase agreement caps its risk clause "to the fullest extent permitted by law." The qualifiers point at the Electronic Fund Transfer Act and its regulation, and the regulation has more to say than either contract lets on.

The scope comes first. Regulation E covers electronic fund transfers from consumer accounts, debit card and ACH payments among them. It does not cover wire transfers

through Fedwire or similar systems "used primarily for transfers between financial institutions or between businesses" (12 C.F.R. § 1005.3(c)(3)). Wells Fargo's Section 5 says the same of consumer wires and remittances. Article 4A of the Uniform Commercial Code, in turn, does not apply to a funds transfer any part of which is governed by the Act (UCC § 4A-108). The boundary runs by transaction type, not by who holds the account. The same agent, in the same account, may move a debit payment under Regulation E and a wire under Article 4A, and the agreement's security-procedure clause and indemnity reach both. A credential-sharing agent, a bank-approved read-only connector, a scoped payment token, a debit or ACH transfer, and a wire each sit under a different authorization rule and leave a different record. The contracts treat them as one "agent."

The definition comes second. The regulation's definition of an unauthorized transfer, at 12 C.F.R. § 1005.2(m), excludes any transfer initiated by a person who was furnished the access device by the consumer. The exclusion holds until the consumer notifies the institution that transfers by that person are no longer authorized. The official interpretation, comment 2(m)-2, answers the case of a human delegate: "If a consumer furnishes an access device and grants authority to make transfers to a person (such as a family member or co-worker) who exceeds the authority given, the consumer is fully liable for the transfers unless the consumer has notified the financial institution that transfers by that person are no longer authorized." Comment 2(m)-3 draws the other line: an unauthorized transfer "includes a transfer initiated by a person who obtained the access device from the consumer through fraud or robbery." A consumer who knowingly furnished credentials to an agent that then overreached sits on the first side of that line. A consumer whose credentials were taken by a fraudulent agent, or by a manipulated one, may sit on the second, and that is the distinction Wells Fargo's fraud-or-deceit sentence does not draw. Section 15(f)'s notice trigger follows the structure of comment 2(m)-2 exactly.

The industry has argued the point before any regulator has. The Consumer Bankers Association, the trade group for retail banks, has twice put the position on record, as advocacy rather than interpretation. Its January 2026 white paper states that sharing credentials or API keys with an agent "may be deemed furnishing an access device." Its

August 2026 response to House Financial Services Committee Democrats states that providing them to the agent "constitutes 'furnishing' the access device." Whether a credential, a token, a delegated login, or an instruction to a provider counts as furnishing a device to a person is the classification the facts of the first dispute will test.

The statute limits what these contracts can take away. Section 1693l of Title 15 states that "no writing or other agreement between a consumer and any other person may contain any provision which constitutes a waiver of any right conferred or cause of action created by this subchapter." Section 1693g(b) states that "in any action which involves a consumer's liability for an unauthorized electronic fund transfer, the burden of proof is upon the financial institution to show that the electronic fund transfer was authorized." If it was not, the institution must show that the conditions of consumer liability were met. Section 1005.11 of the regulation requires the institution to investigate when a consumer gives notice of an error, within set time limits. The CFPB's Electronic Fund Transfers FAQs state that an agreement cannot restrict a consumer's rights beyond what the law allows, and that "any contract or agreement attempting to do so is a violation of EFTA." They also state that an institution may not require a police report or other documentation before it begins an investigation, and that private network rules less protective than Regulation E do not change the institution's obligations. A deemed-authorization clause is the bank's evidence position at the start of a dispute. It is not the end of one. In any action over a consumer's liability for a transfer the consumer says was unauthorized, the burden of proving authorization rests with the institution by statute. No record the customer keeps or fails to keep moves it.

What remains open is narrower than the trade press has made it, and it is still open. Comment 2(m)-2 speaks of a "person," and the regulation defines a person at § 1005.2(j) as "a natural person or an organization." Whether software acting for a consumer is a person who can be furnished an access device is undecided. No controlling decision addressing this fact pattern was identified in the sources reviewed for this piece as of November 5, 2026. Three readings compete, each a hypothesis about how the regulation applies rather than a settled doctrine, and this piece keeps all three. One holds that the organization operating the agent can be the person who received the device, which could support the bank's position where that organization in fact held the

credentials and initiated the transfer. A second holds that an agent acting within the authority the consumer gave it initiates the transfer as the consumer's own instrument. On that view the question of personhood is never reached, and only an agent acting outside that authority raises it. A third holds that the question is open and not bank-favoring by default. The comment was written for family members and co-workers, not for software whose provider disclaims the accuracy of its own action log. The question Hill's example turns on is this one. The facts that decide it are the authority the consumer granted, whether the agent exceeded it, whether the consumer received any benefit, and whether a manipulated input or a hijacked system stands between the two.

Where the First Wave Landed

The AI Risk Economy traced the first market response to AI exposure in May and September 2026. Carriers moved the risk off their books by form before anyone tried to price governance. W.R. Berkley's absolute AI exclusion of June 2024, which the paper reads against directors and officers coverage, and Verisk's optional endorsement for commercial general liability, effective January 1, 2026, are the instruments the paper reads. The paper's counter is the same force turned the other way: pressure that makes bypassing governance cost more than complying with it. Liability attaches the cost of AI failure to the business, and insurance turns that cost into coverage terms at every renewal. The paper also marks the limit of what a contract can do: "Allocation returns to the party that controlled the decision, and insurance then prices what contracts can no longer move."

The Liability Map sorted AI legal exposure into three channels and called the third, contract and insurance, the quietest channel and the one repricing AI risk fastest. The agent providers' terms, the bank contracts, and the carrier forms all belong there, and they are not the same instrument. An exclusion decides what a carrier covers and leaves the insured's liability to a claimant where it was. A term of service asserts who owes, as between two parties, and holds only to the extent a statute allows. What the three share is drafting. Each was written by the party it protects, inside existing law but ahead of any rule written for agents. Each assigns the risk to the party downstream that cannot see the agent and has the hardest time showing what it did.

That is where this episode adds something the paper did not have in September. The paper's record thesis was written for enterprises. A deployer that keeps a structured, portable, auditable record of who decided what, on what evidence, at which point, has something an underwriter can read and a court can review. Section 15 extends the same demand to an individual account holder, with none of the tooling. The provider on the other side of the same transaction states that its own record of the agent's actions may be inaccurate. The customer is asked for evidence that neither counterparty necessarily holds in usable form. The Liability Map's figure caption put the three channels' shared demand in one line: a producible record that a named human governed the AI. The paper also warns that the record cuts both ways, because "the same record that shows reasonable care can show notice, a deviation from stated controls, or prior knowledge." The revocation record Section 15 asks for is also the record that shows the customer authorized the agent in the first place, and an authorization record proves scope and duration as well as consent.

Take a hypothetical in which a customer tells an agent to pay a $30 subscription. The agent reads a manipulated page and sends $3,000 instead. The customer revokes the agent at the provider that evening and calls the bank the next morning. Five records decide what follows. The first is the instruction and its scope, which the customer may or may not have kept. The second is the credential the agent held, which the bank's channels recorded as the customer. The third is the provider's action log, which Instinct's terms say may be inaccurate. The fourth is the bank's record of the call, which Section 17(d) makes the event that ends the customer's responsibility under the contract. The fifth is the bank's transfer log. Under Section 15, the loss sits with the customer from the transfer until the morning call, to the extent the law permits. Under comment 2(m)-2, it sits there if the agent is a person who was furnished the device and acted within the authority given. Under comment 2(m)-3, it may not, if the manipulated page counts as fraud against the customer. Under Lead Bank's waterfall, it sits with the provider whose safeguards failed. The example is invented. The five records are not, and the party with the weakest hold on them is the one every contract in this piece names.

Checkpoint-Based Governance draws the line the OCC's examiners are now told to look for. Someone approving outputs faster than they can read them is in the loop, and none

of the people the page lists is a governor; the failure mode, in its words, is presence without accountability. The framework defines AI Governance as a named human with binding authority to accept, modify, or reject AI output at a defined checkpoint. It distinguishes that from human-in-the-loop, which it calls presence without binding authority. The OCC's reviewer who can evaluate, reject, modify, pause, or halt comes close to that standard, and its warning that bulk approval hollows out review names how the standard slips. The two documents arrived at that line separately; the framework is the author's, and the OCC's resource adopts nothing. Measured against that line, the consumer in Section 15 cannot hold the checkpoint. A consumer who banks through a credential-sharing agent holds none of the instruments Lead Bank describes. There is no scoped credential to issue, no provider log the consumer can read, no token to revoke at the source, and no way to compel the agent to honor a revocation. The consumer has decisional authority and no technical hold on the infrastructure that executes the decision. The agent layer that sorts agents by where the named human stands, which the author calls HAIA Agents, rests on one position about accountability: an AI system does not act on its own. A human programmed it, or a human failed to exercise oversight and accountability over it. The consumer who hands over credentials is the human who stands closest to the agent, and the provider who shipped it is the human who programmed it. Section 15 answers that sentence by proximity. Lead Bank's waterfall and the OCC's Section 2.9 answer it by control. Which answer a court adopts, if a court is asked, and whether the Regulation E question lands the way the bank's drafting assumes, will be decided on facts that do not exist yet.

What This May Mean for Disabled Users Who Count on AI as a Tool

One group of customers meets these clauses differently. The author's companion papers on Artificially Intelligent Disability Assistance hold that general-purpose AI can function as assistive technology when a person with a disability uses it to increase, maintain, or improve a functional capability. The test is the author's, and it is governance-neutral: "A tool with a human checkpoint qualifies. So does a product that validates its own output, and so does an agent that carries out a sequence of steps on the person's behalf" (When

AI Becomes Assistive Technology). A customer who banks through an agent because a screen, a keyboard, or a sequence of steps is the barrier is using the agent as a tool in that sense. For such a customer, where no equally effective accessible alternative exists, a categorical block removes a route of access rather than a convenience, and a deemed-authorization clause prices that route at every error the agent makes.

The companion paper on policy, When AI Policy Becomes an Accessibility Barrier, asks what follows when an institution's facially neutral AI rule restricts a use that functions as a disability tool. It proposes a model clause that routes such use through the institution's accommodation process rather than its prohibition. The agreements read for this piece carry no such route. Wells Fargo's agreement contains no accommodation language at all. Chase's agreement carries "unless an accommodation is required by law" on its device requirement and nowhere on its agent prohibition, so the clause was available to the drafter and appears on one requirement and not the other. Whatever accommodation policies the two banks keep outside these agreements, the agent clauses do not point to them, and that placement is the fact the companion paper's policy-conflict gate would examine. Federal law lists a bank among the service establishments that are public accommodations (42 U.S.C. § 12181(7)(F)) and bars discrimination in the full and equal enjoyment of their services (§ 12182). Whether a duty follows for a particular agent, whether another secure accommodation would serve, and which defenses apply are fact-specific questions for the companion paper's actor-by-actor screen, and this piece states the statutory facts and stops. Among the materials reviewed, Spain offers the explicit statutory contrast. Its Ley 4/2026, in force since October 23, 2026, gives persons with disabilities the right to use the support products they need. It provides that "en ningún caso podrán ser obligadas a prescindir de dichos productos o a desactivarlos," especially in their dealings with public administrations and in the use of goods, products, and services made available to the public (Real Decreto Legislativo 1/2013, art. 7.6, as amended by Ley 4/2026). The provision is untested, and whether an AI login agent is a support product within its meaning, and how a bank's security limits apply, are questions its text does not answer. The practical question is the one the companion paper's screen asks: what secure channel lets this customer complete the same banking task without handing general credentials to an outside agent, and

does the bank offer it? The AI Risk Economy already asks whether AI policies that prohibit assistive use create exposure of their own, and the bank clauses give that question a concrete address.

What Can Be Done With This Now

The Factics discipline pairs each fact with a tactic and a measurable outcome, and three parties can act on this record today. The measures below are proposed, not observed; no dispute data, loss series, or agent population exists yet to fill them.

A consumer with an agent faces two providers whose terms assign the agent's consequences to the user and two banks whose terms deem every agent instruction authorized. The tactic follows the contracts' own mechanics. Share no credentials outside the channels the bank approves, and use read-only connectors where the bank offers them. Where an agent is used at all, keep the instruction and its scope in writing outside the agent, revoke at the provider, and keep the provider's confirmation. Then notify the bank through the channels its agreement names; for Wells Fargo, Section 17(d) names a 24-hour telephone number and a postal address, and a secure message counts only if the bank confirms it does. Keep the date, the time, and any reference number, and report any transfer that looks wrong at once, which starts the investigation Regulation E requires whether or not an agent was involved. The outcome is a bank-side record of notice that predates any disputed transaction. Notice fixes the chronology and preserves rights the statute already grants; it does not move a burden the statute has already placed. The measures are the share of revocations with both a provider record and a bank acknowledgment, and the time between the two.

A bank faces an examiner resource that lists what a reviewer needs and what an external agent control looks like. The tactic is to name the reviewer for every agent-facing process, define what that person sees before approving, and log every override. It also means building the revocation path the bank can act on when notice arrives, which Hunton's alert recommends as well. The measures are the median and the slowest five percent of the time from a customer's revocation notice to terminated agent access, and the share of agent transactions that post after an acknowledged notice. A third is the

share of agent disputes resolved from the bank's own log rather than the customer's screenshot. Section 2.14 says documentation for a bank's own systems should let someone other than the developer audit them, and a revocation the bank cannot find in its own records sits poorly beside that objective.

Counsel and standards bodies face three sets of terms, one waterfall, and no rule. The tactic is to date the record: October 27, 2026 for Chase's agent terms, November 24, 2026 for Wells Fargo's, January 1, 2027 for Chase's device requirement, and October 23, 2026 for Spain's statute. Add the CFPB's reconsideration of its Section 1033 open banking rule, at OIRA by Hunton's November 3 account. The outcome is a timeline that shows which private allocation was in force when the first consumer dispute reaches a regulator, and whether the statute that bars waivers was raised in it.

The party that writes the terms decides who keeps the receipt. For now that party has been a provider and a bank, and the receipt belongs to the customer. The counter-pressure continues above the contracts. A statute bars waivers of the rights it grants, and a burden of proof sits with the institution in any action over consumer liability. An examiner's resource asks banks how their own vendor contracts place model losses, a competitor's proposal follows control, and claims could arise from customers who use these tools because they must. The first contested transfer will test what authority the customer granted, where the agent exceeded it, which records exist, when the bank was told, and which parts of the contract the law lets stand.

References

Board of Governors of the Federal Reserve System, Federal Deposit Insurance Corporation, & Office of the Comptroller of the Currency. (2026, April 17). *Revised guidance on model risk management* (SR 26-2). https://www.federalreserve.gov/supervisionreg/srletters/SR2602.pdf

Business Wire. (2026, April 14). *American Express debuts Agentic Commerce Experiences (ACE) developer kit and announces industry-first protection for registered agent purchases* [Press release]. https://www.businesswire.com/news/home/20260414231355/en/

Campbell, K. (2026, October 22). Gould: Banks preparing for rogue AI agents for months. *American Banker*. https://www.americanbanker.com/news/gould-banks-preparing-for-rogue-ai-agents-for-months

Consumer Bankers Association. (2026a, January). *Agentic AI payments: Navigating consumer protection, innovation, and regulatory frameworks* [White paper]. https://consumerbankers.com/wp-content/uploads/2026/01/CBA-Agentic-Symposium-White-Paper-2026-01v2.pdf

Consumer Bankers Association. (2026b, August 14). *Re: House Financial Services Committee Democrats request for information on artificial intelligence* [Letter to Ranking Member Maxine Waters]. https://consumerbankers.com/wp-content/uploads/2026/08/FINAL_CBA-Response-to-HSFC-Democrats-RFI-on-AI-in-Financial-Services-August-2026.pdf

Consumer Financial Protection Bureau. (n.d.). *Electronic fund transfers FAQs*. Retrieved November 5, 2026, from https://www.consumerfinance.gov/compliance/compliance-resources/deposit-accounts-resources/electronic-fund-transfers/electronic-fund-transfers-faqs/

Crosman, P. (2026a, October 13). Bankers grapple with liability when AI agents mess up. *American Banker*. https://www.americanbanker.com/news/bankers-grapple-with-liability-when-ai-agents-mess-up

Crosman, P. (2026b, November 2). Wells Fargo wants customers to prove they've revoked AI agents. *American Banker*. https://www.americanbanker.com/news/wells-fargo-wants-customers-to-prove-theyve-revoked-ai-agents

Electronic Fund Transfer Act, 15 U.S.C. §§ 1693g(b), 1693l. https://www.law.cornell.edu/uscode/text/15/1693g ; https://www.law.cornell.edu/uscode/text/15/1693l

Electronic Fund Transfers (Regulation E), 12 C.F.R. §§ 1005.2(j), 1005.2(m), 1005.3(c)(3), 1005.11, and Official Interpretation comments 2(m)-2 and 2(m)-3. https://www.consumerfinance.gov/rules-policy/regulations/1005/2/ ; https://www.consumerfinance.gov/rules-policy/regulations/1005/3/

Freya, A. (2026, October 23). US banks block AI agents from customer accounts over fears they could move money without permission. *International Business Times*. https://www.ibtimes.co.uk/us-banks-restrict-third-party-ai-agents-1827385

Hunton Andrews Kurth. (2026, November 3). *Major banks move to block or limit customer-directed AI agents* [Legal update]. https://www.hunton.com/insights/legal/major-banks-move-to-block-or-limit-customer-directed-ai-agents

JPMorgan Chase & Co. (2026a). *Terms of use: AI agent or agentic AI notices*. Chase.com. Retrieved November 5, 2026, from https://www.chase.com/digital/resources/terms-of-use

JPMorgan Chase Bank, N.A. (2026b, October). *Digital Services Agreement* (last updated October 2026), sec. 9 and definitions. Retrieved November 5, 2026, from https://static.chasecdn.com/content/dam/legal-agreements/library/en/dsa_la/versions/dsa_la.pdf

Lead Bank. (2026, September 29). *Moving money with agents* [Company engineering post]. https://www.lead.bank/blog/moving-money-with-agents

Ley 4/2026, de 1 de octubre, por la que se modifica el texto refundido de la Ley General de derechos de las personas con discapacidad y de su inclusión social (BOE-A-2026-20528; in force October 23, 2026). Boletín Oficial del Estado. https://www.boe.es/eli/es/l/2026/10/01/4/con ; consolidated text of Real Decreto Legislativo 1/2013, art. 7.6: https://www.boe.es/buscar/act.php?id=BOE-A-2013-12632

Morris, C. (2026, September 29). Meta's own fine print says you're responsible for every purchase its Muse AI agent makes with your card. *Moneywise*, via Yahoo Finance. https://finance.yahoo.com/technology/ai/articles/metas-own-fine-print-says-210000032.html

Office of the Comptroller of the Currency. (2026a, November 4). *Super intelligence: Supervisory framework* (OCC Bulletin 2026-54). https://www.occ.gov/news-issuances/bulletins/2026/bulletin-2026-54.html

Office of the Comptroller of the Currency. (2026b, November 4). *OCC supervisory framework for bank use of super intelligence*, secs. 2.3, 2.9, 2.10, 2.14. https://www.occ.gov/news-issuances/news-releases/2026/nr-occ-2026-97a.pdf

Spear Street Technology, Inc. (2026, August 26). *Terms of service*, secs. 3 and 10. Instinct. https://instinct.com/terms

Title III of the Americans with Disabilities Act, 42 U.S.C. §§ 12181(7)(F), 12182. https://www.law.cornell.edu/uscode/text/42/12181 ; https://www.law.cornell.edu/uscode/text/42/12182

Uniform Commercial Code § 4A-108. https://www.law.cornell.edu/ucc/4A/4A-108

Volkova, M. (2026a, October 13). JPMorgan warns of risks to consumers who rely on AI agents. *American Banker*. https://www.americanbanker.com/news/jpmorgan-warns-of-risks-to-consumers-who-rely-on-ai-agents

Volkova, M. (2026b, November 3). Banks navigate AI risks solo as regulators play catch-up. *American Banker*. https://www.americanbanker.com/news/banks-navigate-ai-risks-solo-as-regulators-play-catch-up

Waller, C. J. (2026, September 29). *Payments in the age of AI agents* [Speech]. Sibos 2026, Miami, FL. Board of Governors of the Federal Reserve System. https://www.federalreserve.gov/newsevents/speech/waller20260928a.htm

Wells Fargo. (2026). *Online Access Agreement* (upcoming version, effective November 24, 2026), secs. 2(b), 2(d), 5, 8(b), 15, 17(c), 17(d), 18, 19, 20(c), 21. Retrieved November 5, 2026, from https://www.wellsfargo.com/online-banking/online-access-agreement/upcoming/

Works by the author, with the sources supporting each

Puglisi, B. C. (2026, May 24). *The AI Risk Economy: Why insurance cannot price what governance cannot prove* (Working paper). basilpuglisi.com. https://basilpuglisi.com/ai-risk-economy-insurance-governance/

- Latest edition: Puglisi, B. C. (2026, September). *The AI Risk Economy: The Economic Override and the Liability Counter-Pressure*. Zenodo. https://doi.org/10.5281/zenodo.20824299
- W.R. Berkley Corporation. (2024, June). *Form PC 51380 00 (06-24)* [Absolute AI exclusion endorsement].
- Verisk Analytics, Insurance Services Office. (2026). *Form CG 40 47 01 26* [Commercial general liability endorsement, effective January 1, 2026].
- Crowell & Moring (2026), on Colorado SB 26-189, as cited in the paper.

Puglisi, B. C. (2026, June 13). *The liability map: The three channels through which AI creates legal exposure*. basilpuglisi.com. https://basilpuglisi.com/ai-liability-map-three-channels/

- *Moffatt v. Air Canada*, 2024 BCCRT 149.
- Directive (EU) 2024/2853 on liability for defective products (transposition due December 9, 2026).
- Regulation (EU) 2024/1689 (EU Artificial Intelligence Act), Article 99.

Puglisi, B. C. (2025). *Governing AI: When capability exceeds control*. ISBN 9798349677687.

- Latest page: Puglisi, B. C. (2026, September 11). *Checkpoint-Based Governance: Presence is not authority* [Canonical framework page]. basilpuglisi.com. https://basilpuglisi.com/cbg/
- Stanford HAI and RegulatingAI (2024), on the voluntary safety commitments, as cited in Chapter 2.
- Wilson, K., & Caliskan, A. (2024). *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*.

Puglisi, B. C. (2026, October 13). *HAIA Agents: From the definition problem to running code*. basilpuglisi.com. https://basilpuglisi.com/haia-agents-running-code/

- Latest page: Puglisi, B. C. (2026, October 13). *HAIA Agents: Who acts, and who answers* [Framework page]. basilpuglisi.com. https://basilpuglisi.com/agents/

Puglisi, B. C. (2026, September 24). *When AI becomes assistive technology: Artificially Intelligent Disability Assistance (AIDA)*. basilpuglisi.com. https://basilpuglisi.com/aida-ai-assistive-technology/

- Deposit: Zenodo. https://doi.org/10.5281/zenodo.22908271
- Assistive Technology Act, 29 U.S.C. § 3002.
- United Nations. (2006). *Convention on the Rights of Persons with Disabilities*.
- Adnin, R., & Das, M. (2024). *Proceedings of the 26th International ACM SIGACCESS Conference on Computers and Accessibility*.

Puglisi, B. C. (2026, September 24). *When AI policy becomes an accessibility barrier: ADA exposure from restrictions on Artificially Intelligent Disability Assistance (AIDA)*. basilpuglisi.com. https://basilpuglisi.com/ai-policy-accessibility-barrier/

- Deposit: Zenodo. https://doi.org/10.5281/zenodo.22940224
- Americans with Disabilities Act, 42 U.S.C. §§ 12102, 12182, 12189, 12201, 12203.
- *US Airways, Inc. v. Barnett*, 535 U.S. 391 (2002).
- Schilke, O., & Reimann, M. (2025). *Organizational Behavior and Human Decision Processes*, 188, 104405.

Disclaimer

I am not a lawyer, and this article does not provide legal advice. This is thought research and governance analysis based on public sources, cited materials, and human-AI review. It is intended to help executives, practitioners, insurers, and governance teams think more clearly about AI risk, liability exposure, and documentation practices. Readers should not rely on this article as a legal opinion, compliance determination, or substitute for qualified counsel. Any organization facing a legal, regulatory, contractual, or insurance question should consult its own attorney, broker, or professional adviser before acting.

The Other AI: Audio Briefings on Augmented Intelligence and AI Governance
Spotify | Apple Podcasts | Amazon Music | YouTube Playlist

Basil C. Puglisi, MPA
A Human-AI Collaboration

#AIassisted using the HAIA Ecosystem | CC BY-NC-SA 4.0

Free for personal, educational, and noncommercial research use with attribution. Commercial exploitation, paid productization, and enterprise commercialization require separate permission and licensing.

Suggested citation: Puglisi, B. C. (2026). The Liability Counter-Pressure Reaches the Banks: AI Agent Risk Is Being Assigned to Consumers. basilpuglisi.com.