

AI has industrialized financial scams, with recovery rates near zero

Sep 25, 2026 · 65 items read from 62 outlets

40

stories clustered from the day's intake

Today, by the numbers

65 ITEMS READ	62 OUTLETS	19 SIGNALS PUBLISHED
51 PAPERS ANALYZED	40 STORIES	

TAKEAWAY

40 stories clustered from 65 items across 62 outlets.

Markets at press time

- **Nvidia 225** +0.2%
- **Microsoft 517** +3.9%
- **Alphabet 344** +0.4%
- **Meta 750** -3.6%
- **Amazon 250** +0.3%
- **TSMC 452** +0.2%
- **Semis (SOXX) 575** +1.5%
- **Utilities (XLU) 39.33** -0.1%

TAKEAWAY

An unofficial feed, priced at edition time.

AI has industrialized financial scams, with recovery rates near zero

- **Why it matters** Scam losses hit roughly \$20 billion in 2025 while banks detect only 10% to 15% of losses and recover 2% to 5%, so prevention is the only viable defense. The report flags misaligned liability across banks, social platforms, and telecoms as the reason prevention stays underfunded.
- **By the numbers** Victim losses were estimated at roughly \$20 billion in 2025, up 26% year over year.
- **Coverage** 1 outlet
- **Go deeper** [Read the source](#)

EU regulators: the quantum deadline may beat commercial quantum

- **Why it matters** European financial regulators warn Q-Day could arrive before quantum computers are commercially useful and want post-quantum cryptography adopted by end of 2026. Harvest-now-decrypt-later attacks make encrypted data stolen today a future liability.
- **By the numbers** Regulators urge adoption of post-quantum cryptography by the end of 2026.
- **Coverage** 1 outlet
- **Go deeper** [Read the source](#)

The Clearing House picks Quant for a 25-bank tokenized deposit network

- **Why it matters** A settlement network for tokenized deposits across 25 major U.S. banks, connected to existing rails like RTP and CHIPS, is a concrete step toward regulated on-chain money. It reaches participating banks in the first half of 2027, setting a shared standard banks will have to plan around.
- **By the numbers** The Clearing House processes over \$2 trillion daily.
- **Coverage** 1 outlet
- **Go deeper** [Read the source](#)

Snorkel AI's \$3.5B raise shows training data is its own value chain

- **Why it matters** An eighteenfold revenue surge to \$375M annualized on a shift to hybrid synthetic-plus-expert datasets confirms frontier model development still runs on proprietary, verified data, not public web scraping. Data provenance and quality assurance are emerging as a distinct spend layer in AI infrastructure.
- **By the numbers** Snorkel raised \$350 million at a \$3.5 billion valuation with \$375 million in annualized revenue.
- **Coverage** 1 outlet
- **Go deeper** [Read the signal](#)

Alibaba claims China's most powerful AI chip and a 20GW data center bet

- **Why it matters** The Zhenwu V900, clusterable into units of 500,000, is a concrete non-US vendor capability claim backed by a 2032 global capacity target. The open question is whether Alibaba can sustain that capex without US supply chains.
- **By the numbers** Alibaba targets 20 gigawatts of data center capacity by 2032.
- **Coverage** 1 outlet
- **Go deeper** [Read the signal](#)

Banks are fighting over small-business bankers, and AI is the hiring weapon

- **Why it matters** Business banker openings at major competitors jumped 105% as banks chase business owners for cross-selling. BofA is pairing aggressive hiring with AI tools to speed underwriting, and targeting 30% client share, up from 27.6%.
- **By the numbers** Job openings at major competitors rose 105% over the past year.
- **Coverage** 1 outlet
- **Go deeper** [Read the source](#)

THE COLUMN

The Week in Money Nobody Can Recover

DAILY EDITION • Sep 25, 2026

The Week in Money Nobody Can Recover

■ **Let us close** the week with an accounting theme: everywhere you look, someone is spending money to defend against a bill that has not arrived yet.

The Week in Money Nobody Can Recover

- **The scam economy** has an allocation problem. The sharpest number in the Thomson Reuters report is not the \$20 billion in losses, it is the 2% to 5% recovery rate. When recovery is functionally zero, fraud stops being a reimbursement expense and becomes an industry, which is what generative AI has made it. The report's diagnosis, misaligned liability among banks, platforms, and telecoms, is really a claim that nobody owns the externalized cost, so nobody invests in preventing it. What would strengthen that reading: any regulatory move assigning liability upstream, say to platforms that host the lures. What would weaken it: banks quietly improving detection well above that 10% to 15% range, which would suggest the market can fix this without a mandate.

The Week in Money Nobody Can Recover

- **Everyone is prepaying** future liabilities, including cryptographers. European regulators warn in the Q-Day advisory that harvest-now-decrypt-later attacks mean data stolen today is already compromised in 2032, so they want post-quantum cryptography adopted by end of 2026. It is the same structure as the fraud story in reverse: pay now to avoid a bill later, with the payment due before anyone can prove the threat is real. Meanwhile the actual plumbing is being rebuilt too. The Clearing House's tokenized deposit network, spanning 25 banks and bridging to RTP and CHIPS by mid-2027, is the rare infrastructure story where the standard arrives before the fight over the standard. That coordination is worth more than the tokenization itself.
- **The data layer** keeps getting paid. Snorkel's \$3.5 billion raise is the week's clearest evidence that verified training data is a distinct line item, not a byproduct. Eighteenfold revenue growth to \$375 million annualized tells you enterprises will pay for provenance, which supports the broader view that training-data IP liability is now a quantifiable balance-sheet cost for AI vendors. If verification is a cost center buyers demand, the vendors who cannot document clean data get excluded from procurement, and the industry segments on paperwork. The counterpoint is that synthetic data could eventually make the whole verification layer cheap to substitute; Snorkel's own hybrid model hedges that, which is telling.
- **One honest wrinkle:** all this prepaying assumes the spending produces revenue. The same week brings Alibaba's Zhenwu V900 and a 20-gigawatt bet against a backdrop where announced aggregate capex materially exceeds what will be deployed. Betting on 2032 capacity is the ultimate pay-now-for-maybe-later trade.
- **The banks fighting** over small-business bankers with AI underwriting tools at least bill someone this year.

Links from the column

- [Thomson Reuters report](#)
- [Q-Day advisory](#)
- [The Clearing House's tokenized deposit network](#)
- [Snorkel's \\$3.5 billion raise](#)
- [training-data IP liability is now a quantifiable balance-sheet cost for AI vendors](#)
- [Alibaba's Zhenwu V900 and a 20-gigawatt bet](#)
- [announced aggregate capex materially exceeds what will be deployed](#)

Things happen

- **Bank of America to hire 1,000 more apprentices, invest \$150M in workforce development**

foxbusiness.com · T3

- **AI adoption in Chinese enterprises boosts productivity but raises ...**

ilo.org

- **Peer-reviewed review frames AI data center grid integration as requiring multi-layer energy storage coordination**

sciencedirect.com

- **St. Louis Fed study links industry-level AI use to US post-pandemic productivity recovery**

frbsf.org

- **OpenAI proposes global AI safety standards and incident tracking to US-China dialogue partners**

axios.com

- **Blizzard union contract establishes AI usage rules and layoff terms**

hrexecutive.com

- **Korea HBM exports to Malaysia surge fivefold as advanced packaging capacity diversifies**

biz.chosun.com

Things happen

- **China's Supreme People's Court issues first adjudication guidance for AI disputes, requiring training-data evidence from developers**

iam-media.com

- **OpenAI initiates early-phase third-party safety evaluations for AI models**

bloomberg.com

- **OpenAI and Anthropic CEOs to address UN Security Council on AI safety amid model autonomy incidents**

cnn.com

- **California Executive Order N-9-26 directs study of mandatory AI kill switch and onsite lab auditors**

ppc.land

- **ACM study finds GitHub Copilot efficiency gains do not translate to throughput increases**

cacm.acm.org

- **Anthropic releases Claude Opus 5.5 with 40% lower inference cost and higher agentic reliability**

anthropic.com

- **OpenAI reduces GPT-6 Sol and Luna API costs by 50 percent with improved factuality**

techcrunch.com

Things happen

- **IC-Link joins TSMC 3DFabric Alliance, securing European co-design access to CoWoS capacity**

techtimes.com

- **CFR report details China's persistent critical-mineral leverage post-Busan and gaps in Western de-risking**

cfr.org

- **NIST CAISI finds Z.ai's GLM-5.3 is most cyber-capable open-weight model to date, but lags US frontier by about four months**

nist.gov

- **DOJ Statement of Interest Argues Against Mandatory Licensing for AI Training Data**

news.bloomberglaw.com

- **EU Watchdog Makes Digitization A Priority For Regulators - Law360**

law360.com · T2

- **Lovable's annualized revenue crosses \$600M as vibe coding takes off**

techcrunch.com · T2

- **Decoding the New Premium Traveler - Skift**

skift.com · T2

Things happen

■ [Bank of America plans to double AI budget next year](#)

[ciodive.com](#) · T2

■ [ElevenLabs' CEO on margins, IPO timing, and telling customers they're talking to a bot](#)

[techcrunch.com](#) · T2

What builders are reading

- [Using LLMs to trace alchemical knowledge and decode 17th century letters](#) 155 points · 35 comments
- [Classified Estimates Show the NSA Is Paying Billions to Test AI Models](#) 76 points · 30 comments
- [U.S. appeals court upholds designation of Anthropic as supply chain risk](#) 54 points · 11 comments
- [Microsoft Abandons Personal AI Chatbot Race with Copilot Reboot](#) 32 points · 23 comments
- [Jevmem - automatic project memory for Claude Code, built on Jev](#) 24 points · 9 comments

Research

- **IndicBankBench: Evaluating Safety and Reliability of Language Model Assistants in Indian Retail Banking** RBI and other financial regulators get a reproducible instrument for auditing whether deployed banking assistants confirm before writes and avoid fabricated identifiers, which maps directly onto FREE-AI governance requirements. It documents that agentic reliability under nonzero temperature is materially lower than single-run benchmark scores suggest, quantifying the eval-to-reliable-deployment gap in a high-stakes domain. If banking assistants need human oversight on 40%+ of interactions, automation saves in retail banking support remain bounded near-term rather than substitutive.
- **Where Does Exactly-Once Live? Model, Harness, and Tool-Contract Effects on Duplicate Side Effects in LLM Agents** Shows that model-side verification reasoning has a provable ceiling: no model can be exactly-once on ambiguous faults without contract support, redirecting capability effort toward tool contracts and middleware. Teams deploying agents for operational write tasks (billing, deploys, notifications) need idempotency-aware tooling and cannot rely on better models or harness swaps alone.
- **Same Bit Width, Different Outcomes: Post-Training Quantization of Text-to-Speech Across Architectures** Engineers deploying quantized speech or generative models on-device learn that bit width is not a sufficient specification, and that component ablation plus per-layer GPTQ is the actual path to working 4-bit deployments. Vendors of on-device TTS and edge inference tooling can treat the 0.60x latency 4-bit Supertonic path as evidence that aggressive quantization is shippable, but only with per-model validation.
- **Sharp Limits for Honest Uncertainty in Hard-Budget Repeated Evaluation** Teams running repeated agent or model evals get a principled answer to how to split budget between new tasks and repetitions, potentially getting materially tighter intervals (about 30% median width reduction) from the same inference spend.

Research

- **Human-AI-Powered Hypothesis Testing: Cost-Aware Selective AI Scoring and Sequential Human Escalation** Speaks to whether AI-assisted quality review replaces humans wholesale or restructures the work into selective escalation of hard cases, keeping fewer but more targeted human review hours. Reframes LLM-as-judge from a capability question (make the judge better) to a statistical design question (control residual judge error formally), which changes how eval pipelines should be built.
- **When Fancy Eviction Fails: Rethinking Cache Replacement For LLM Prefix Reuse** Serving-infrastructure engineers get a workload characterization showing recency is the right foundation for prefix-cache eviction, with compute-aware eviction and block-level granularity as the productive directions. Inference providers' unit economics depend on prefill compute, and this bounds how much cache-policy sophistication can reduce serving cost relative to simple LRU.
- **No More Free Lunch: Corpus Task Complexity Matters as Corpora Grow** Architecture teams choosing sparse or hybrid attention should know that low-CTC benchmark suites systematically understate the cost of those choices on harder corpus tasks. Knowledge workers hoping models will audit large document sets for contradictions or unmatched items face a task class where current long-context scaling demonstrably fails.
- **AI-Moderated Interviews for Market Research and Digital Twins Calibration** AI-moderated interviews offer a roughly 5x cost reduction over human moderators with comparable depth, threatening the economics of traditional qualitative research agencies and validating the revenue case of AI-native interview startups. Human moderators charging \$200 to \$400 per hour are matched on content quality by AI at a fraction of the cost, with only the emotional-engagement advantage left to defend their pricing. Shows a measurable gap between fluent conversational capability and useful simulation: AI-moderated richness does not yet transfer into better quantitative prediction of real consumer responses.

Blind spots

- [Fortune](#)
- [Fortune](#)
- [Cheniere Energy \(LNG\) Q4 2025 Earnings Call Transcript](#)
- [UNGRFAU:SP - United CIO Growth Fund - Bloomberg.com](#)
- [FFAI: Faraday Future Intelligent Electric Inc Stock Price Quote](#)
- [ISTK: Intellistake Technologies Corp Stock Price Quote - Bloomberg](#)

Sources

- **finance.yahoo.com** 6 items, tier 3
- **ffnews.com** 5 items, tier 2
- **techcrunch.com** 5 items, tier 2
- **bankingdive.com** 4 items, tier 2
- **federalreserve.gov** 2 items, tier 1
- **americanbanker.com** 2 items, tier 2

TAKEAWAY

62 outlets read today

Sources

- **paymentsdive.com** 2 items, tier 2
- **federalregister.gov** 2 items, tier 1
- **cnbc.com** 2 items
- **bloomberg.com** 2 items
- **genfinity.io** 1 items
- **financefeeds.com** 1 items

TAKEAWAY

62 outlets read today

C L O S E

Every link resolves to a stored record or its source.

Sep 25, 2026