

ARCHITECTING MULTI-ENTITY LEDGER SYSTEMS FOR SCALABLE AND COMPLIANT GLOBAL PAYMENT PLATFORMS

Satheesh Kumar Kumara Chinnaian
Independent Researcher, USA

Abstract

Global payment platforms have grown into extraordinarily complex financial ecosystems, ones that touch dozens of legal entities, hundreds of currency pairs, and numerous regulatory perimeters, often within the lifecycle of a single transaction. This technical review examines how multi-entity ledger architectures can be designed to meet that complexity, with particular focus on customer liability management, payables and receivables tracking, revenue recognition, transaction cost monitoring, loss accounting, and cash management reconciliation. Beyond structural design, the review explores how embedded control frameworks, self-healing exception pipelines, and trend-based anomaly detection can meaningfully reduce operational overhead while improving financial accuracy. Practical diagnostic examples are included, including how a rising transaction cost ratio can signal that an external processor has silently risk-flagged a merchant's traffic due to missing critical data fields. Visual dashboards and architecture diagrams support these concepts throughout. The article uses peer-reviewed and practitioner literature from the fields of fintech, distributed systems, and financial governance

Keywords: Multi-Entity Ledger, Transaction Cost Anomaly, Self-Healing Finance, Revenue Recognition, Cash Reconciliation, Closed-Loop Validation, Payment Platform Intelligence, Event-Driven Architecture

1. Introduction

1. The Structural Complexity Of Financial Records In Global Payment Ecosystems

There is a particular kind of complexity that emerges not from any single difficult problem but from the collision of many individually manageable ones. Global payment platforms live precisely in that space. At their core, they are accounting systems, tracking who owes what to whom, in which currency, under which legal framework, at what moment in time. But the scale and multi-jurisdictional nature of modern platforms has stretched traditional accounting architectures well past their original design limits.

The fintech revolution, as documented in the broader literature [1], does not merely layer new technology onto existing financial processes, it fundamentally reconstitutes the architecture of financial services themselves. That reconstitution is nowhere more consequential than in the ledger. A payment platform serving merchants across twenty countries, processing thirty currencies, routing through six processor relationships, and operating under four distinct legal entity structures simply cannot function on a single-entity, centralized ledger. The accounting model has to match operational reality.

What makes multi-entity ledger design genuinely difficult is not any individual requirement, it is their simultaneity. A single card authorization might touch a customer liability account, a merchant payable, a revenue accrual, a processor cost accrual, and a currency conversion entry, across three legal entities, in under two hundred milliseconds. Each of those entries must be correct, balanced, traceable, and auditable. And that sequence must repeat, without variance, for hundreds of millions of transactions per day.

Research into multi-entity transaction analysis [2] demonstrates that conventional transaction graphs become analytically intractable at scale without structured entity disambiguation. The entity context of a financial record is not metadata, it is load-bearing. Strip it out or treat it as an afterthought, and the entire accounting structure becomes unreliable.

The implication for ledger architecture is direct. Financial records in global payment ecosystems must be designed from the ground up as multidimensional objects, carrying not just a debit, a credit, and an amount, but a full context envelope: legal entity, currency, settlement layer, regulatory classification,

and originating event reference. This model is the foundation on which everything else in this review is built.

The diagram below illustrates how a production multi-entity ledger organizes these dimensions, from the canonical event bus through to reconciliation, self-healing, and operational dashboards.

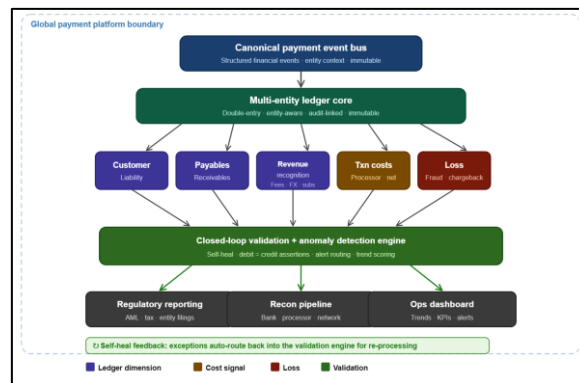


Figure 1. Multi-entity ledger system

2. Core Architectural Principles Of Multi-Entity Ledger Systems

2.1 Separating the Processing Layer from the Accounting Layer

One of the more consequential design decisions in building a payment ledger is the one made earliest and hardest to undo: whether to keep transaction processing and accounting in the same system. Platforms succeed in part because they decompose traditional financial monoliths into purpose-built components [3], and ledgers benefit from exactly that decomposition.

Processing systems are built for speed, availability, and customer responsiveness. Accounting systems are built for correctness, immutability, and auditability. These are genuinely different engineering objectives with varying consistency requirements. Conflating them routinely produces a system that satisfies neither adequately.

The canonical financial event model sits at the boundary between the two. When the processing system completes a transaction, authorization, capture, reversal, settlement, it emits a structured financial event. The ledger consumes that event and translates it into double-entry records, applying entity context, currency treatment, and regulatory classification. The processing system never writes directly to the ledger. The ledger does not influence processing decisions. The interface between them is a well-defined, version-controlled event contract.

2.2 Entity-Aware Financial Records

Compliance requirements apply to specific legal entities, not to the platform as a whole [4]. A platform may appear unified from the outside, but from a regulatory perspective, it is a collection of licensed entities, each with its own capital requirements, reporting obligations, and jurisdictional constraints.

Every financial record must therefore carry an entity identifier that is enforced by the ledger system itself, not inferred by downstream consumers. Cross-entity financial relationships, intercompany loans, intragroup settlements, and shared cost allocations, require explicit bilateral ledger entries. Entity-aware ledger structures are not an optional enhancement; they are the technical implementation of regulatory reality.

2.3 Immutability and Audit Traceability

Decentralized and distributed financial audit trail architectures [5] make a compelling case that traditional mutable ledger structures create systemic audit risk. When a correction applied without a paper trail is indistinguishable from an error or manipulation, the entire ledger loses its evidentiary value. The governing principle is straightforward: once an entry is posted, it cannot be deleted or overwritten. Corrections travel as documented reversal-and-reentry pairs, creating an unbroken chain of financial history that any authorized party can independently verify.

3. The Seven Financial Dimensions Of A Payment Ledger

A production-grade payment ledger models at least seven distinct financial dimensions simultaneously. They are not separate systems, they are different interpretive lenses applied to the same underlying event stream. The interplay between them is where most of the intriguing financial management happens.

3.1 Customer Liability

Customer funds held on the platform, wallet balances, prefunded accounts, and escrow positions, are liabilities on the platform's balance sheet. The customer is a creditor of the platform. Every deposit increases the liability; every disbursement reduces it; every pending or failed transaction sits in an uncertain liability state until it resolves.

Work on central bank digital currency design [6] highlights an analogous liability management challenge: large-scale digital payment platforms face structurally similar pressures to those banks face managing demand deposits. The implication is that platforms operating at meaningful scale need the same rigor around liability monitoring, reserve management, and liquidity stress testing that banks apply to deposit books, and their ledger architectures need to support that rigor natively.

3.2 Payables and Receivables

Merchant settlement obligations accumulate as accounts payable throughout each settlement cycle. The platform owes merchants the net proceeds of transactions, less applicable fees, reversals, and reserve holdbacks. Receivables arise from disputed fees, chargebacks in recovery, and deferred income.

Accurate aging for both is non-negotiable. Unaged payables overstate obligations. Unrecognized receivables understate assets. Either distorts the picture that treasury, finance leadership, and regulators use for decision-making, and in a platform processing billions of dollars of throughput, even small aging errors compound into material misstatements quickly.

3.3 Revenue Recognition

Revenue in a payment platform is multi-sourced and recognition-event-specific. Interchange income is typically recognized at settlement. FX spread income is recognized at conversion. Subscription revenue follows a time-based schedule. Getting the recognition triggers right, and enforcing them consistently across entities and jurisdictions, is one of the genuine accounting discipline challenges of running a platform at scale.

The foundational logic of double-entry accounting [7] applies here in its most classical form: the complexity of payment revenue recognition is not managed by building special-purpose accounting machinery for every revenue type. It is managed by defining clear recognition rules that map any financial event to the correct revenue category through a consistent and auditable ledger framework.

3.4 Transaction Cost and Loss

Transaction costs include processor fees, card network assessments, fraud-screening tooling costs, and dispute processing fees. The loss ledger captures fraud write-offs, unrecovered chargebacks, and operational errors resulting in financial shortfalls. The ratio of total transaction costs to revenue is one of the platform's most strategically sensitive metrics, explored in depth in Section 6 with worked diagnostic examples.

3.5 Operational Expenses

Beyond direct transaction costs, payment platforms carry a distinct layer of operational expenditure that the ledger must track and allocate with equal discipline. This dimension encompasses staff costs, technology infrastructure, software licensing, vendor contracts, and facility overhead, expenses that do not vary directly with transaction volume but are nonetheless essential to the platform's financial picture.

In a multi-entity structure, operational expenses introduce an additional challenge: shared costs incurred at the group level must be allocated across legal entities in a manner that is commercially reasonable, consistently applied, and defensible under transfer pricing rules. The ledger must support entity-level cost attribution, not merely aggregate expense reporting, so that each entity's standalone

profit and loss position accurately reflects its true cost of operation. Misallocated or unallocated operational expenses distort entity-level profitability, undermine regulatory capital calculations, and create audit exposure during intercompany reviews.

4. Cash Management Reconciliation And Insights

Reconciliation is the discipline of verifying that what the ledger records actually occurred, as confirmed by the independent records of every financial counterparty the platform touches. In a multi-entity platform, this means continuous comparison of internal ledger balances against bank statements, card network settlement files, processor daily reports, and partner institution records across every entity-currency pair in the platform's operating footprint.

Automated financial reconciliation systems [8] show that platforms that move from batch overnight reconciliation to continuous near-real-time pipelines achieve measurable improvements in error detection speed and operational efficiency. The architectural shift is from reconciliation as a downstream reporting activity to reconciliation as an embedded control, running alongside financial processing rather than after it.

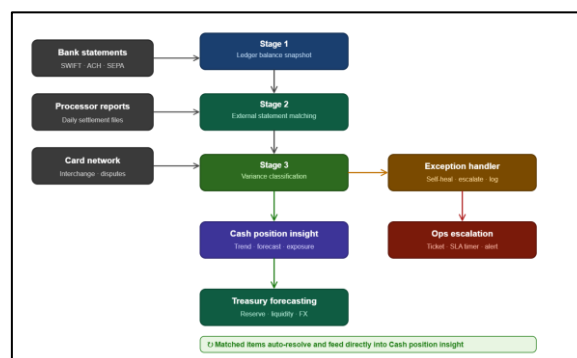


Figure 2. Three-stage cash management reconciliation pipeline

The pipeline above ingests external financial statements from banking partners, card networks, and processors; matches them against internal ledger snapshots; classifies variances; and routes exceptions, either to automated self-healing routines or to human escalation queues with SLA timers. Matched items feed a module for cash position insights that continuously updates treasury forecasting across all entity-currency pairs.

5. Self-Healing Architecture: Reducing Operational Overhead Through Embedded Controls

The term "self-healing" has a specific and important meaning in the ledger domain. A self-healing ledger is one that can detect financial exceptions, classify them by probable cause, and resolve them, either automatically through pre-approved rule execution or through structured escalation, without requiring ad hoc human investigation for every single instance.

Blockchain-enabled traceability research [9] demonstrates that combining automated anomaly detection with structured resolution workflows reduces operational overhead while simultaneously improving audit quality. The principle transfers directly to managing exceptions in payment ledgers.

5.1 Exception Classification: Three Tiers, Three Responses

Not all reconciliation breaks are equal, and treating them as if they were, routing every exception to a human analyst regardless of nature, is both wasteful and counterproductive.

Tier 1 exceptions are auto-resolvable. These include timing differences (a settlement posting one minute before midnight internally but received by the bank one minute after), known FX rounding differences within tolerance bands, and duplicates the platform has already processed under a different external reference. These are resolved automatically, documented in the audit trail, and closed without human involvement.

Tier 2 exceptions match a known pattern but require rule-assisted resolution. A processor fee arriving 3% above the contracted rate may fall within a pre-approved operational tolerance band. The system

posts an adjustment entry, documents the variance, closes the exception, and flags the cumulative pattern for periodic review. Tier 3 exceptions are novel, they do not match any known pattern, or they exceed materiality thresholds requiring human judgment. These are escalated with full context: the specific break, a suggested resolution approach, the relevant ledger history, and a countdown SLA timer. The analyst reviews a pre-analyzed situation rather than starting from scratch.

5.2 Embedded Controls as Real-Time Assurance

Platform governance at scale is fundamentally different from traditional corporate governance, the speed and volume of platform operations make after-the-fact oversight structurally insufficient [10]. The same logic applies directly to ledger controls: if a control runs only at the end of the day, it finds errors that have already propagated through twelve hours of downstream processing.

Embedding controls within the ledger pipeline, validating each financial event before it is committed, transforms the control function from a retrospective audit to a real-time assurance mechanism. A debit/credit balance assertion firing at event ingestion catches an imbalanced entry before it touches the ledger. An exposure limit check for the entity firing at commitment catches a potential breach before it affects reported balances.

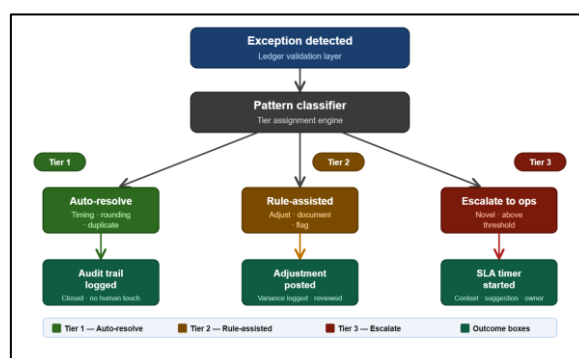


Figure 3. Self-healing exception classification and routing flow

6. Anomaly Detection: Reading The Signals In Financial Trends

A mature ledger platform does not just store financial data, it generates intelligence from it. Three trend anomaly classes are particularly informative for payment platform operators, and each carries a diagnostic vocabulary that experienced teams learn to read the way a clinician reads a vital signs chart.

The interactive dashboard below presents a representative twelve-month view of all three signals together: revenue trend, transaction cost ratio, and cash position. Note the October spike in the cost ratio, which breaches the 2.8% alerting threshold, this is the scenario unpacked in detail in Section 6.1.

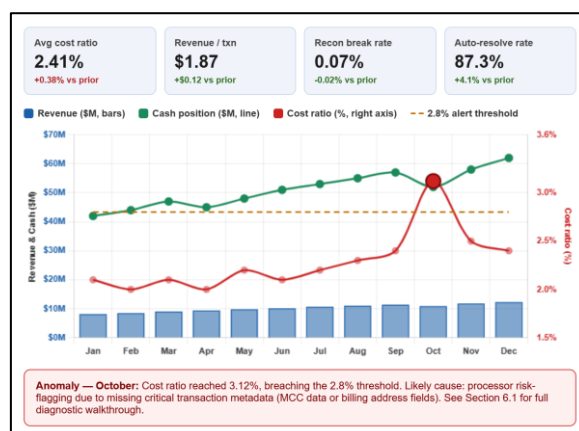


Figure 4. Twelve-month financial trend dashboard, revenue, cost ratio, and cash position

6.1 Transaction Cost Ratio: The Most Information-Dense Signal

Digital footprint analytics and fintech credit research [11] establish that behavioral signals embedded in transaction data carry far more diagnostic information than their surface appearance suggests. The same holds true for cost ratio trends. A rising transaction cost ratio is almost never simply a pricing event, it is a symptom, and the symptom can be mapped to a cause with reasonable reliability once the diagnostic pattern is understood.

The table below summarizes five specific cost anomaly patterns with their probable root causes and recommended investigation paths.

Cost anomaly pattern	Probable root cause	Investigation action
Sudden spike on one processor route only	Processor has risk-reclassified merchant traffic, may be applying a higher interchange tier silently	Pull fee breakdown from that processor; compare MCCs before and after spike onset
Gradual creep across all routes simultaneously	Authorization request data quality degrading, missing fields losing platform rate eligibility	Diff authorization payload completeness rate before and after creep began
Spike isolated to cross-border corridors	FX margin widening or correspondent bank fee schedule change	Check FX rate feeds and correspondent fee contracts
Spike correlated with rising chargeback volume	Elevated fraud driving dispute processing fees and network assessments	Review fraud scoring thresholds and 3DS enrollment rates
Spike following a platform integration change	New integration omitting fields required for lowest-rate interchange eligibility	Compare authorization request payloads before and after the change

Table 1. Transaction cost anomaly diagnostic framework with root cause mapping

The October spike visible in Figure 4, where the cost ratio reaches 3.12% against a 2.8% threshold, is a textbook illustration of the first pattern. In a real platform, this would trigger an automated investigation workflow: pulling the processor fee breakdown by MCC and corridor for that period, comparing against the contracted fee schedule, and cross-referencing with the platform's change log. If certain MCCs suddenly attract higher rates, the most likely explanation is that the processor's risk-scoring system flagged those transactions, often because critical data fields such as billing address verification data, device fingerprinting tokens, or merchant category descriptors are absent or malformed in the authorization request. The processor does not typically communicate this change; the cost ratio spike is the only visible signal.

6.2 Revenue Trend Anomalies

Transaction-level financial data [12] carries substantial predictive and diagnostic value that is only accessible through a properly structured revenue ledger with sufficient dimensional granularity. A revenue-per-transaction decline alongside stable or growing volume almost always reflects mix shift, more debit volume, more domestic corridors, and more lower-fee merchant categories. That is not necessarily a problem, but distinguishing a deliberate strategic mix shift from fee schedule misconfiguration or billing engine errors requires the revenue ledger to be filterable by transaction type and merchant cohort simultaneously.

A sudden revenue spike not accompanied by volume growth is a red flag. It may represent a data ingestion error, a double-billing event, an incorrect FX rate applied to revenue recognition, or a timing anomaly in accrual. Finding the cause quickly matters: an overstated revenue figure that enters a regulatory report creates a materially larger problem than one caught internally.

6.3 Cash Position Trends and Liquidity Intelligence

Real-time cash flow visibility [13] fundamentally changes the decisions that platform operators are able to make, not because the underlying financial position has changed, but because the information arrives early enough to act on. A GBP entity trending toward its reserve floor with four days of runway is a solvable liquidity management problem. One discovered on the morning of the breach is a crisis.

The platform's anomaly detector should alert treasury when the projected cash position for any entity-currency pair is forecast to breach its minimum reserve within a configurable forward window. The projection logic should incorporate scheduled settlement outflows, expected inflows from authorized transactions, and any known timing anomalies, such as a processor file delivery delay that reduces the effective settlement receipt window.

7. Operational Intelligence Dashboards

The value of a multi-entity ledger is not fully realized in the accuracy of stored balances, it is realized in the quality of the decisions those balances enable. Effective operational dashboards aggregate signals from all six ledger dimensions and present them in role-appropriate views: a cash position view for treasury, a revenue and cost margin view for finance leadership, an exception and self-healing status view for the reconciliation operations team, and a regulatory exposure view for compliance.




Figure 5. Live exception management and anomaly alert dashboard

The dashboard above reflects what a reconciliation team lead would see on a normal business day. The bulk of exceptions, timing differences, FX rounding, and duplicates, are resolved without human involvement. The anomaly alert panel surfaces the three trend signals from Section 6 alongside a routine operational note about a delay in processor file delivery that falls within a pre-approved tolerance band requiring no action.

8. Ensuring Financial Integrity Through Closed-Loop Validation And Reconciliation

Closed-loop validation confirms that every financial event produces a balanced and internally consistent set of ledger entries before it is committed. Reconciliation adds a second layer by verifying those committed entries against independent external counterparty records. Together, they establish the three pillars of financial data integrity that most regulatory frameworks require: completeness, accuracy, and timeliness.

Event-driven financial architectures built on stream processing infrastructure [14] can achieve the throughput required for global-scale payment platforms while preserving the ordering and consistency guarantees that financial ledgers demand. The enabling design choices are well understood: partition ledger data by entity and currency to distribute write load; separate write-path operations (requiring strong consistency) from read-path analytics (tolerating eventual consistency); and use event streams as the integration backbone between ledger processing, reconciliation, and reporting systems. Each subsystem can then scale independently without contending with the others.

9. Scalability Strategies For High-Volume Ledger Platforms

Digital payment ecosystems often process hundreds of millions of transactions per day. Ledger systems must scale horizontally while preserving strict transactional guarantees within each partition boundary. The consequence of the event-driven architecture described in Section 8 [14] is that the reconciliation pipeline processes at its own pace without blocking ledger writes, the reporting layer aggregates without affecting either, and the self-healing engine consumes the exception stream without contending with the main ledger processing path. What binds these subsystems together is not shared infrastructure, it is a shared, well-defined, and version-controlled event contract.

10. Regulatory Compliance And Audit Transparency In Multi-Entity Financial Platforms

The reconceptualization of financial regulation in the fintech era [15] has produced a regulatory environment that both shapes and is shaped by the platforms it governs. The implication for multi-entity ledger design is that compliance is not a fixed constraint satisfied at a point in time, it is a moving requirement that an architecture must be able to adapt to continuously.

A ledger storing financial records with full dimensional metadata, maintaining complete and immutable audit trails, enforcing entity-scoped access controls, and capable of filtering and aggregating across any combination of entity, currency, jurisdiction, and reporting category is not merely compliant with today's requirements. It is positioned to adapt to the needs of tomorrow. Specific regulatory obligations vary by jurisdiction and entity type, but major frameworks, including IFRS 9, Basel III, PSD2, and AML/KYC regimes, share common requirements: financial data must be traceable to originating transactions, corrections must be documented rather than deleted, historical financial states must be reconstructable, and reporting must be consistent across entities within a group.

A properly designed multi-entity ledger satisfies all of these requirements by construction, not as an overlay added after the fact, but as a natural consequence of its foundational architectural decisions. The integration of scalable processing, embedded controls, continuous reconciliation, and trend-based anomaly detection produces a platform that does more than satisfy regulatory requirements. It generates genuine insight into the financial health and risk posture of the business, enabling better decisions, earlier interventions, and more confident engagement with the regulators and auditors who depend on the accuracy of its records.

Conclusion

Global payment platforms have moved well beyond the point where a single-entity, centralized ledger can serve as a credible financial backbone. The operational realities of multi-jurisdictional operations, concurrent legal entity structures, and transaction volumes measured in the hundreds of millions per day demand an entirely different architectural philosophy, one where the ledger is not a passive record-keeper but an active, intelligent, and self-correcting financial system.

This review has traced that philosophy across its core dimensions. Customer liability, payables, receivables, revenue recognition, transaction costs, and loss accounting are not isolated accounting categories, they are interdependent lenses through which the same underlying financial events must be interpreted simultaneously and consistently. A ledger architecture that treats any one of these dimensions as secondary will inevitably produce blind spots that surface at the worst possible moments: during a regulatory examination, a liquidity stress event, or an operational incident that requires rapid financial triage.

The case for self-healing exception management is equally clear. Platforms that rely on overnight batch reconciliation and ad hoc human investigation for every break are carrying a structural operational risk that compounds with scale. Embedding controls within the processing pipeline, classifying exceptions at ingestion time, and routing them to pre-defined resolution workflows, auto-resolution for known patterns, rule-assisted adjustment for tolerance-band variances, and SLA-timed escalation for novel breaks, reduces both the cost and the latency of financial error detection in ways that batch approaches fundamentally cannot match.

Perhaps the most practically valuable contribution of this review is the diagnostic framework around trend-based anomaly detection. The transaction cost ratio, revenue-per-transaction trend, and cash position forecast are not merely reporting metrics; they are early warning signals with specific and interpretable diagnostic vocabularies. A sudden cost ratio spike on a single processor route, for instance, is rarely a pricing event. It is far more likely to indicate that the processor's risk-scoring system has silently reclassified a subset of the platform's traffic, often triggered by missing or malformed data fields in the authorization request that the platform had no direct visibility into. Recognizing that signal for what it is and building investigation workflows that systematically decode

it is the difference between reactive financial operations and genuinely intelligent financial management.

The broader implication is architectural. A ledger platform designed around canonical event models, entity-aware record structures, immutable audit trails, continuous reconciliation, and trend intelligence does not merely satisfy today's regulatory requirements. It creates the institutional infrastructure for sustained financial clarity, enabling treasury teams, finance leaders, operations managers, and compliance officers to make better decisions faster, grounded in a single authoritative source of financial truth that they can trust. As the regulatory landscape continues to evolve and transaction volumes continue to grow, that infrastructure becomes not a competitive advantage but a fundamental operating requirement for any platform that intends to operate with credibility at a global scale.

References

- [1] Peter Gomber, et al., "On the Fintech Revolution: Interpreting the Forces of Innovation, Disruption, and Transformation in Financial Services," *Journal of Management Information Systems*, 2018. Available: <https://www.tandfonline.com/doi/full/10.1080/07421222.2018.1440766?scroll=top&needAccess=true>
- [2] Yan Wu, et al., "A BTN-Based Method for Multi-Entity Bitcoin Transaction Analysis and Influence Assessment," *ACM Digital Library*, 2024. Available: <https://dl.acm.org/doi/epdf/10.1145/3686168>
- [3] Rainer Alt, et al., "FinTech and the transformation of the financial industry," *Electronic Markets*, 2018. Available: <https://link.springer.com/article/10.1007/s12525-018-0310-9>
- [4] Johnathan E Andrews, et al., "Regulatory Frameworks for Digital Payment Systems," *ResearchGate*, 2025. Available: https://www.researchgate.net/publication/400780432_Regulatory_Frameworks_for_Digital_Payment_Systems
- [5] Mark Graham, "Decentralized Financial Audit Trail Architecture for Enterprise Resource Planning Platforms," *ResearchGate*, 2026. Available: https://www.researchgate.net/publication/401795224_Decentralized_Financial_Audit_Trail_Architecture_for_Enterprise_Resource_Planning_Platforms
- [6] Itai Agur, et al., "Designing central bank digital currencies," *Journal of Monetary Economics*, 2022. Available: <https://www.sciencedirect.com/science/article/abs/pii/S0304393221000520>
- [7] Toshio HARA, "A Review of the Double-Entry Accounting System and Accrual Accounting in the Public Sector," *Government Auditing Review*, 2006. Available: <https://www.jbaudit.go.jp/english/exchange/pdf/e13d01.pdf>
- [8] Faysal Khan and Tahmina Akter Bhuya Mita, "Automated Financial Reconciliation Systems for Enhancing Efficiency and Transparency in Enterprise Accounting Workflows," *ResearchGate*, 2024. Available: https://www.researchgate.net/publication/400471145_Automated_Financial_Reconciliation_Systems_for_Enhancing_Efficiency_and_Transparency_in_Enterprise_Accounting_Workflows
- [9] Piera Centobelli, et al., "Blockchain technology for bridging trust, traceability and transparency in circular supply chain," *Information & Management*, 2022. Available: <https://www.sciencedirect.com/science/article/abs/pii/S0378720621000823>
- [10] Mark Fenwick, et al., "The End of 'Corporate' Governance: Hello 'Platform' Governance," *ResearchGate*, 2019. Available: https://www.researchgate.net/publication/331337702_The_End_of_'Corporate'_Governance_Hello_'Platform'_Governance
- [11] Tobias Berg, et al., "On the Rise of FinTechs: Credit Scoring Using Digital Footprints," *ResearchGate*, 2020. Available: https://www.researchgate.net/publication/345403114_On_the_Rise_of_FinTechs_Credit_Scoring_Using_Digital_Footprints
- [12] Anjan V. Thakor, "Fintech and banking: What do we know?," *Journal of Financial Intermediation*, 2021. Available: <https://www.sciencedirect.com/science/article/abs/pii/S104295731930049X>

- [13] Hua Song, et al., “Blockchain-enabled supply chain operations and financing: the perspective of expectancy theory,” *International Journal of Operations & Production Management*, 2023. Available: <https://www.emerald.com/ijopm/article-abstract/43/12/1943/148656/Blockchain-enabled-supply-chain-operations-and?redirectedFrom=fulltext>
- [14] Oyejide Timothy Odofin, et al., “Designing Event-Driven Architecture for Financial Systems Using Kafka, Camunda BPM, and Process Engines,” *International Journal of Scientific Research in Science, Engineering and Technology*, 2024. Available: <https://ijsrset.com/index.php/home/article/view/IJSRSET25121178/IJSRSET25121178>
- [15] Douglas W. Arner, et al., “FinTech, RegTech, and the Reconceptualization of Financial Regulation,” *Northwestern Journal of International Law & Business*, 2017. Available: <https://scholarlycommons.law.northwestern.edu/cgi/viewcontent.cgi?article=1817&context=njilb>