PAYMENTS AND INFRASTRUCTURE

The Authorization Gap

Agentic AI in financial services is stalled by regulation, not capability

By Diego Molina — Mediavertical

Agentic AI has already moved from pilot to production inside most banks, but the ceiling on how far it can go is not a modeling problem — it is a gap in the law over who is permitted to authorize a payment.

Adoption is mainstream, not experimental

More than seven in ten banking firms are already using agentic AI to some degree, according to EY's 2026 Global Financial Services Regulatory Outlook. The share of finance teams using agentic AI grew by more than 600% year over year, reaching 44% in 2026 per a Wolters Kluwer survey — a pace of adoption that has outrun almost every forecast made in 2024 and 2025.

Core banking providers have followed the demand rather than led it. Fiserv's agentOS, an operating system purpose-built for agentic AI in banking, already has six financial institutions co-developing use cases, with two running agents in beta ahead of a targeted broad release this August.

The engineering problem is largely solved

Deployment has concentrated exactly where you would expect a mature capability to land first: procedural, auditable workflows. Fraud detection is the clearest case — 87% of financial institutions already deploy some form of AI- or ML-driven fraud scoring, with card networks and payment processors evaluating transactions and triggering step-up authentication in real time.

Vendors have converged on the same design principle: governed environments where agent decisions are traceable and auditable end to end. FIS, for instance, has partnered with an AI lab to bring agentic tooling into financial-crimes workflows, with early bank deployments built around exactly that kind of auditability rather than raw autonomy.

The real constraint is legal authorization, not model quality

The friction point in agentic payments sits in the law, not the model. Strong Customer Authentication rules under PSD2 in the EU and UK require clear human authorization for payment orders, and no current

regulatory mechanism treats an AI agent as equivalent to a human payer. That is a structural gap, not a temporary one — it exists regardless of how capable the underlying model is.

The IMF has begun treating this as an infrastructure question in its own right, framing autonomous agents as part of critical digital infrastructure that requires governance mechanisms allowing supervisory intervention even while agents operate independently. Trust, not technology, is the obstacle financial institutions cite most often: roughly 87% name trust as the biggest barrier to agentic payment adoption, and nearly 78% expect fraud exposure to rise as agentic commerce scales.

What this means for payments builders and operators

The platforms that win this cycle will be the ones designed around auditability and human-checkpoint architecture from the first line of code, not retrofitted with compliance controls after the fact. Regulatory clarity — how the EU AI Act's obligations land through the rest of 2026, and how UK regulators resolve the authorization question — will set the pace of adoption more than any single model release.

For anyone building payments or treasury infrastructure today, the practical takeaway is sequencing: governance design is the prerequisite, not an add-on that follows model selection.

If the bottleneck on agentic payments is legal authorization rather than model capability, should payments infrastructure be built around today's rules — or around the assumption that authorization frameworks will eventually have to change to accommodate autonomous agents?