

Agentic Finance

Building the
foundations for
trusted, bounded
autonomy

CONTENT

- 06** Voices in this article
- 07** Introduction
- 08** The shift to compliant payment
- 10** Why tokenised money matters
- 12** The Stack: Six layers that have to work
- 16** What the respondents are seeing today
- 18** How respondents are tackling the challenges
- 21** Three barriers that remained unsolved
- 24** Looking ahead: Where will adoption first arrive
- 27** Conclusion: Bounded autonomous finance

Information as of 14 October 2026

Foreword

For decades, the financial services industry has focused on making money move faster.

From electronic banking and card networks to real-time payments and API-driven financial services, each wave of innovation has reduced friction in how value is transferred, settled and managed. Yet throughout this evolution, one fundamental assumption has remained largely unchanged: humans decide when money should move, and systems execute those instructions.

Agentic payments challenge that assumption.

They describe a future in which AI agents not only analyse information or recommend actions, but increasingly initiate, execute and optimise financial transactions on behalf of individuals and organisations within predefined mandates. This represents one of the most profound shifts in financial services architecture since the advent of modern payment infrastructure.

What makes this shift possible is the convergence of two technologies advancing at extraordinary speed. The first is artificial intelligence, which is evolving beyond information retrieval and content generation towards reasoning, planning and autonomous decision-making. The second is the emergence of tokenised and programmable forms of money, enabling value to move continuously, conditionally and at machine speed. Together, these developments are laying the foundations for an economy in which autonomous agents become active participants in financial activity.

Agentic AI creates autonomous economic actors capable of making decisions. Tokenised money, whether in the form of stablecoins, tokenised bank deposits or future forms of digital cash, provides the settlement layer that agentic finance requires. It allows value to be embedded into workflows, linked to programmable conditions and settled with the speed, granularity and certainty that autonomous systems require. As AI agents increasingly participate in commerce and financial operations, the availability of trusted, scalable and programmable forms of money will become a critical enabler of adoption.

Yet one conclusion emerged consistently from the organisations that contributed to this research: the future is not one of unrestricted autonomy.

Across settlement infrastructure providers, payment networks, identity platforms, financial applications and emerging agent ecosystems, there was broad consensus that adoption will succeed only if autonomous systems operate within clearly defined boundaries. The market is not building open-ended autonomous finance. Rather, it is building what this report describes as bounded autonomous finance; a model in which AI agents act within explicit mandates, governed by permissions, controls, auditability and human accountability.

This distinction matters.

The central challenge facing agentic payments is no longer whether the technology can work. Increasingly, the more important questions concern trust. How do we verify the identity of an AI agent? How do we ensure it operates within approved limits? Who remains accountable when autonomous decisions create unintended outcomes? And how do institutions provide the transparency and auditability necessary for customers, regulators and stakeholders to place confidence in machine-executed financial activity?

These questions surfaced repeatedly throughout our research, regardless of where participants sat within the emerging agentic payments ecosystem.

This report was commissioned to deepen our understanding of these realities. Through structured engagement with organisations spanning settlement infrastructure, payments, identity, orchestration, credit and agent-native applications, we sought to capture insights from those actively shaping the market rather than merely observing it.

Their perspectives reveal a market that is both earlier and more advanced than many realise. Earlier, because critical questions around identity, liability and interoperability remain unresolved. More advanced, because practical deployments and real-world experimentation are already taking place today.

Several themes stand out.

Machine-to-machine payments are emerging as one of the earliest and most compelling applications of agentic finance. Enterprise use cases spanning procurement, treasury and financial operations are advancing rapidly, supported by the clear efficiency gains that autonomous execution can unlock. At the same time, applications involving greater discretion, judgment and loss-bearing responsibility continue to require more mature governance frameworks before wider adoption can occur.

Taken together, these findings reinforce an important lesson: the future of autonomous finance will be determined not only by technological capability, but equally by the strength of the trust frameworks that surround it.

At UOB, we believe innovation delivers its greatest value when it is pursued with both ambition and responsibility. Throughout our history, progress has come not simply from adopting new technologies, but from embedding them within systems that customers, regulators and society can trust.

Agentic payments represent a compelling opportunity to reimagine how value moves in an increasingly digital economy. Realising that opportunity, however, will require collaboration across financial institutions, financial market infrastructure operators, technology providers, regulators and the broader ecosystem. A whole-of-ecosystem approach is needed to establish the foundations of trust, governance and accountability that sustainable adoption demands, particularly around identity, liability and interoperability.

It is our hope that this report contributes meaningfully to that effort. By bringing together the perspectives of those building the next generation of financial infrastructure, we aim to provide practical insights into one of the most significant developments emerging at the intersection of artificial intelligence and finance.

The transition from human-directed transactions to trusted autonomous execution will not happen overnight. But the foundations are being laid today, and the choices we make now will shape how the next era of financial services evolves.



Alvin Eng

Managing Director & Head
Enterprise AI, Innovation Group



Yip Kah Kit

Executive Director
Head of Blockchain and Digital Assets

Voices in this article

This piece draws on written submissions from eight respondents operating across the payments, infrastructure and digital-asset stack. These companies are privately held and not listed in the public domain.

UOB Venture Management (UOBVM) has investments in certain participating companies. Their views are their own and do not necessarily reflect the views of UOB.

Company facts, figures, client references and metrics cited throughout this article are self-reported by the respective companies.



Bluwhale

an AI-powered financial platform that connects traditional and digital assets, using AI agents and blockchain technology to help users manage and optimise their finances

Eigenteam

provides data aggregation and AI-powered analytics solutions to help businesses improve operational performance and decision-making

Kasu Finance

is a digital lending platform that connects investors with real-world private credit opportunities, using blockchain technology to facilitate business lending

NetMind.AI

an AI platform providing access to AI models, GPU computing and tools for building and deploying AI applications

OKX

a global cryptocurrency platform offering digital asset trading, a self custody wallet and access to decentralised applications and blockchain services

Polyhedra Network

a blockchain and AI infrastructure platform that uses zero-knowledge proof technology to enable secure, scalable and verifiable AI applications and cross chain connectivity

Transak

a global payments infrastructure platform that enables users and businesses to move between traditional currencies, stablecoins and other digital assets

Velo Labs

connects traditional finance and Web3 through cross-border payments, digital asset settlement and PayFi infrastructure

Introduction

Imagine a not so distant future where a corporate treasurer wakes to find part of the company's cash already at work.

Overnight, an AI agent had watched balances across currencies, spotted idle funds, checked the treasury rules it had been given, chosen a settlement route, converted part of the surplus into tokenised money, moved it across borders and parked the result in an approved yield instrument.

Nobody keyed in a payment instruction by hand. No one was woken for a 2 a.m. approval. The treasurer had set the objective, the risk limits, the permitted instruments and the escalation thresholds. The agent worked inside those lines.


The promise of agentic payments, is what makes financial institutions uneasy. The transaction is no longer just automated after a human has decided. The decision to initiate, route and optimise it is increasingly handed to AI agent.

That does not mean AI agent is getting unrestricted access to money. The strongest message from the eight companies whose submissions inform this piece runs close to the opposite. What the market is building is bounded autonomous finance: agents operating within explicit mandates, on programmable money, with controls, audit trails and accountability built around them.

Agentic payments, meaning transactions initiated, executed and optimised by autonomous agents working within defined mandates, rank among the larger architectural shifts in financial services since real-time payments arrived.

Two technologies are converging to drive it. One is AI that can coordinate complex financial workflows. The other is tokenised money that settles continuously, around the clock, with the programmability, availability and speed that autonomous systems require. Technical feasibility is no longer the question. What matters now is governance, trust and accountability: who authorised the agent, what it may do, how permissions are enforced, and who is responsible when something goes wrong.

Inside a bounded mandate, the overnight treasury run step by step



Source: UOBVM, UOB Innovation Group (UOBIG)


01 The shift to compliant payment

From can an agent pay to can it pay compliantly

Every wave of financial technology has promised to make money move faster. Correspondent banking was enhanced by global financial messaging networks such as SWIFT. Real-time payment systems compressed settlement windows, card networks digitised commerce at scale, and API-based fintech made financial services programmable from the application layer. Through all of it, one assumption held: humans decided when money should move, and systems carried out the instruction.

Agentic payments break that assumption. Now humans set the policy, the objective and the boundaries, and the agent watches conditions, weighs options and executes inside the mandate. Oversight does not disappear. It shifts from approving each transaction to governing the rules the agent runs on.

Figure 1 – Bounded autonomy is a loop, not a switch



Source: UOBIG, UOBVM

As Figure 1 illustrates, bounded autonomy is not a transfer of control from humans to machines. It is a governance loop. Humans define objectives, limits and permitted actions. Agents execute within those boundaries. Controls enforce the mandate and record every step. When uncertainty arises or thresholds are breached, authority returns to a human. The question is therefore not whether agents are acting autonomously, but whether that autonomy remains observable, enforceable and accountable.

The clearest sign that this is real has less to do with whether agents can pay than with the question companies say they now get asked. Polyhedra put it most sharply: the conversation has moved from “Can an agent pay?,” which was always trivially yes, to “Can an agent pay in a way my compliance officer will sign off on?,” which is mostly still no. The concrete tell, it adds, is that risk and procurement are now in the room, not just innovation labs.

The market evidence sits underneath that shift. Industry participants across payments, financial infrastructure and technology ecosystems have begun introducing and testing agent-oriented payment capabilities. Respondents consistently reported that conversations have shifted away from basic technical feasibility towards practical questions around accountability, spend controls, identity verification and liability. What was once primarily the domain of innovation teams is increasingly attracting the attention of risk, compliance and operational stakeholders responsible for governing adoption at scale. NetMind describes buyers who no longer ask whether agentic payments are possible, but ask instead who is liable, how spend is bounded, and how an agent’s identity is verified. Transak reads the same shift as a change of address inside the institution, moving from the innovation team that treated it as a curiosity to the payments, to the risk and compliance teams that now have to put controls around it. For OKX, the tell is executive attention: the question has moved from “Is this real?” to “Where do we sit in this?”

NetMind captures the implication most clearly. Everyone expected intelligence to be the bottleneck. Models improved faster than that, and governance became the scarce thing. Different companies arrive at the same place from different layers. What makes the moment matter is the kind of shift it is. The challenge is no longer moving money faster. It is deciding how much transaction authority institutions are willing to delegate to software.



KEY TAKEAWAYS

- Major payment networks and industry consortia have introduced or begun piloting agent-oriented payment capabilities. This is no longer solely an innovation-lab discussion.
- Institutional stakeholders have largely moved beyond asking whether agentic payments are possible. The live questions are identity, control limits and liability.
- The shift underway is from automating transaction execution to delegating bounded transaction authority.

02 Why tokenised money matters

Programmable, always-on settlement for machines

The overnight treasury scenario described earlier depends on more than just intelligent agents. It also depends on money that can move, settle and respond to instructions at machine speed. Most money today is already digital, yet it sits in account-based systems built around human patterns: business hours, batch processes, minimum sizes and manual review at the key points. Autonomous agents need something else. An agent may pay for compute every few seconds, settle a cross-border obligation outside banking hours, or run many small, conditional payments one after another. Actions like these tend to be high-frequency, low-value and machine-triggered, and traditional rails can carry them only with enough friction to make the economics or the timing unattractive.



Tokenised money, whether stablecoins, tokenised deposits or tokenised central bank money, changes the operating model. This potential is increasingly recognised by policymakers and market infrastructure operators.

In a [joint report](#) to the G20, the BIS and CPMI noted that tokenisation could enable new forms of programmability, integrated settlement and more efficient financial market infrastructure, while also introducing new governance and interoperability considerations.

In each case, value moves on infrastructure built for programmability, composability and near-instant settlement. Those properties matter because agents do more than transfer value. They coordinate workflows. Programmability lets conditions live inside the payment. Composability lets a payment trigger a further on-chain action without extra calls. Atomic settlement means dependent steps either complete together or not at all, which matters when an agent has to know one step is final before the next can proceed.

OKX was precise on this, pointing to programmability, composability and deterministic settlement finality, and noting that at machine timescales, where transactions may occur continuously and in rapid succession, latency can accumulate across workflows and materially affect overall performance.

Every respondent then raised the same caveat independently. Tokenised money does not remove friction. It relocates it to the boundary between tokenised and traditional systems. As more transactions, assets and business processes move into digital and on-chain environments, the efficiency benefits of tokenised money are likely to become more pronounced.

Friction at the boundaries becomes less significant when a larger share of the transaction lifecycle can remain within programmable payment environments. The moment value enters or leaves that environment, operational frictions such as limited operating hours, liquidity constraints and manual processes can reappear.

Transak is explicit about this at the fiat-to-token edge, and does not flatter the rails: “anyone who tells you these rails are already frictionless is not being straight with you.” Kasu sees the same in private credit, where lenders may fund in stablecoins while borrowers still operate in fiat, and treats that conversion as deliberate risk reduction rather than a flaw. Velo sees this across its ASEAN corridors, where tokenised money enables faster movement of value, but on-ramps, off-ramps, licensing requirements and local banking and payment partners continue to determine whether a payment reaches its final destination.

The more likely path is that tokenised money becomes an increasingly important settlement layer for machine-speed, cross-border and high-frequency flows, complementing rather than replacing existing forms of money and payment infrastructure.

Polyhedra adds an observation that often gets overlooked. Institutional users do not necessarily want every transaction visible on public infrastructure. A bank does not want competitors observing its settlement activity, and a treasury desk does not want sensitive positions or transaction intentions becoming visible on shared transaction infrastructure before settlement is completed. Privacy, in Polyhedra’s view, is not the opposite of compliance. What institutions increasingly want is confidentiality paired with selective disclosure, so that a transaction can be proven valid without exposing sensitive commercial information. For tokenised money to work at institutional scale, privacy and verifiability must coexist.

The significance of tokenised money is not that it makes payments digital, because payments are already digital. Its significance is that it turns money into software, allowing settlement, conditions and execution to operate at machine speed. For agentic finance, that distinction matters.



KEY TAKEAWAYS

- Tokenised money supports programmability, atomic settlement and composability, which suit continuous machine-speed payments better than traditional rails.
- Friction does not disappear. It relocates to the boundary where value crosses between tokenised and traditional money.
- Tokenised money is particularly well suited to machine-speed flows where programmability and continuous settlement matter.

03 The Stack: Six layers that have to work


Six capability layers, and how they work together

Agentic payments do not run on a single technology. They run on a stack, and six capability layers have to work together before autonomous payments reach institutional scale. This layered approach is consistent with emerging policy research.

The [IMF](#) has argued that agentic payment systems can be understood through separate layers covering intent and orchestration, authorisation and control, and settlement, highlighting the importance of maintaining governance and accountability between decision-making and execution.

OKX is a useful illustration of why the stack is turning multi-layered in practice. Its own effort spans OKX Pay, where verified users authorise agents through bounded spending mandates, an Agentic Wallet designed for agent-native asset holding and execution, and an Agent Payments Protocol that provides a machine-to-machine rail for developers. Those three pieces address different parts of the stack, user authorisation, agent-native execution and machine payment rails, and OKX sees them converging rather than running as separate tracks. The written submissions from participating companies map onto these layers and, more usefully, illustrate how different participants are approaching the same challenges from different parts of the stack.

Figure 2 – Six layers of the agentic payments ecosystem and where value is likely to accrue



Source: UOBVM, UOBIG, UOB research and synthesis of written submissions from participating respondents

Figure 2 illustrates why agentic payments are not a single-product opportunity. Institutional deployment requires multiple layers to mature simultaneously. A failure in settlement can stop execution, but failures in identity, permissioning or accountability can prevent adoption altogether. In practice, the readiness of any deployment is likely to be determined by its weakest layer rather than its strongest.

While all six layers are required for institutional-scale adoption, respondents generally expect differentiation to migrate upward from settlement and execution toward trust, identity, permissioning, orchestration and user control. As payment infrastructure becomes increasingly commoditised, value creation is likely to concentrate in the layers that govern autonomy rather than the layers that move value.

LAYER 1

Payment, clearing and settlement



This is where intent meets money, and where the most important distinction in the whole piece lives. Transak sits at the boundary between ordinary money and tokenised money, and it describes three converging approaches to keeping agent payments compliant: a public register of agents, inherited permission where a person is verified once and the agent then acts within limits handed down from that check, and point-of-payment enforcement that screens every transaction as it happens.

Transak's own model combines the second and third, and its critique of the first two is pointed. A registry tells you who an agent is but not whether a given payment should be allowed right now, and an old permission grant does not catch new risk if the agent is later compromised.

The lesson they share is that a settled payment proves value moved, not that it should have moved.

Polyhedra makes the same point from the other direction. Its confidential settlement primitive executes a valid, authorised spend flawlessly, but does not decide whether the spend should happen. As Polyhedra puts it, a system can settle correctly on top of a bad authorisation decision, which is "garbage in, confidentially settled garbage out."

Velo expects AI capabilities to become embedded within settlement infrastructure itself, helping determine the optimal routing, timing and execution path for transactions based on factors including cost, speed, liquidity and risk. As greater decision-making moves into the settlement layer, governance, accountability and auditability become more important, not less.

LAYER 2

Agent intelligence and orchestration



The consistent message here is that the market overrates raw intelligence and underrates orchestration. NetMind argues that intelligence alone is insufficient. An agent may be capable of making a good decision, but institutional deployment requires multiple agents, tools and systems to coordinate consistently and reliably.

Eigenteam observes that operational issues often arise not because an individual agent fails, but because different agents interpret the same information differently as it moves through a workflow. Small inconsistencies can accumulate across multiple steps and only become visible later during audits, investigations or dispute resolution.

As a result, orchestration, consistency and traceability are increasingly viewed as important as intelligence itself.

LAYER 3

Treasury, liquidity and credit

Treasury and liquidity workflows are relatively well suited to bounded autonomy because they are highly structured, rules-based and supported by established controls.

Credit presents a more difficult challenge because it involves judgment, risk allocation and the potential for financial loss. Kasu argues that institutions will require evidence across multiple economic cycles, clear accountability structures and a responsible licensed party before autonomous credit decisioning can be widely adopted.

For the foreseeable future, respondents generally expect a human-plus-agent model, where agents automate data collection, analysis and routine execution while humans remain accountable for lending decisions and credit risk outcomes.

LAYER 4

Trust, identity and authorisation

This may be the most important layer, as every other layer depends on knowing who an agent is, who authorised it and what it is permitted to do. Respondents approached this challenge from different perspectives.

OKX focuses on bounded mandates, where users define spend limits, permitted transaction types, asset classes and expiry conditions before authority is delegated to an agent.

Bluwhale argues that trust should extend beyond traditional one-time KYC checks towards continuous monitoring of agent behaviour. Under this approach, spending limits and escalation thresholds can be adjusted dynamically based on changing risk profiles, making trust an ongoing governance process rather than a point-in-time verification exercise.

NetMind emphasises the importance of portable agent identities that can be recognised across platforms, wallets and institutions.

Polyhedra approaches the challenge through cryptographic verification. Its zkML framework is designed to provide evidence that an AI agent followed approved rules, processes and authorised data inputs when performing a task. This allows institutions to verify that actions were carried out in accordance with established controls, even when the underlying data or model logic cannot be fully disclosed.

Together, these approaches highlight that trust in agentic finance extends beyond identity verification to encompass permissions, oversight, accountability and verifiable execution.

LAYER 5

Interoperability and liquidity



Agents will need to operate across multiple payment networks, digital asset platforms, wallets and currencies. Respondents therefore view interoperability as a critical but still immature layer of the ecosystem.

While technological solutions for connecting networks continue to improve, practical challenges remain around liquidity availability, regulatory requirements, operational resilience and settlement certainty.

Several respondents noted that cross-network transactions can introduce additional points of failure and complexity. As a result, interoperability is not simply a technical challenge but also a liquidity, governance and risk-management challenge.

LAYER 6

End-user agent ecosystems



This is the surface people touch, and here trust takes a different form. As agents act in the background and people watch fewer screens, knowing who is responsible matters more, not less.

Bluwhale makes permissioning the product through its Agent Store, where users allocate funds to specific agents and revoke access instantly, and it finds that people delegate more authority than expected when permissions are granular, transparent and easily reversible.

Eigenteam notes that escalation rules become critical as a result, since thresholds that are too tight destroy efficiency while thresholds that are too loose let risk accumulate unnoticed.

Velo argues that trust becomes more important as agents increasingly operate in the background. Users need visibility into permissions, approvals, dispute resolution and the institution responsible when money moves.



KEY TAKEAWAYS


- No single layer solves institutional adoption alone. The weakest layer caps what an institution can safely deploy.
- A settled payment proves value moved, not that it should have moved, so compliance and permissioning belong at the point of transaction.
- Trust and identity are the most foundational and least standardised layers in the stack.

04 What the respondents are seeing today

Earlier and more advanced than the narrative suggests

Read together, the eight submissions paint a market that is both earlier and more advanced than the public narrative suggests. Earlier, because liability frameworks, identity standards and interoperability models remain unsettled. More advanced, because respondents are already deploying real-world solutions and encountering governance, compliance and operational challenges in production rather than in laboratories.

Figure 3 – Relative deployment readiness across agentic finance use cases



Source: UOBVM, UOBIG, UOB research and synthesis of written submissions from participating respondents

Figure 3 highlights a recurring pattern across the submissions. Adoption tends to appear first where activities are highly structured, liability is clear and outcomes are easy to measure. As responsibility, discretion and loss-bearing increase, autonomy becomes harder to deploy. This is why machine payments are progressing faster than autonomous credit, despite both being technically feasible.

The strongest consensus is that machine-to-machine payments are emerging as the first practical application of agentic finance. Nearly every respondent identified payments for compute, AI inference, APIs, data services and other digital resources as the most immediate use case. These transactions are frequent, low-value, highly structured and generally carry lower operational and risk-management complexity than many retail financial activities, making them well suited to autonomous execution.

Enterprise workflows form the second area of agreement. Respondents consistently highlighted procurement, invoice processing, reconciliation, treasury operations and routine cross-border transfers as workflows where agents can deliver measurable efficiency gains. These activities are governed by clear rules, generate extensive audit trails and operate within established control frameworks, making them natural candidates for bounded autonomy.

A second point of consensus is that governance has overtaken intelligence as the industry's primary challenge. Several companies observed that advances in AI capabilities have progressed faster than expected, while control frameworks have lagged behind. The central questions are no longer whether an agent can make a payment, but who authorised the action, what boundaries applied, how compliance is enforced and who is responsible when outcomes diverge from expectations.

The submissions also converge on the role of tokenised money. Respondents generally view stablecoins, tokenised deposits and other forms of programmable money as important enabling infrastructure rather than a complete solution. Faster settlement, programmability and continuous availability can improve efficiency, but seamless interoperability across payment rails, liquidity pools and financial ecosystems remains critical for adoption at scale.

Credit decisioning emerged as the clearest example of where the industry remains cautious. While respondents expect automation to increase across underwriting, monitoring and portfolio management processes, they consistently distinguished between moving money and allocating risk. Institutions may become comfortable allowing agents to execute payments within defined mandates long before they delegate loss-bearing credit decisions to autonomous systems.

Finally, contributors broadly agreed that trust is becoming the industry's most important layer. Whether approached through identity verification, behavioural monitoring, cryptographic proofs, permissioning frameworks or auditability tools, the common objective is the same: ensuring that autonomous actions remain attributable, controllable and accountable.



KEY TAKEAWAYS

- Machine-to-machine payments are emerging as the leading near-term use case for agentic finance.
- Enterprise workflows such as treasury, reconciliation and procurement are the next likely area of adoption.
- Governance, accountability and authorisation have become more important than model intelligence.
- Tokenised money is enabling infrastructure, where seamless interoperability across payment and settlement ecosystems is critical to achieving greater efficiency.
- Trust, identity and auditability are becoming the industry's most important common priorities.

05 How respondents are tackling the challenges

Different approaches to governance, trust and control

The industry broadly agrees on the barriers to adoption. The more interesting question is how different participants are attempting to address them. While the submissions reveal broad agreement on where the market is heading, they also highlight different approaches to solving the challenges that still separate experimentation from institutional adoption. Across the ecosystem, innovation is increasingly concentrated not in payment execution itself, but in the layers that govern trust, identity, control and accountability.

Although the approaches differ, they share a common objective: making autonomous financial actions observable, controllable and accountable without removing their efficiency advantages.



Governance and control

The most consistent design principle across the submissions is that autonomy must remain bounded.

Rather than pursuing unconstrained financial agents, respondents are building systems that operate within clearly defined mandates, spending limits and escalation policies. OKX, for example, approaches governance through bounded permissions that allow users to define transaction limits, authorised asset types and expiry conditions before authority is delegated to an agent.

NetMind similarly focuses on non-custodial execution, where agents can coordinate actions but final control remains with the user.

Eigenteam approaches the challenge from an orchestration perspective. Its experience suggests that many operational failures do not originate from a single agent making an incorrect decision, but from multiple agents interpreting information differently across a workflow. Governance therefore extends beyond transaction approval to ensuring that actions remain traceable, explainable and reproducible across an entire chain of decisions.

Taken together, these approaches suggest that the future of agentic payments will be governed less by what agents can do and more by the constraints within which they operate.



Trust, identity and authorisation

If governance determines what an agent may do, trust determines whether institutions are willing to let it act at all.

Several respondents identified identity and authorisation as the least mature layer of the stack. Bluwhale argues that traditional identity verification was designed for human users conducting occasional transactions, not software acting continuously on behalf of users. Its approach combines KYC with ongoing behavioural assessment, allowing trust levels and permissions to evolve as activity patterns change over time.

Polyhedra tackles the same problem through cryptographic verification. Its zero-knowledge proof infrastructure is designed to provide evidence that an authorised process was followed without requiring participants to disclose sensitive underlying information. In this model, institutions gain stronger assurance that actions were executed correctly while preserving confidentiality.

NetMind highlights another emerging requirement: portability. For agentic payments to scale across institutions, wallets and applications, agent identities and permissions will eventually need to move beyond proprietary ecosystems. An agent trusted in one environment may ultimately need to present verifiable credentials in another.

Despite the differing approaches, all seek to answer the same institutional question: who authorised the action, what permissions are applied, and how can it be independently verified afterward?



Interoperability and settlement

Many of the remaining challenges arise not within individual systems, but at the boundaries between them.

Transak focuses on the connection points between fiat and tokenised money, where identity verification, regulatory compliance and banking infrastructure continue to introduce operational friction. While programmable money improves settlement efficiency, the movement of value into and out of regulated financial systems still requires trusted intermediaries and compliance processes.

Velo addresses interoperability through cross-border settlement networks designed to move value across jurisdictions. Its experience shows that speed alone does not ensure usability. Liquidity availability, local market access, regulatory acceptance and strong compliance capabilities remain equally important considerations. Velo argues that regulatory compliance can serve as an important trust enabler for users, partners, financial institutions and regulators.

Polyhedra approaches interoperability through cryptographic infrastructure intended to support secure communication and value transfer across chains. Even so, interoperability remains one of the least mature areas of the ecosystem, with liquidity fragmentation, differing standards and cross-network coordination continuing to create complexity.

The emerging consensus is that moving value within a network is increasingly straightforward; moving value reliably across networks remains the greater challenge.



Treasury, credit and risk management

The clearest distinction between today's opportunities and tomorrow's ambitions appears in the area of credit.

Most respondents expressed confidence that agents will increasingly automate treasury operations, liquidity management and payment execution within predefined mandates. These activities benefit from clear objectives, measurable outcomes and well-established control frameworks.

Credit is the hardest layer, and Kasu argues that it demands a different standard of proof. Kasu argues that any meaningful move toward autonomous lending requires more than model performance. It requires demonstrated results through complete economic cycles, clearly assigned accountability and institutions willingness to bear losses when decisions prove wrong. As Kasu puts it, "Institutions underwrite scars, not backtests. Follow the liability and the losses, not the autonomy."

For this reason, many respondents expect a human-plus-agent model to dominate for the foreseeable future. Agents may automate data gathering, monitoring, analysis and execution, while responsibility for material underwriting and capital allocation decisions remains with accountable institutions and individuals.

The implication is straightforward: autonomous capital movement is likely to arrive well before autonomous capital allocation.



KEY TAKEAWAYS

- Respondents are focusing on governance, identity and control rather than unconstrained autonomy.
- Different companies are approaching trust through bounded permissions, behavioural monitoring, cryptographic verification and portable credentials.
- Interoperability remains a challenge spanning liquidity, regulation, compliance and settlement infrastructure.
- Treasury and operational workflows appear closer to adoption than autonomous credit decisioning.
- The greatest opportunities for innovation increasingly sit in the trust and governance layers that surround payments, rather than in payment execution itself.

06 Three barriers that remained unsolved

Liability, identity and interoperability

Despite rapid progress across the ecosystem, three barriers remain unresolved. Importantly, all three sit at the intersection of technology, governance and institutional accountability. They are not problems that better models or faster settlement alone can solve.

As Figure 4 shows, the industry's constraints increasingly sit in the trust layers of the stack rather than the execution layers.

Figure 4 – What each layer does, how ready it is, and the key risks that remain

01

Liability



Role - Determines accountability when autonomous actions result in losses or disputes

Current state - No consistent industry framework

Areas to improve - Regulatory, contractual and legal clarity

02

Identity & authorisation



Role - Determines who an agent is and what it is allowed to do

Current state - Fragmented and platform-specific

Areas to improve - Portable identity, standardised permissions, continuous authorisation

03

Interoperability



Role - Allows agents to operate across networks, institutions and jurisdictions

Current state - Liquidity and compliance fragmentation remain

Areas to improve - Common standards, trusted interoperability frameworks and deeper liquidity

Source: UOBVM, UOBIG, UOB research and synthesis of written submissions from participating respondents

Liability



The most fundamental unanswered question is liability.

When an autonomous agent initiates a payment that causes financial loss, responsibility is not always clear. Liability could potentially sit with the user who delegated authority, the institution that approved the mandate, the software provider that designed the agent, the wallet or payment provider that executed the transaction, or some combination of the above.

The challenge becomes more significant as agents gain greater discretion. An institution may be comfortable accepting responsibility for a payment executed within a narrowly defined mandate. It is far less clear who bears responsibility when an agent interprets instructions incorrectly, operates on incomplete information or behaves unexpectedly within otherwise authorised boundaries.

Technology can help establish accountability through signed mandates, immutable audit logs and verifiable execution records. These controls can make liability traceable and assignable. They cannot determine how liability should ultimately be allocated. That remains a legal, contractual and regulatory question that the industry has yet to answer consistently.

As several respondents observed, autonomous payments can only scale once participants know not merely what happened, but who is responsible when something goes wrong.

Identity & authorisation



The second unresolved challenge is agent identity.

Financial institutions have decades of experience verifying human customers and legal entities. They do not yet have a broadly accepted framework for verifying autonomous software acting across multiple institutions, platforms and jurisdictions.

Today, most agent identities remain platform-specific. Permissions granted in one ecosystem are rarely recognised in another. This limits interoperability and increases operational complexity as agents move between applications, wallets and financial institutions.

The challenge extends beyond identification. Institutions also need confidence that permissions remain valid over time. An agent that was authorised yesterday may behave differently tomorrow because circumstances, data inputs or risk conditions have changed.

The industry is gradually moving toward combinations of identity credentials, bounded mandates, behavioural monitoring and continuous authorisation. However, no common standard has yet emerged. Until portable identity and permissioning frameworks mature, many deployments are likely to remain confined to controlled environments and closed ecosystems.

Interoperability



The third challenge is interoperability.

While much attention focuses on blockchain bridges and technical connectivity, respondents consistently described interoperability as a broader problem encompassing liquidity, compliance, settlement certainty, FX conversion and regulatory alignment.

Agents may operate seamlessly within individual systems, but institutional workflows rarely remain inside a single network. Value must move across different payment rails, jurisdictions, asset classes, compliance frameworks and liquidity pools. Each transition introduces potential friction, delays or uncertainty.

Liquidity fragmentation remains a particular concern. An agent may identify an optimal route, but the route is only useful if counterparties, assets and liquidity are available at the moment execution is required. Similarly, settlement finality may appear straightforward to a human operator but can be considerably more difficult for an autonomous system to evaluate across multiple networks.

As a result, many of the risks in agentic payments do not arise within individual layers of the stack. They emerge at the boundaries between them.

The common thread - What links all three challenges is that they are fundamentally trust problems rather than technology problems.

The industry already knows how to move money autonomously. The harder challenge is determining who may authorise those movements, how trust is established across institutional boundaries, and who remains accountable when autonomous systems interact with the real economy.

These challenges do not invalidate the agentic payments thesis. Rather, they help define the priorities for the next phase of industry development. Across the ecosystem, participants are converging around several common approaches, including bounded mandates, stronger identity frameworks, continuous oversight, cryptographic verification and greater interoperability across payment and settlement networks. While no single solution addresses every challenge, progress across these areas is gradually laying the foundations for trusted autonomous finance at institutional scale.



KEY TAKEAWAYS

- Liability remains unresolved because technology can identify actions but cannot determine who ultimately bears responsibility for losses.
- No widely accepted framework yet exists for portable agent identity and cross-platform authorisation.
- Interoperability extends beyond technical connectivity into liquidity, compliance, regulation and settlement certainty.
- The greatest risks emerge at the boundaries between systems rather than within individual layers.
- The remaining barriers to adoption are primarily trust and governance challenges, not limitations in payment technology.


07 Looking ahead: Where will adoption first arrive

Where respondents expect adoption to arrive first

Looking ahead, respondents broadly expect adoption to follow a gradual and uneven path rather than a single industry-wide transformation. The next three years are likely to see bounded autonomous finance expand first in areas where outcomes are measurable, mandates can be tightly defined, and accountability remains clear.

Progress is unlikely to be determined solely by advances in AI. Instead, the pace of adoption will depend on how quickly the industry resolves questions around identity, control, liability and interoperability. Use cases that operate within established governance frameworks are therefore expected to move ahead more quickly than those involving significant discretion, judgement or loss-bearing responsibility.

Figure 5 – Machine-to-machine leads, enterprise follows, credit lags by years



Source: UOBVM, UOBIG, UOB research and synthesis of written submissions from participating respondents

Respondents generally expect machine-to-machine payments to be among the earliest large-scale deployments of agentic finance. Payments for computing resources, data services, APIs and other digital inputs are highly structured, repetitive and governed by clearly defined rules, making them well suited to autonomous execution. They also tend to involve lower risk and clearer transaction parameters than many retail financial activities. Over time, many respondents expect these machine-native payments to become embedded within digital platforms and applications, operating largely in the background rather than as a visible payment category.

Beyond machine-to-machine payments, enterprise operations are expected to form the next major wave of adoption. Treasury management, procurement, invoice processing, reconciliation and recurring cross-border payments already operate within established governance and control frameworks. As organisations become more comfortable delegating routine operational decisions to AI agents, these workflows are likely to provide a practical pathway for scaling bounded autonomy within financial institutions and large enterprises.

Cross-border payments also feature prominently in several forecasts. Recent industry initiatives have reinforced this direction. For example, [BIS Project Agorá](#) demonstrated the potential for programmable multi-currency settlement using tokenised commercial bank deposits and tokenised central bank reserves, highlighting how tokenisation could support more efficient cross-border payment workflows while preserving the safety and integrity of settlement systems. Respondents generally agree that programmable settlement, tokenised money and continuously available payment infrastructure will improve the economics and speed of cross-border transactions. At the same time, most expect progress to be constrained by factors outside the payment rail itself, including liquidity availability, regulatory requirements, local market access and fiat-to-token conversion. The opportunity is significant, but adoption will depend as much on ecosystem maturity as on technology.

Consumer-facing adoption is expected to emerge more gradually. While respondents generally agree that AI agents will increasingly assist consumers in managing routine financial activities, there is less consensus on the speed of adoption. Questions surrounding liability, dispute resolution and customer trust remain more challenging when financial decisions directly affect individuals. For this reason, respondents expect early consumer use cases to focus on tightly bounded activities such as budgeting, cash-flow management and predefined savings or investment mandates, where permissions are explicit and easily withdrawn.

Autonomous credit decisioning remains the most distant opportunity. Across the submissions, respondents consistently distinguished between executing financial decisions and making them. While AI is expected to play a growing role in underwriting, risk assessment and portfolio monitoring, institutions remain cautious about delegating loss-bearing decisions to autonomous systems. Most respondents expect human oversight to remain central wherever credit risk and regulatory accountability are involved, even as automation becomes increasingly sophisticated.

Over the longer term, respondents expect competitive advantage to migrate away from payment execution itself and towards the surrounding trust infrastructure. As payment rails, settlement capabilities and orchestration tools become more widely available, greater differentiation is likely to come from identity frameworks, permissioning models, auditability, interoperability and regulatory credibility. In other words, success may depend less on how autonomous an agent is and more on how confidently institutions can govern and trust it.



KEY TAKEAWAYS

- Machine-to-machine payments represent the strongest area of industry consensus for near-term adoption.
- Enterprise workflows such as treasury, procurement and reconciliation are likely to be the next major deployment wave.
- Consumer adoption is expected to emerge gradually through tightly bounded, low-risk financial activities.
- Autonomous credit decisioning remains a longer-term prospect due to liability, governance and loss-allocation considerations.
- Competitive advantage is shifting away from payment execution and toward trust, identity, permissioning and accountability.

08 Conclusion: Bounded autonomous finance

Trust applied to autonomous economic activity

Back to the earlier example of the Autonomous Treasurer, in a mature agentic payments environment, the agent does not simply move money while people sleep. It operates within a framework of mandates, permissions, controls and accountability established by the institution it serves. The payment layer settles continuously and compliantly. The intelligence layer coordinates workflows and escalates uncertainty. Treasury processes optimise liquidity within defined boundaries. Trust and identity frameworks determine what the agent is authorised to do, while interoperability infrastructure enables value to move across networks, jurisdictions and forms of money. Together, these layers make autonomous action possible without removing institutional oversight.

When all six layers work together, the treasurer does not wake to a mystery. The Autonomous Treasurer wakes to an optimised position, a complete audit trail and a clear record of what the agent did, why it was permitted and where accountability resides.

The most important finding from this research is that the industry is not building toward unconstrained autonomous finance. Across different business models, technologies and market segments, respondents consistently described a future built on bounded autonomy: explicit mandates, revocable permissions, continuous monitoring, auditable decision-making and clear accountability.

The technology required to initiate and execute payments autonomously is advancing rapidly. The harder challenge lies elsewhere. Liability frameworks remain unresolved. Agent identity standards remain fragmented. Interoperability across financial institutions, jurisdictions and settlement networks remains incomplete. These are not limitations of AI capability. They are institutional challenges that require coordination between technology providers, financial institutions, regulators and policymakers.

The future of agentic payments will therefore be determined less by how intelligent agents become and more by how effectively institutions govern them. The organisations that succeed will not necessarily be those that build the most autonomous systems. They will be those that make autonomy trustworthy, controllable and accountable at scale.

Agentic payments are not ultimately a story about artificial intelligence. They are a story about trust. As finance enters an era where software can increasingly act on behalf of people and institutions, the enduring challenge remains the same as it has always been: ensuring that authority, accountability and confidence move together with money.

Encouragingly, many of the building blocks needed to address these challenges are already emerging. Across the ecosystem, institutions are experimenting with bounded mandates, portable identity frameworks, continuous authorisation models, cryptographic verification techniques and trusted interoperability frameworks. Together with stronger auditability, human oversight and industry collaboration, these developments are helping lay the foundations for autonomous financial systems that can be trusted, governed and scaled responsibly. The future of agentic payments will not be defined by unrestricted autonomy, but by the industry's ability to combine technological innovation with robust governance, accountability and interoperability. If that balance can be achieved, bounded autonomous finance has the potential to become a foundational component of the future financial system.

Glossary

Agentic Payments	Payments initiated, executed or optimised by AI agents acting within predefined mandates and control frameworks.
Agentic AI	AI systems capable of planning, reasoning and taking actions autonomously to achieve specified objectives.
Bounded Autonomous Finance	A model in which AI agents operate autonomously within predefined mandates, permissions, controls and accountability frameworks.
Autonomous Agent	Software capable of making decisions and executing actions with limited human intervention within authorised boundaries.
Agent Identity	A verifiable digital identity that establishes who an AI agent is and links its actions to an authorised owner, organisation or mandate.
Delegated Authority	Permission granted by a user or organisation allowing an AI agent to perform specified actions on its behalf.
Permissioning	The rules, controls and authorisations that determine what actions an agent may perform.
Continuous Authorisation	A model in which permissions are continuously monitored and reassessed based on behaviour, risk conditions and context.
Tokenised Money	A digital representation of money issued or recorded on programmable infrastructure such as a distributed ledger.
Stablecoin	A digital asset designed to maintain a stable value relative to a reference asset such as a fiat currency.
Tokenised Deposit	A commercial bank deposit represented in tokenised form on a programmable infrastructure.
Programmability	The ability to embed rules, conditions or logic directly into financial assets or payment processes.
Composability	The ability for different applications, services or financial transactions to interact and build upon one another automatically.
Atomic Settlement	A transaction mechanism in which multiple related steps either all complete successfully or none of them complete.
Settlement Finality	The point at which a transaction becomes irrevocable and cannot be reversed.
Machine-to-Machine (M2M) Payments	Payments executed directly between software systems, applications, devices or AI agents without human intervention.
Digital Credentials	Verifiable digital records used to prove identity, permissions or authorisations.
Interoperability	The ability for different payment networks, platforms, blockchains and institutions to interact and exchange value seamlessly.
Liquidity Fragmentation	A situation where available liquidity is distributed across multiple platforms or networks, making efficient execution more difficult.
Non-Custodial	A model in which users retain direct control of their assets and cryptographic keys rather than relying on a third party.
Audit Trail	A verifiable record of actions, decisions and transactions that supports transparency, governance and accountability.
Zero-Knowledge Proof (ZKP)	A cryptographic technique that allows one party to prove something is true without revealing the underlying information.
zkML	The application of zero-knowledge proofs to machine learning systems, enabling verification that an AI model or process executed as intended without exposing sensitive data.
Orchestration	The coordination of multiple agents, systems, tools and workflows to achieve a desired outcome reliably and consistently.
Liability Framework	The legal, contractual and regulatory arrangement that determines responsibility when losses, errors or disputes occur.

References

Davidovic, S. and Tourpe, H. (2026). How Agentic AI Will Reshape Payments. IMF Note No. 2026/004. International Monetary Fund.

Bank for International Settlements (BIS) and Committee on Payments and Market Infrastructures (CPMI) (2024). Tokenisation in the Context of Money and Other Assets: Concepts and Implications for Central Banks. Report to the G20.

Bank for International Settlements (BIS) and Institute of International Finance (IIF) (2026). Project Agorá: A Shared Programmable Platform for Wholesale Cross-Border Payments.

Meet the team

Authors



Nikhil Chogle

SVP, Blockchain and Digital Assets
Nikhil.Chogle@uobgroup.com



Charles Kok

Director, UOB Venture Management
Charles.KokHC@uobgroup.com



Justin Wang

Director, Innovation Research Lab
Justin.Wang@uobgroup.com



Do Duc Tai

Data Scientist, Innovation Group
Do.DucTai@uobgroup.com

Editors



Heng Koon How, CAIA

Head of Markets Strategy
Heng.KoonHow@uobgroup.com



Tan Lena

Business Data Designer
Lena.Tan@uobgroup.com

Disclaimer

This publication is strictly for informational purposes only and shall not be transmitted, disclosed, copied or relied upon by any person for whatever purpose, and is also not intended for distribution to, or use by, any person in any country where such distribution or use would be contrary to its laws or regulations. This publication is not an offer, recommendation, solicitation or advice to buy or sell any investment product/securities/instruments. Nothing in this publication constitutes accounting, legal, regulatory, tax, financial or other advice. Please consult your own professional advisors about the suitability of any investment product/securities/ instruments for your investment objectives, financial situation and particular needs.

The information contained in this publication is based on certain assumptions and analysis of publicly available information and reflects prevailing conditions as of the date of the publication. Any opinions, projections and other forward-looking statements regarding future events or performance of, including but not limited to, countries, markets or companies are not necessarily indicative of, and may differ from actual events or results. The views expressed within this publication are solely those of the author's and are independent of the actual trading positions of United Overseas Bank Limited, its subsidiaries, affiliates, directors, officers and employees ("UOB Group"). Views expressed reflect the author's judgment as at the date of this publication and are subject to change.

UOB Group may have positions or other interests in, and may effect transactions in the securities/instruments mentioned in the publication. UOB Group may have also issued other reports, publications or documents expressing views which are different from those stated in this publication. Although every reasonable care has been taken to ensure the accuracy, completeness and objectivity of the information contained in this publication, UOB Group makes no representation or warranty, whether express or implied, as to its accuracy, completeness and objectivity and accept no responsibility or liability relating to any losses or damages howsoever suffered by any person arising from any reliance on the views expressed or information in this publication.



Right By You

Head Office
80 Raffles Place
UOB Plaza
Singapore 048624
Telephone: (65) 6533 9898
Facsimile: (65) 6534 2334

www.uobgroup.com