When the Bot Signs the Wire: Why US “own-risk” waivers for AI banking shatter against emerging rules in Abu Dhabi, London, Hong Kong, and Singapore.

November 23, 2026

Introduction: Follow the Net Interest Margin, Then Read the Fine Print

In global banking, whenever chief executives invoke “customer safety” in unison on third-quarter earnings calls, prudent observers immediately check the net interest margin. Throughout October and November 2026, the largest financial institutions in the United States and Britain erected sudden technical and contractual barriers against third-party artificial intelligence agents, autonomous browser tools and personal assistants capable of logging into accounts, paying bills, and reallocating cash. Publicly, bank leadership framed the crackdown around rogue software and cybersecurity hygiene. Privately, the physics of modern bank balance sheets tell a simpler story.

Retail banking profitability rests on inertia. When households leave cash sitting in checking accounts yielding roughly 0.1%, banks lend or park those balances at 4% to 5% and pocket the spread. An autonomous agent programmed to optimize household yield does not suffer from human procrastination; it sweeps idle cash into higher-yielding treasury bills or money-market funds while its owner sleeps. Apollo Global Management chief

economist Torsten Slok warned in September 2026 that synchronized balance-sweeping could trigger an “agentic bank run” at machine speed, with industry estimates placing annual US deposit-margin erosion near $47 billion. Across the Atlantic, McKinsey estimated that AI-driven switching threatens £700 billion in low-yielding British retail deposits, where moving just 5% to 10% of balances would slash UK bank deposit profits by a fifth or more.

Faced with software that threatens both authentication perimeters and deposit spreads, US banks split into two camps in late 2026. JPMorgan Chase, Bank of America, Citigroup, and Charles Schwab opted for perimeter prohibition, blocking credential-sharing agents like Meta’s Muse outright while funneling third-party apps toward read-only open-banking pipes that can view balances without moving money. Wells Fargo chose a far more revealing path. Effective November 24, 2026, a newly minted Section 15 of Wells Fargo’s *Online Access Agreement* leaves the door ajar for AI agents, provided the customer signs up for an extraordinary legal bargain: total assumption of operational, fraud, and manipulation risk, paired with an affirmative evidentiary duty to prove revocation when shutting an agent down.

For cross-border financial institutions, corporate treasurers, and sovereign wealth allocators operating across London, Dubai, Hong Kong, and Singapore, Wells Fargo’s contractual experiment raises a decisive question. Can a regulated bank contractually blindfold itself to automated account activity while dumping the tail risk of machine error onto its depositors, or does that posture shatter against mandatory regulatory frameworks and emerging judicial precedent, exemplified by the Middle East’s first AI liability ruling in **Abu Dhabi Commercial Court Case No. 762 of 2026**?

The Contractual Baseline: Dissecting Wells Fargo’s Section 15 and the US Fault Line

Effective November 24, 2026, Wells Fargo's updated *Online Access Agreement* governs retail and commercial digital channels alike, spanning Wells Fargo Online, Wells Fargo Business Online, Wells Fargo Mobile, and Wells Fargo Advisors, all reinforced by mandatory binding arbitration and a class-action waiver. Rather than banning autonomous tools, Section 15, titled *"Third-Party Artificial Intelligence Agents or Agentic AI"*, constructs an eight-part liability shield from subsections 15(a) through 15(h):

- **Definitional Dragnet , 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"*. Read literally, the inclusion of *"non-autonomous action"* sweeps in deterministic treasury scripts, ERP reconciliation macros, and scheduled bookkeeping connectors alongside frontier language models.
- **Attribution Without Verification , Sections 15(b)–(d):** Disclaims bank control or vetting of third-party tools, strips bank privacy-policy protection from shared data, and decrees that every instruction transmitted by an authorized agent counts as the customer's personal act, with zero obligation on Wells Fargo to perform secondary verification. Under Section 2(d), if an account delegate or employee unleashes an agent, the primary account holder owns the resulting transactions, *"including those you did not want or intend"*.
- **Contractual Blindness , Section 15(e):** Declares in stark capitals that agent use is *"AT YOUR OWN RISK"*, covering software bugs and manipulated inputs such as third-party prompt injections, while stipulating: *"We do not monitor, screen, or review AI Agent activities"*.
- **The Revocation Burden , Section 15(f):** Places the entire mechanical and evidentiary burden of cutting off an agent on the customer. Account holders must revoke

access *“properly and timely”*, rotate compromised Wells Fargo credentials, notify the bank via the channel specified in Section 17(d), and retain documentary proof, screenshots, vendor termination receipts, and call logs, to produce upon bank demand. Until Wells Fargo receives formal notice, the customer remains on the hook.

- **Unilateral Kill-Switch , Sections 15(g)–(h):** Bars use by minors under 18 and reserves Wells Fargo’s right to mandate agent registration or block any tool overnight without notice or consequential liability.

Viewed through a commercial lens, Wells Fargo has drafted a free option on the agentic ecosystem. If fintech developers build sticky consumer workflows, Wells Fargo captures the transaction volume while JPMorgan and Citi block the door; if a prompt-injected browser agent wires funds to a fraudster, Wells Fargo points to Section 15(d) and demands the customer’s revocation screenshots.

The Institutional Mirror: Abu Dhabi Case No. 762 of 2026 and UAE Algorithmic Accountability

Step outside American arbitration clauses into the onshore courts of the United Arab Emirates, and the premise of Wells Fargo’s Section 15(e), *“We do not monitor, screen, or review AI Agent activities”*, runs straight into a wall of institutional duty.

On June 11, 2026, the **Abu Dhabi Commercial Court of First Instance** delivered a landmark ruling in **Case No. 762 of 2026** (ABCFI 762/2026), establishing the first judicial precedent in the UAE and wider Middle East on institutional liability for AI-driven banking risk systems. The dispute arose from unauthorized credit card transactions executed via Apple Pay and one-time password (OTP) authentication channels. Despite the cardholder immediately notifying the bank of the fraud, the bank failed to halt the bleed and sought to hold the

customer liable for the resulting card debt.

Evaluating the record under UAE Federal Decree-Law No. 35/2022 on Evidence and Federal Decree-Law No. 42/2022 on Civil Procedure, the Abu Dhabi Commercial Court discharged the customer's card liability and awarded compensation against the financial institution. Crucially, the Court treated structural deficiencies in the bank's artificial intelligence systems and architecture not as an excusable technological quirk, nor as a risk transferable by standard account terms, but as a fundamental institutional breach of the banker's duty of care.

The judgment does not stand in isolation. It operationalizes a dense UAE regulatory lattice governing automated finance:

1. **Capital Markets & SCA Rulebook Codification:** Under Securities and Commodities Authority Decision No. 13/RM/2021, as updated by Capital Market Authority Decision No. 13/RM/2026, UAE financial regulators explicitly require licensed institutions to govern four core algorithmic failure modes: *Model Drift*, *Hallucination* or illogical outputs, *Explainability Loss*, and *Algorithmic Bias*.
2. **Central Bank of the UAE (CBUAE) Supervisory Guardrails:** In February 2026, the CBUAE issued its *Guidance Note on Responsible Adoption of AI and ML*, mandating comprehensive AI model inventories, bias stress-testing, consumer human-review rights, and a strict boundary confining fully autonomous, unreviewed execution to low-risk, non-material processes.
3. **Civil Code Policing of Unfair Exclusions:** Under UAE civil transactions principles, AI possesses no independent legal personality, and onshore courts retain statutory discretion to pare back or nullify overbroad limitation-of-liability clauses in adhesion contracts that purport to absolve a regulated custodian from monitoring obvious fraud vectors.

Put bluntly: as case law stands, in Abu Dhabi or Dubai, a bank cannot deploy automated payment rails, pocket the operational cost savings of digitized servicing, insert a clause declaring *"we do not monitor or screen"* machine-driven transactions, and expect a commercial judge to enforce the disclaimer when an account is emptied in thirty seconds.

Global Financial Hubs Compared: Will London, Dubai, Hong Kong, and Singapore Adopt the Wells Fargo Playbook?

While banks across every major financial capital share Wall Street's hostility toward credential-scraping bots, **the four premier non-US financial hubs, London, Dubai, Hong Kong, or Singapore, can cleanly import Wells Fargo's "permit credentialed agents at the customer's sole risk" model.** Instead, regulatory physics in all four hubs is forcing banks toward a two-stage sequence: **block unverified credential-sharing today, then channel agents onto regulated, auditable API rails tomorrow**.

Financial Hub	Current Bank Stance on AI Agents (Late 2026)	Governing Regulatory & Liability Architecture	Viability of Wells Fargo's "Sole Customer Risk" Model
London (UK)	**Precautionary Perimeter Block:** HSBC, Barclays, Lloyds, NatWest, and Nationwide moved in Nov 2026 to block/restrict transactional AI agents ahead of Meta Muse's UK rollout.	FCA Consumer Duty & Senior Managers Regime; UK Strong Customer Authentication (biometric app sign-off); PSR mandatory APP scam reimbursement; FCA developing trusted agentic finance framework.	**Low for Retail / Moderate for Corporate:** FCA treats failure to catch obvious fraud anomalies as a senior-management breach; biometric SCA physically blocks password-scraping agents anyway.

Dubai / UAE	**Controlled Sovereign Buildout:** Emirates NBD advocates bank-built agentic frameworks and hosts GPT models on onshore UAE servers; 53% of DIFC firms pilot agents, only 14% deploy live.	**ABCFI Case No. 762/2026** (institutional liability for AI fraud-detection failures); CBUAE Open Finance Regulation & Feb 2026 AI Guidance; SCA/CMA Decision 13/RM/2026; DFSA cyber/AI oversight.	**Rejected for Institutional Failures:** Courts hold banks liable for deficient fraud-detection models; third-party agent access is being routed into CBUAE-licensed Open Finance APIs rather than raw passwords.
Hong Kong	**Supervised Sandbox & Hard Human Gates:** Banks restrict external credential-sharing while testing agentic workflows inside regulatory sandboxes.	HKMA **ART Framework** & *Supervisory Guidance* (Ref. B1/15C, Oct 15, 2026); **Project SAGE** & *Discussion Paper on Agentic A.I. Governance* (Nov 2, 2026).	**Incompatible with Supervisory Expectations:** HKMA expects defined delegation scopes, deterministic fail-safes, real human approval windows, and immutable action logs, the exact opposite of “we do not monitor”.

Singapore	**Phased Risk-Proportionate Governance:** DBS, OCBC, and UOB enforce strict anti-credential-sharing controls and app-bound cryptographic tokens.	MAS *Guidelines on AI Risk Management* (finalized Oct 7, 2026; phased Oct 2027 / Oct 2028) covering AI agents; MAS/IMDA *Shared Responsibility Framework* (SRF) for unauthorized/scam flows.	**Preempted by Waterfall Liability:** Under the SRF and MAS AI Guidelines, banks must fulfill real-time fraud surveillance duties first; blanket contractual disclaimers cannot override breached bank duties.

1. London: Biometric Friction and the £700 Billion Deposit Defense

On November 8, 2026, Britain’s five largest lenders, HSBC, Barclays, Lloyds, NatWest, and Nationwide, confirmed precautionary restrictions blocking third-party AI agents from executing customer account transactions, despite zero reported incidents inside UK accounts to date. Structurally, London differs from New York in two decisive respects. First, UK Strong Customer Authentication rules already force per-payment biometric approval inside a bank’s native mobile app, breaking password-based screen-scraping agents before they reach the ledger. Second, under the FCA’s Consumer Duty and Senior Managers and Certification Regime, if a bank’s automated fraud-monitoring system misses glaring anomalies, regulators treat the lapse as an institutional governance failure rather than a retail user’s tough luck. Rather than copying Wells Fargo’s liability dump, the FCA is working with industry to construct a formal agentic finance framework built around verified **agent identity, scoped delegated authority, and clear institutional accountability**, keeping read-only open banking alive while gating payment initiation behind regulated rails.

2. Dubai and the UAE: Sovereign Infrastructure Over Password Scraping

Across the DIFC, ADGM, and onshore UAE banking markets, the regulatory ethos favors sovereign-grade infrastructure over laissez-faire credential sharing. The DFSA's *AI Survey 2026*, published November 2, 2026, revealed that while 71% of authorized DIFC firms now use AI and 53% are testing or piloting AI agents, only 14% have moved agents into production, reflecting cautious supervisory pacing on third-party concentration risk. Leading regional lender Emirates NBD has explicitly argued that banks must construct their own governed agentic frameworks rather than surrendering the interface to outside bots, going so far as to run OpenAI models on dedicated servers inside the UAE. Combined with the CBUAE's Open Finance Regulation, which standardizes tokenized API access, and the judicial warning shot of **Case No. 762 of 2026**, UAE banks have zero incentive to tolerate raw password-sharing agents while gambling on an untested contractual waiver.

3. Hong Kong: Project SAGE and Deterministic Guardrails

Hong Kong has moved faster than almost any peer jurisdiction to write an explicit architectural blueprint for agentic banking. Following its October 15, 2026 rollout of the **ART Framework** (Accountability, Resilience, and Trust) and updated *Supervisory Guidance* (Ref. B1/15C), the HKMA launched **Project SAGE** (*Securing Agentic Finance with Governance and Enablement*) alongside its *Discussion Paper on Agentic A.I. Governance in the Banking Sector* on November 2, 2026, with industry comments closing February 1, 2027. Four HKMA principles directly invert the Wells Fargo approach: agents must act strictly within bounded delegations of authority; failed deterministic checks must automatically block or escalate the transaction; human-on-the-loop approvals must afford genuine time to reject actions before settlement; and banks must preserve granular audit trails showing prompts, permissions, and approvals.

Under Project SAGE, revoking or bounding an agent becomes a cryptographic infrastructure feature audited by the bank, not a screenshot scavenger hunt outsourced to the depositor.

4. Singapore: Waterfall Accountability and the October 2026 MAS Mandate

Singapore's Monetary Authority closed Consultation Paper P017-2025 on October 7, 2026 by issuing final, binding-expectation *Guidelines on Artificial Intelligence Risk Management* applicable to all financial institutions and explicitly encompassing generative AI and AI agents. Phased into force on October 7, 2027 (governance, AI inventories, and materiality assessments) and October 7, 2028 (full lifecycle and infrastructure controls), with targeted MAS industry engagement on agentic AI slated for 2027, Singapore treats autonomous software as a supervised operational risk. Moreover, Singapore's *Shared Responsibility Framework* (SRF) allocates digital fraud losses through a strict "waterfall" mechanism: financial institutions sit on the top rung and must prove they met real-time transaction-surveillance, cooling-off, and notification duties before passing losses downstream. While deliberate password sharing by a retail user still breaches user-protection duties in Singapore, a bank that adopted Wells Fargo's *"we do not monitor or screen"* posture would fail the first rung of MAS's waterfall immediately.

The Legal Reality: Structural Issues Along the Fault Line

Beneath the immediate friction over blocked logins and updated terms of service lies a deeper set of structural problems that neither bank boilerplate nor developer optimism has yet resolved. The first is the **collapse between identity and delegation**. Legacy online banking was built on a binary premise: whoever presents valid credentials is the customer. Autonomous software breaks that assumption. As long as read-only data aggregation and payment-execution authority travel over the same shared credentials, account holders,

particularly corporate treasurers operating outside retail consumer-protection caps, face asymmetric exposure whenever an automated workflow misinterprets an instruction or falls prey to external manipulation. Separating balance visibility from the power to move funds is no longer a mere technical preference; it is the boundary line where liability begins.

Closely tied to delegation is the **evidentiary mismatch around revocation**. Contractual terms that obligate a customer to cut off an agent and prove after the fact that termination was timely assume a clean, linear paper trail. Software permissions rarely fail so neatly. Persistent session tokens, cached credentials, and millisecond execution windows mean that future disputes will turn less on what a depositor intended than on which party holds verifiable, time-stamped records of when authority ended and whether the underlying rail actually honored the cutoff. Where revocation remains an informal tug-of-war over passwords rather than an instantaneous cryptographic switch, evidentiary ambiguity will consistently favor the party that drafted the agreement.

For institutions operating across borders, however, the sharpest tension is the **territorial limit of contractual blindness**. A disclaimer stating that a bank neither monitors nor screens agent-driven activity may deter claims in domestic consumer arbitration, but it does not travel well. Across Abu Dhabi, London, Hong Kong, and Singapore, courts and supervisors increasingly treat anomaly detection, deterministic pre-settlement checks, and clear delegation boundaries as non-delegable institutional functions. Relying on passive “own-risk” waivers in jurisdictions that hold banks accountable for the blind spots of their own automated risk engines invites the worst of both worlds: operational complacency paired with unenforceable legal shelter.

Ultimately, neither blanket bans nor one-sided liability dumps represent a durable commercial equilibrium. Banks have every incentive to capture the transaction volume of automated

commerce, provided they do not surrender low-cost deposits to unpoliced yield-sweeping along the way. The lasting resolution therefore points toward governed “Know Your Agent” infrastructure, dedicated rails where an agent’s identity, spending ceiling, and counterparty permissions are verified before settlement, and where liability tracks actual operational control rather than fine print. Until those rails mature, the gap between what autonomous software can execute and what banking law will excuse remains the central risk on the ledger.

Mahmoud Abuwasel serves on the arbitral rosters of the American Arbitration Association (AAA), the Dubai International Arbitration Centre (DIAC), the Korean Commercial Arbitration Board (KCAB), and other international arbitral institutions, and is featured on the AAA-ICDR Blockchain and Web3 Disputes Panel. Advising extensively on high-stakes crypto, digital-asset, and algorithmic disputes, the earliest proving ground for litigation over autonomous execution, delegated credentials, and cryptographic revocation, he is the author of UAE Crypto Litigation and a contributing author to Kluwer Law International’s Crypto Asset Disputes. Operating across the fault line between international arbitration and mandatory onshore judicial supervision, including before the UAE courts and major global financial centers, Wasel & Wasel counsels clients on the governance, enforcement, and liability architecture of automated and cross-border finance.

Author: Mahmoud Abuwasel
Title: Partner – Disputes
Email: mabuwasel@waselandwasel.com
Profile: https://waselandwasel.com/about/mahmoud-abuwasel/

Lawyers and consultants.
Tier-1 services since 1799.
www.waselandwasel.com
business@waselandwasel.com