



Autonomous Agents in International Payment Flows

Ethan Williams
Asst Professor

Abstract

Cross-border payments are a core component of the global economy, relying on a complex network of banks and financial institutions to process transactions that range from a few seconds to several days via static APIs. Calls to modernize this system are widespread, yet current approaches offer only marginal improvement. Proposals built around digital identifiers and tokenization typically focus on changes to payment messaging formats or settlement mechanisms. Autonomous financial agents offer a different perspective on the problem, operating across institutional silos and handling transactions individually rather than in batches. In agent-based cross-border payment orchestration, individual agents act as digital representatives of the services they offer: service availability, terms, and costs are negotiated dynamically, and each transaction is managed end-to-end, from initiation through settlement. While this orchestration model offers meaningful advantages, fully realizing it raises concerns that merit closer examination.

Balancing the trade-off between minimizing overall latency and ensuring straightforward verification of compliance with regulations and service-level agreements requires a well-designed initiation-and-negotiation protocol. Because agents function as service gateways across the transaction lifecycle — from initiation to final settlement — certain steps normally handled through explicit exchange must instead be inferred by interacting agents. Future research should address the development of transaction-mechanism capacity suited to agent-based systems, ongoing measurement of latency and transaction volume within current architectures, and the real or perceived risks of misuse in agent-based orchestration. Resolving these issues would support broader adoption of autonomous agents in financial services and lay the groundwork for their eventual use in cross-border payments.

Keywords: Cross-Border Payments, Autonomous Financial Agents, Agent-Based Orchestration, Digital Identifiers, Tokenized Services, Transaction Negotiation Protocols, Latency Optimization, Regulatory Compliance Verification, Service-Level Agreement Enforcement, Multi-Institution Payment Networks, Transaction Lifecycle Management, Agent-Mediated Settlement, Dynamic Service Discovery, Interoperable Financial Silos, Payment Message Modernization, Compliance-Aware Orchestration, Decentralized Payment Coordination, Risk Of Agent Abuse, Real-Time Transaction Monitoring, Next-Generation Payment Architectures.

1. Introduction

Cross-border payments are financial transactions where money is moved from a payer's bank account in one country to another bank account located in a different country. Central banks, commercial banks, other licensed banks (such as money service businesses and other licensed PayPal-like payment platforms) are considered the primary stakeholders for cross-border payments. Payments can be either buy-side or sell-side; the intent will be to build a service that enables financial institutions to transmit their

customer requests to complete a payment without putting the funds at risk. Any payment demand can be executed when there is sufficient liquidity to complete the customer transaction. Currently, cross-border payments are built using static API implementation using the GET and POST protocols. The only transaction-related data is the request URL where specific transaction details are sent in the body. Consequently, messages between the wallet and the backend follow different schemas and access patterns.



A customer wallet sends a request to the backend, which in turn sends the data to the respective payment gateway. The requirement to pass structured data solely through the URL path or body does not allow additional enrichment for better governance and views. As a result, transaction status updates cannot be sent directly to the customer wallet. Additional costs may arise as the payee's bank must send separate SWIFT messages to the payee with information or manage an alternative interface. The centralized nature of APIs also brings with it the risk of data breaches and the associated costs.

1.1. Introduction to Cross-Border Payment Innovation

Cross-border payments by definition use two or more currencies, involve correspondents in different jurisdictions, engage multiple parties, and span national boundaries. These payments allow companies to purchase goods and services from international vendors, enable people to send money to family and friends living abroad, and allow travelers to pay for expenses incurred outside their home countries. Although cross-border payments comprise only around 2% of the global payment volume (as compared with 75% for domestic payments), they are heavily regulated and costly to process. Total cross-border payment volumes are nonetheless significant—around \$156 trillion per year.

Cross-border payment with transactions can involve more than two parties in the process, generating massive amounts of data and requiring orchestration to sequence and synchronize the activities and closely monitor progress. Operations can include monitoring transaction progress and being notified of deviations, rejections, and, in some cases, failures. In addition, the activity can be centralized by messaging or flow control systems such as SWIFT or by operational call centers but is mostly decentralized. Although cross-border payments with transactions are more costly, complex, and slower than domestic payment, efforts have been made to modernize payment gateways and payment messaging to improve cost and speed. A shift has also been proposed in payments, with the virtualization and digitalization of cash providing the basis for digital payment-enabled Central Bank Digital Currencies (CBDCs). However, operational and regulatory

considerations can complicate the introduction of these solutions, especially when integrating cross-border digital payment systems that span multiple jurisdictions.




Fig 1: Navigating the \$156 Trillion Friction: Orchestration Challenges and CBDC Integration in Global Cross-Border Payment Systems


2. Background: Static APIs and Current Cross-Border Payment Architectures

MSG-READ-ONLY sections are generated in-line in response to the user query and transformed from human-readable text into LaTeX markup style suitable for integration into the document.

Static APIs drive the data exchanges necessary to close currency and liquidity gaps present when payment systems denominated in two currencies interact. Current actor positions and the corresponding direction of the payment flits are extracted from the market. These data elements are exchanged through APIs using protocols such as MC-PROXY or SWIFT and follow a routing pattern similar to those used by READ commands. Static APIs provide the least change to current infrastructures, allowing gateway providers to add support for a new currency without internal development effort. However, strong dependencies on external actors exert frictions that increase costs and latency.

Architecture models that are mainly based on static API designs attempt to improve cross-border processes by applying specific modernization principles. However, these models are limited by the conservative nature of their approach; they enable only minimal changes to the existing

actor ecosystem and add only a single new capability: the use of static APIs. Hence, static API-based models do not mitigate well-known pains such as the burden of coordinate actors, maintaining oversight about payment flows, the risk of error during exception handling, the difficulty of monitoring compliance along the flow, the need to trust third-party vendors of a payment service, and the cost, latency, and risk exposure inherent in using intermediaries.



Equation 1 — “SLA single scoring function” (counterparty selection)

Step-by-step construction

Step 1 — define risk elements

Let the SLA/risk vector for counterparty j be:

$$\mathbf{r}_j = (r_{j,1}, r_{j,2}, \dots, r_{j,k})$$

Example components (consistent with the paper’s themes):

- $r_{j,1}$: expected delay risk
- $r_{j,2}$: exception-handling quality risk
- $r_{j,3}$: auditability / explainability risk
- $r_{j,4}$: regulatory exposure risk
- etc.

Step 2 — normalize each component

To combine heterogeneous metrics, normalize into $[0,1]$:

$$\tilde{r}_{j,i} = \frac{r_{j,i} - \min(r_i)}{\max(r_i) - \min(r_i)}$$

Step 3 — weight them

Let $w_i \geq 0$ and $\sum_i w_i = 1$. Then define the SLA risk score:

$$S_j = \sum_{i=1}^k w_i \tilde{r}_{j,i}$$

Step 4 — add monetary terms (fees + penalties)

Let fee quote be f_j . Let expected penalty be $P \cdot p_{nc,j}(v)$.

Define total expected cost:

$$C_j(v) = f_j + \lambda L_j(v) + P \cdot p_{nc,j}(v) + \mu S_j$$

where:

- $L_j(v)$ is predicted latency if choosing j
- λ converts minutes to “money-equivalent” (or utility weight)
- μ weights SLA non-monetary risk score

Step 5 — select counterparty

$$j^* = \operatorname{argmin}_j C_j(v)$$

2.1. Overview of Static API Limitations in Cross-Border Payments

Static APIs impose constraints that hinder the effectiveness of orchestrating cross-border payment transactions. A static API is defined as a system or service that only supports a specific, well-defined set of transactions in an unchanging manner. In contrast, a modern software architecture often supports business processes and transactions that are defined at runtime. Typically, these support multiple workflows and allow alternative paths of execution. In the case of cross-border payments, these include alternative messaging flows, gateway destinations, and settlement paths, among others.

The static API approach to cross-border payments relies on payment messages being directed from one payment gateway to another. The flow of cross-border payments can be likened to the flow of goods in the global supply chain: payments are moved from one financial institution to



another via payment gateways, with messages and funds being exchanged according to pre-established rules and standards. Given that payment gateways are part of the formal global financial system, they both report specific flows to the regulatory authorities and maintain an adequate level of compliance with anti-money laundering (AML), counter-terrorist financing (CFT), and sanction regulations. Despite the design of the cross-border gateway API and the national payment systems that support it, certain key variables need to be synchronized for the cross-border payment to be operational. These include, at a minimum, the following messages: the identity of the originators, the amounts and currencies, the beneficial owners, the purpose of the transaction, the originator and beneficiary banks, and all other operational conditions related to the transaction. Static APIs create bottlenecks when other variables need to be handled—those that are not usually part of such transaction but require flexibility (e.g., exception handling or time delays when a message is temporarily blocked). Because all interaction must comply with the rules established by the payment gateways, the absence of a dedicated agent with the capability of handling operational exceptions in a timely manner can generate significant latency.

3. Limitations of Traditional Orchestration in Global Transactions

The orchestration of cross-border payments is inherently complex, stochastic, and capillary. Hence, its implementation across the entire route remains challenging in terms of coordination and real-time status visibility. Certain events, such as unexpected delays, typically require manual intervention and reconciliation, increasing costs and the risk of errors. Organizational boundaries between intermediating parties in transaction routes, combined with the asymmetry of information flow and fulfillment, often pose additional hurdles for processing. Data operational considerations fundamentally limit the amount of information exposed to the whole transaction route. Sovereign regulations restrict the use of user data outside national borders, leading to data fragmentation.

Institutional aspects that require compliance with regulations within the organizational ecosystem that operates a payment route enforce a distinction between a transaction and its execution. Regardless of the route, it is still a violation of the terms that rule its execution. The necessity of collective compliance rules introduces an additional regulatory burden for all of the counterparties involved in a transaction, limiting globalization when transacting with countries with radically different regulatory policies. The need for different compliance obligations for counterparties in a transaction may even increase transaction costs and processing time, making them less attractive.

Moreover, the operation of static APIs creates lock-in conditions with a payment service provider, especially for users on the sell-side of the transaction. A payment service integrates a new payment route solely to enable transactions for its users, maintaining the status quo for all other transaction routes. As a consequence, the provisioning of that payment route is only beneficial for the users needing it, whereas the other users still face the same transaction costs and delays with static APIs. Hence, despite the availability of multiple payment routes, they are not fully leveraged.

3.1. Challenges in Traditional Global Transaction Orchestration

The architecture of global transaction orchestration is primarily determined by the orchestration of agents and the residue of the initiator and responders. All other components are required to either transparently facilitate operations with these orchestration agents or require (advanced) controls and conditions to primarily fulfil a limited set of control functions (e.g., anti-money laundering). Local and regional payment service providers have focused on rolling the dynamic transaction data into more coherent approximate construction of business objectives (e.g., through price indications), searches for desirable counterparties, or cumulative data channels that smooth regulatory checks but have not solved the fundamental problems in global coordination. Consider the traditional challenges and the role of the infrastructure supplier behind the communication messages



in cross-border payments. Cross-border orchestration has some of the hardest coordination problems: All participants in cross-border payments have acknowledged that, while it might not currently be needed in most, the economy breaks down once a payment between two individuals in different countries goes wrong, and that dealing with an exception might be complex, slow and expensive, and should avoid the need for other contemporaneous and future payments to also be intercepted. The foundational problems are transparency and audit behind the communication messages: the only topics where there is systemic (non-voluntary) risk are those related to settlement.




Fig 2: Systemic Settlement Risks and Agent Orchestration: A Framework for Transparency and Audit in Global Transaction Architectures

4. Emergence of Autonomous Financial Agents

Despite ambiguity in the literature, the terminology of the payments and fintech community is familiar and an autonomous agent is a construct that can be placed at the

intersection of those two domains. An autonomous financial agent collects funds from its owner and attempts to satisfy as many payment requests for others as possible, at low cost and with minimal risk. Such payment requests may cover a large number of use-cases and may be obligation-based (contractual) or demand-based, with non-fulfilment being either acceptable (non-critical) or associated with penalties (critical). An agent may act on the request of a single party (1:1) or of many parties (1:n). These requests can be local (intra-region) or global (cross-border). An agent can handle risk through negotiation and accept automatic off-loading to another agent when risk exposure exceeds a predefined level.

Transactions cannot be processed without affected parties knowing, and at least authorizing, any charge on ledgers under their control. To fulfill requests an agent must first obtain or negotiate the conditions of a financial contract with each counterparty. These negotiations must take into account what data will be exchanged during the transaction life-cycle between the counterparty agents, so as to detect inconsistencies. Above all, negotiated contracts must contain SLA (Service Level Agreement) information that connects various non-monetary risk elements into a single scoring function. This function is essential when agents select the best counterparty in order to minimize monetary and reputational loss. Since such selection must not be made in excess of risk analytical capabilities, any request being handled by more than one party creates a natural risk-sharing mechanism. When an agent's risk exposes it to possible loss it automatically off-loads the request – temporarily or permanently – by offering the service to another agent with uncommitted capacity or capability. Decision points emphasize the autonomous decision-making task of agents (e.g., risk-filtering, contracting, counterparty selection). Initiation and negotiation processes are described in detail, including aspects of the initial contract that specifically apply to initiation, together with risk filtering when demand exceeds supply, and counterparty discovery when agents need to find suitable partners.

Equation 2 — Latency vs verification trade-off (protocol objective)

- minimizing latency
- easy verification of compliance with regulations and SLAs

Step-by-step derivation

Step 1 — split latency

$$L(v) = L_0 + T_{\text{ver}}(v)$$

where L_0 is baseline orchestration time (routing, messaging, settlement mechanics), and $T_{\text{ver}}(v)$ is extra time for verification.

Step 2 — model verification time

A common smooth model is convex (strong checks cost more and more):

$$T_{\text{ver}}(v) = a + bv^2, \quad a, b \geq 0$$

Step 3 — model non-compliance probability

More verification reduces risk; a standard diminishing-returns model:

$$p_{\text{nc}}(v) = p_0 e^{-kv}, \quad p_0 \in (0,1), \quad k > 0$$

Step 4 — expected penalty

$$\mathbb{E}[\text{Penalty}(v)] = P p_0 e^{-kv}$$

Step 5 — define one optimization objective

Minimize “latency + risk cost”:

$$\begin{aligned} J(v) &= L(v) + \alpha \mathbb{E}[\text{Penalty}(v)] \\ &= (L_0 + a + bv^2) + \alpha P p_0 e^{-kv} \end{aligned}$$

Step 6 — find best v


Differentiate and set to zero:

$$\frac{dJ}{dv} = 2bv - \alpha P p_0 k e^{-kv} = 0$$

So the optimum satisfies:

$$2bv = \alpha P p_0 k e^{-kv}$$

Expected compliance/SLA penalty decreases with stronger verification (illustrative)



4.1. Definitions and Core Capabilities

An autonomous agent is defined as an entity that undertakes a task for another agent in a decentralized system according to a set of rules or a contract. Financial autonomy is the ability of an agent to make decisions about sources and uses of funds, whereas functional autonomy requires that an agent can independently undertake a conceptual step in the task, thus allowing for specialization. Both forms of autonomy can span multiple decisions over time. Autonomous agents can operate without central entities orchestrating simple tasks, yet coordination of more complex tasks remains necessary because of constraints such as hard-to-predict horizontal interoperability issues. Autonomous agents can use computational primitives such as contracts, negotiation, reconciliation, and dispute resolution—also called the building blocks of autonomy—at different degrees of independence. Some tasks, such as credit risk assessment, must be centralized owing to data aggregation constraints. Regional regulatory constraints may force data segregation, leading to multiple autonomous systems that are still separated in the vertical dimension. Each agent is responsible for its own compliance and operates in a sandbox whenever possible. In the event of a failure, every agent ensures that its governance rules are compatible with safe failure propagation to the entire cross-border system. Every executed contract is auditable, while explainable AI addresses explainability issues.




Stage	Static API latency (min)	Agent-based latency (min)	Improvement factor (Static/Agent)
Initiation	5	1	5.0
Negotiation	30	5	6.0
Execution	240	30	8.0
Confirmation	60	5	12.0
Settlement	720	60	12.0
Total	1055	101	10.4×

4.2. Governance, Compliance, and Risk Management

Autonomous financial agents must be capable of operating with appropriate levels of governance, compliance, and risk management. The decisions they take during transaction orchestration must be auditable; their actions should be explainable, with clear reason codes if they are rejected; and agents must operate in accordance with regulatory requirements, which may differ across jurisdictions. In systems with low latency or high volumes, custodial and other third-party responsibilities should be segregated into different agents, with custody of client funds residing with a separately registered payment institution or bank known to possess a license to operate at the required settlement layer.

These requirements can be addressed as follows. First, agents can be designed to log their actions and decisions in immutable form, with verification keys provided to a trusted third party, thus enabling third-party audit while maintaining the agents' inherent capacity for autonomous operation. Second, agents can be deployed in a way that third parties can probe them and seek explanations for rejected offers and other actions (up to a suitable limit). Third, agents' operations with clients can be constrained and enforced with the use of (non-positive) contracts. In the case of cross-border payments, this opens the pathway to enabling compliance by design. A Regulatory Compliance Agent could acquire the necessary licenses and oversee all parties in the payment operation—service providers, gateways, banks—ensuring that all regulatory requirements with respect to customer identification, money laundering

detection, risk taking, and costs are met. All contracts, including risk filters, could be registered and visible to this third party and used for automatic sanction or embargo checks. Custody management of client or institutional funds can be allocated to a separate Payment Custodian Agent, trusted by all the other agents and registered as a licensed financial institution.



5. Architectural Paradigms for Cross-Border Orchestration

Two architectural paradigms emerge when considering the orchestration of complex actions in cross-border transactions. The first is an abstract, multi-layered stack consisting of the following components: a central layer that coordinates and steers the various participants; orchestration engines that are specific to the function and control the processes of individual stakeholders; and terminal connections that set up the transaction with users, payment systems, and liquidity providers. The second is a service-oriented approach supported by a microservices architecture, where external developers can create new capabilities for the transaction processes and institutions can offer their services and assets as APIs.

The orchestration of autonomous financial agents and other automated decision-makers is best viewed as a central layer. Such systems are charged with acting in the best interest of constituents—initiating transactions, identifying

real-time pricing, selecting service providers, spotting market anomalies for settlement optimization, coordinating multiple participants to complete a transaction, managing exceptions, and more. The services provided to these agents are defined by the execution of specific contracts, with well-defined input information and exit criteria, and their addressability follows a request/reply pattern. Any established standards for the structuring of messages, especially message schemas, can be reused and supported, including harmonization of data in explicit message transformations. Other aspects that require definition and standardization are the creation of identities and cryptographic sets, the specification of security measures, and the management of data access rules. Many of these are already covered by the ISO 20022 Mapping for the Cross-Border Payments and Reporting plus Additional Requirements Message Set.




Fig 3: Autonomous Financial Orchestration: Multi-Layered Stacks and Microservice Architectures in ISO 20022-Compliant Cross-Border Systems

5.1. Agent-Based Orchestration Layer

The agent-based orchestration layer supports the management of autonomous agents and serves as their primary interface for interaction with the underlying data and service environment. It performs four main functions: (1) managing the lifecycle of individual agents, (2) mediating interactions between agents during transaction processing, (3) reconciling data from various sources, and (4) linking to legacy systems of record and control.

The orchestration layer does not define solution logic or provide specialized service functions; rather, it enables other layers (e.g., business process, choreography) and third parties to build value-added services. Like other layers, it must be designed with properties such as user experience, security, availability, and performance in mind. Fault tolerance is crucial, not least because agent-based approaches are still unproven and many agents might not be subject to constant or rigorous operational controls. Various levels of service and autonomy within autonomous-agent approaches will require programmability across a spectrum—from proactive control through explicit choreography and orchestration to true autonomy with agent-based approaches. This also implies potential agent interoperation with legacy systems, support for externally managed agents (e.g., for auditing), and appropriate levels of central coordination.

Verification strength v	Total latency (min)	Non-compliance probability	Expected penalty (\$)
0.0	45.0	0.2500	1250.00
0.1	45.6	0.1852	925.95
0.2	47.4	0.1372	685.82
0.3	50.4	0.1016	508.04
0.4	54.6	0.0753	376.63
0.5	60.0	0.0558	278.95
0.6	66.6	0.0413	206.58
0.7	74.4	0.0306	153.05
0.8	83.4	0.0227	113.36
0.9	93.6	0.0168	83.97
1.0	105.0	0.0124	62.25

5.2. Protocols and standards for Interoperability

A number of considerations are critical for agents design decisions: the properties required for agents' operations to be efficient and effective; the necessary terminology for implementation; and the protocols and standards guaranteeing agents' interoperability through both hardware and software that enable their joint operation.




Both the messages exchanged between the participating parties and the semantics of the messages are to be specified. Data needs to be harmonized according to the infrastructure on which the agents operate, addressing issues that are essential for message and algorithm execution. Infrastructure component policies (identity management and crypto certificate issuing, key management, access APIs, etc.) are considered for agents' software implementation and operation. Such policies are covered by common standards, which should preferably be re-used. The main focus here is on software protocol standards from the software layer perspective. Interoperability between agents warrants the joint performance of the multi-party automated execution of transaction life cycles, considering several aspects and addressing associated problems. Message schemas and bothersome issues are detailed here. Lack of compliance between parties generates delays (loss of efficiency) or leaving such problems unattended can expose participants to fraud. It is common and necessary to maintain such properties in any electronic transaction; therefore, it is necessary to guarantee that agents check such properties and services operate in compliance with regulations. The problem should be reduced to one party not complying after execution and thus submitting the system to risk. Auditable agents exposed the problem and provided explanation and accountability features. They comply with regulations.

6. Transaction Lifecycle with Autonomous Agents

The autonomous agent-based model of transaction orchestration provides a new lens for examining the transaction lifecycle with implications for all stages, from initiation through negotiation and execution to confirmation and settlement. Each stage is linked by relevant decision points that confirm, shape or alter the data requirement-gathering process and agent actions. Figure 10 illustrates the stages of the lifecycle and associated autonomous decisions.

In the initiation stage, an entity wishing to send a payment or other asset generates a contract together with any accompanying SLA and risk filters required by the nature of the asset, the underlying business purpose, any regulatory requirements, and the risk appetite of the initiator. The contract is then published to the matching service where agent entities acting in the capacity of senders can discover and evaluate the contract. Dashboards allowing manual confirmations are feasible at this stage for agents assisting humans.



Equation 3 : Core variables

- **Latency** (multi-second to multi-day in current systems)
- **Compliance verification** (AML/CFT/sanctions + contractual/SLA compliance)
- **SLA scoring** to choose counterparties and minimize monetary + reputational loss
- **Negotiation/contracting + risk filters + penalties**

We'll define:

Lifecycle latency

Let the stages be $s \in \{\text{init, neg, exec, conf, settle}\}$.
Let stage latency be $\ell_s \geq 0$.
Then total end-to-end latency:



$$L = \sum_s \ell_s$$

Compliance verification “strength”

We’ll model:

- Verification increases time: $T_{\text{ver}}(v)$
- Verification reduces probability of non-compliance: $p_{\text{nc}}(v)$

Expected penalty / loss

If non-compliance happens (regulatory breach, SLA violation, failed exception handling), let penalty be P (money + expected reputational cost).
Expected penalty:

$$\mathbb{E}[\text{Penalty}] = P \cdot p_{\text{nc}}(v)$$

6.1. Initiation and Negotiation

Cross-border transactions are instantiated and negotiated within the agent-based orchestration layer. An initiating agent submits a transaction request that specifies the originator and beneficiary of the transaction, the relevant amount in the originator's currency, and the expected settlement timing. All Data fields described in Section 2.2 need to be filled in for the transaction to be processed, even if only implicitly. The agent starts the initiation process by choosing the messaging schema for the transaction request based on the originator's and beneficiary's own Data schemas and thereafter, finding an interoperable Data schema with sufficient coverage supported by the counterparty's Data Access Layer. All Data fields can then be filled with the help of the Data Access Layer applied to the originator's identity and its Data Access Layer. Typically, the request will also include a Service Level agreement (SLA) with performance bounds and penalties. The message is propagated downward until a complete transaction description is obtained and the negotiation of contract terms is triggered. The initiation and negotiation stages include several key decision points. Risk filters may be applied on the service contracts proposed by the counterparty agents or on the set of initial counterparts

available to the initiating agent. A party's potential failure in fulfilling the agreement may be accounted for by proper penalty terms at the outset. Counterparty discovery may be based on reputation and trust systems, possibly involving closed or federated trust domains, or ad-hoc. The process captures some of the complexity required to support real-time inter-institutional contracts at the time of business deal-making without having to accommodate for every transaction class and its special features beforehand. Once the transaction contract, possibly integrating multiple service contracts, is formed, it is sent back towards the initiating agent side for validation. The actual embedding contract is next negotiated, ensuring that every contract element—also including timing terms and motives—are verified and validated. If all is considered correct, the transaction leaves the initiation–negotiation phase for the execution stage in the transaction life cycle.

7. Conclusion

The proposed agent-based orchestration layer offers a new approach to automating complex workflows for cross-border transactions among decentralized networks. The benefits of cross-border transaction orchestration using agent-based technology outweigh the associated risks. Automated regulatory compliance, risk assessment, and auditing functions, when plant in agents of different types, relieve individual parties from extensive regulatory burdens and reduce the overhead of auditing access to identity photos and other sensitive information. Nevertheless, research is required to ensure that the increased agent market does not become a source of risk concentration in

the global system.




Fig 4: Distribution of Systemic Outcomes

Future research should also develop technical and operational specifications for payment systems and explore pathways toward governance and regulatory frameworks that establish credibility and trust in agents supporting cross-border transactions among parties lacking a preexisting relationship. For example, a decentralized public-key infrastructure could provide a credible trust anchor for a public-key (crypto)identity scheme. The net outcome of successful research and implementation in these areas would be a set of agents incorporating official KYC/AML lists and other commonly controlled information that banks could use to augment the semi-autonomous capabilities of their payment systems without exposing them to the risk and expense associated with a general proposal for the public adoption of agent technology.

7.1. Future Directions and Implications for Cross-Border Payment Systems

Future Directions and Implications for Cross-Border Payment Systems

The findings suggest that agent-based orchestration provides new opportunities for scaling cross-border payments, but also creates risks that must be addressed through regulation, oversight, and continued evolution of

the payments ecosystem. These benefits and risks merit exploration as implementation becomes possible. Having the flexibility to specify transaction management characteristics such as conditions, service level agreements, risk appetite, and other factors opens up the possibilities for cross-border transactions with different characteristics. Obviously, embracing flexibility comes with the challenge of associating governance and controls that provide for a desired assurance level or risk profile.

The level of support for cross-border payment service standardisation will be an important consideration for payments and fintech architects. Potentially agent-based orchestration could evolve into an approach to measurements of performance, security, and resilience that go beyond traditional Service Level Agreement specifications for transactions and better assess the integration of processes along the transaction lifecycle. Industry initiatives that have been working on defining the structure and reducing the complexity of cross-border transactions could also have a role to play. It remains to be seen whether other stakeholders will embrace agents as a way forward and whether the standardisation of agents and their governing processes will eventually become a regulatory requirement to protect against the risks of a vendor lock-in and a lack of resilience.

8. References

1. Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., & Cao, Y. (2023). ReAct: Synergizing reasoning and acting in language models. International Conference on Learning Representations.
2. Schick, T., Dwivedi-Yu, J., Dessi, R., Raileanu, R., Lomeli, M., Hambro, E., Zettlemoyer, L., Cancedda, N., & Scialom, T. (2023). Toolformer: Language models can teach themselves to use tools. Advances in Neural Information Processing Systems, 36.
3. Liu, X., Yu, H., Zhang, H., Xu, Y., Lei, X., Lai, H., Gu, Y., Ding, H., Men, K., Yang, K., Zhang, S., Deng, X., Zeng, A., Du, Z., Zhang, C., Shen, S., Zhang, T., Su, Y., Sun, H., Huang, M., Dong, Y., & Tang, J.



- (2024). AgentBench: Evaluating LLMs as agents. International Conference on Learning Representations.
4. MANGALAMPALLI, B. M., KOLLA, S. H., APPA RAO NAGUBANDI, D. R., & SEGIREDDY, A. R. (2025). AN INTELLIGENT, REAL-TIME DIGITAL FABRIC FOR HEALTHCARE AND FINANCIAL ECOSYSTEMS USING AUTONOMOUS LEARNING AND GENERATIVE SYSTEMS. *TPM—Testing, Psychometrics, Methodology in Applied Psychology*, 32(S9 (2025): Posted 15 December), 3070-3086.
 5. Wu, Q., Bansal, G., Zhang, J., Wu, Y., Li, B., Zhu, E., Jiang, L., Zhang, X., Zhang, S., Liu, J., Awadallah, A. H., White, R. W., Burger, D., & Wang, C. (2024). AutoGen: Enabling next-gen LLM applications via multi-agent conversation. *Proceedings of the 2024 International Conference on Learning Representations Workshop on Large Language Model Agents*.
 6. Hong, S., Zhuge, M., Chen, J., Zheng, X., Cheng, Y., Wang, J., Zhang, C., Wang, Z., Yau, S. K. S., Lin, Z., Zhou, L., Ran, C., Xiao, L., Wu, C., & Schmidhuber, J. (2024). MetaGPT: Meta programming for a multi-agent collaborative framework. *International Conference on Learning Representations*.
 7. Wang, S., Ding, H., & Chen, H. (2023). Large language models in finance: A survey. *Proceedings of the Fourth ACM International Conference on AI in Finance*, 374–382.
 8. Wu, S., Irsoy, O., Lu, S., Dabrovolski, V., Dredze, M., Gehrmann, S., Kambadur, P., Rosenberg, D., & Mann, G. (2023). BloombergGPT: A large language model for finance. *arXiv preprint arXiv:2303.17564*.
 9. Mangalampalli, B. M., Bandi, V. D. V. K., Kolla, S. K., & Kumar, M. V. K. (2025). Towards Self-Evolving Healthcare Intelligence: Integrating Advanced Learning Systems with Real-Time Clinical Data Pipelines. *Cultura: International Journal of Philosophy of Culture and Axiology*, 22(12s), 464-486.
 10. Dong, Y., Wu, F., Zhang, K., Dai, Y., Zhang, S., Ye, W., Chen, S., & Cheng, Z.-Q. (2025). Large language model agents in finance: A survey bridging research, practice, and real-world deployment. *Findings of the Association for Computational Linguistics: EMNLP 2025*, 17889–17907.
 11. Aldasoro, I., & Desai, A. (2025). AI agents for cash management in payment systems. *Bank for International Settlements Working Paper No. 1310*.
 12. Davidovic, S., & Tourpe, H. (2025). How agentic AI will reshape payments. *IMF Notes*, 2025(001).
 13. Adrian, T., Garratt, R., He, D., & Mancini-Griffoli, T. (2023). Trust bridges and money flows. *Bank for International Settlements Working Paper No. 1112*.
 14. Arner, D. W., Buckley, R. P., Lammer, T., & Zetsche, D. A. (2022). Building regional payment areas: The single rule book approach. *Bank for International Settlements Working Paper No. 1016*.
 15. Anker-Sørensen, L., Passador, M. L., Wehrli, A., & Zetsche, D. A. (2022). DLT-based enhancement of cross-border payment efficiency: A legal and regulatory perspective. *Bank for International Settlements Working Paper No. 1015*.
 16. Auer, R., Haene, P., & Holden, H. (2021). Multi-CBDC arrangements and the future of cross-border payments. *Bank for International Settlements Papers No. 115*.
 17. Mangalampalli, B. M., Kolla, S. K., Bandi, V. D. V. K., Yandamuri, U. S., & Rani, P. S. (2025). Designing Intelligent Healthcare Ecosystems through Adaptive Data Integration and Autonomous Learning Systems. *Vascular and Endovascular Review*, 8(20s), 330-347.
 18. Bech, M. L., Faruqui, U., & Shirakami, T. (2020). Payments without borders. *Bank for International Settlements*.
 19. Lammer, T., Rees, D., Rice, T., & Shirakami, T. (2025). Enhancing cross-border payments: State of play and way forward. *Bank for International Settlements Bulletin No. 119*.
 20. Garratt, R., Koo Wilkens, P., & Shin, H. S. (2024). Next generation correspondent banking. *Bank for International Settlements Bulletin No. 87*.
 21. Van Hee, K., Kosse, A., Wierds, P., & Wijngaard, J. (2025). Let's speak the same language: A formally defined model to describe and compare payment system architectures. *Bank for International Settlements Working Paper No. 1259*.



22. Kabadjova, B. A., Badev, A., Bastos, S. B., Benos, E., Cepeda-López, F., Chapman, J. T. E., Diehl, M., Duca-Radu, I., Garratt, R., Heijmans, R., Kosse, A., Martin, A., Nellen, T., Nilsson, T., Paulick, J., Pustelnikov, A., Rubem do Coutto Bastos, M., & Testi, S. (2023). Intraday liquidity around the world. Bank for International Settlements Working Paper No. 1089.
23. Di Iorio, A., Fitzgerald, E., Lammer, T., & Rice, T. (2024). Regional payment infrastructure integration: Insights for interlinking fast payment systems. Bank for International Settlements.
24. Fitzgerald, E., Illes, A., & Lammer, T. (2024). Steady as we go: Results of the 2023 CPMI cross-border payments monitoring survey. Bank for International Settlements.
25. Kim, S. S., & Yang, S. (2024). Accuracy improvement in financial sanction screening: Is natural language processing the solution? *Frontiers in Artificial Intelligence*, 7, 1374323.
26. Motie, S., & Raahemi, B. (2024). Financial fraud detection using graph neural networks: A systematic review. *Expert Systems with Applications*, 240, 122156.
27. Mangala, N., & Kolla, S. H. (2025). Real-Time Feature Engineering for Streaming AI Workloads Using PySpark and Azure Event Hubs. *International Journal of Multidisciplinary Research in Science, Engineering, Technology & Management*, 1(6), 93-105.
28. Shih, Y.-C., Dai, T.-S., Chen, Y.-P., Ti, Y.-W., Wang, W.-H., & Kuo, Y. (2025). Fund transfer fraud detection: Analyzing irregular transactions and customer relationships with self-attention and graph neural networks. *Expert Systems with Applications*, 259, 125211.
29. Bin Sulaiman, R., Schetinin, V., & Sant, P. (2022). Review of machine learning approach on credit card fraud detection. *Human-Centric Intelligent Systems*, 2, 55–68.
30. Ileberi, E., Sun, Y., & Wang, Z. (2022). A machine learning based credit card fraud detection using the GA algorithm for feature selection. *Journal of Big Data*, 9, Article 24.
31. Benchaji, I., Douzi, S., El Ouahidi, B., & Jaafari, J. (2021). Enhanced credit card fraud detection based on attention mechanism and LSTM deep model. *Journal of Big Data*, 8, Article 151.
32. Dang, T. K., & Tran, T. C. (2021). Machine learning based on resampling approaches and deep reinforcement learning for credit card fraud detection systems. *Applied Sciences*, 11(21), 10004.
33. Khelifa, M., & Gharbi, A. (2021). Identification of non-typical international transactions on bank cards of individuals using machine learning methods. *Procedia Computer Science*, 190, 178–183.
34. Lokanan, M. E. (2023). Predicting mobile money transaction fraud using machine learning algorithms. *Applied AI Letters*, 4(1), e85.