
| RESEARCH ARTICLE

Quantitative and Data-Driven Evaluation of Blockchain-Based Financial Systems: Transaction Efficiency, Transparency, Cost Optimization, and Performance Metrics in Global Markets

Yusuf Oli Rahat

University of New Haven, Business Analytics

Email: yraha1@unh.newhaven.edu

ORCID: <https://orcid.org/0009-0001-8834-7759>

Md Kamrul Islam

University of New Haven, Business Analytics

Email: misla22@unh.newhaven.edu

ORCID: <https://orcid.org/0009-0001-8906-630X>

Shah Farhan Rabbani

University of New Haven, Business Analytics

Email: srabb2@unh.newhaven.edu

ORCID: <https://orcid.org/0009-0001-8434-223X>

Corresponding Author: Yusuf Oli Rahat, **E-mail:** yraha1@unh.newhaven.edu

| ABSTRACT

Blockchain-based financial systems are increasingly evaluated not only as speculative infrastructures but as operational payment, settlement, and record-keeping networks that can be benchmarked against incumbent financial rails. This paper develops a quantitative, data-driven framework for assessing blockchain-based financial systems across four dimensions that matter in global markets: transaction efficiency, transparency, cost optimization, and overall performance resilience. Drawing on evidence from public blockchain networks, payment and remittance statistics, policy experiments, and institutional distributed-ledger pilots, the study synthesizes academic literature with world data from the World Bank, the Federal Reserve Bank of New York, the Bank for International Settlements, Visa, Bitcoin, Ethereum, and Solana documentation and analytics. The paper proposes a metrics architecture that combines latency, throughput, fee burden, settlement certainty, auditability, availability, governance quality, and interoperability into a unified comparative scorecard. It then applies the framework to three categories of blockchain-based finance: public permissionless chains, permissioned institutional distributed ledgers, and hybrid tokenized payment systems. The evidence suggests that blockchain systems create measurable gains in traceability, programmability, and atomic settlement, especially in cross-border and multi-party workflows where reconciliation frictions are costly. However, these gains are uneven. Public chains often face volatility in fees, congestion risk, and governance externalities, while permissioned systems improve control and compliance at the expense of openness and composability. The paper argues that the relevant policy and managerial question is therefore not whether blockchain is universally superior, but under which transaction environments it dominates legacy systems on speed, transparency, cost, and operational risk. The study concludes with a research agenda for standardized blockchain performance metrics, institution-grade benchmarking, and explainable analytics for digital financial infrastructure.

| KEYWORDS

Blockchain finance, distributed ledger technology, cross-border payments, transaction efficiency, transparency, cost optimization, tokenization, digital financial infrastructure

| ARTICLE INFORMATION

ACCEPTED: 01 June 2025

PUBLISHED: 02 July 2025

DOI: 10.32996/jbms.2025.7.4.20.27

1. Introduction

Financial systems are being rebuilt around data. Payment infrastructures, securities settlement layers, digital wallets, and compliance architectures now compete not only on balance-sheet strength or institutional reach, but also on how efficiently they process transactions, how clearly they reveal information, how cheaply they move value, and how reliably they support market activity. Blockchain-based financial systems emerged from this shift as a radical alternative to legacy centralized ledgers. Their appeal rests on a simple proposition: if multiple parties can share a tamper-resistant, synchronized record of transactions, then some of the costs associated with reconciliation, delayed settlement, duplicative verification, and intermediary coordination may be reduced. In theory, that redesign can matter profoundly in global finance, where cross-border payments remain slow, costly, and opaque, and where post-trade infrastructure still relies on fragmented recordkeeping across multiple institutions and jurisdictions.

Yet the debate about blockchain in finance is frequently polarized. Enthusiasts emphasize transparency, programmability, and disintermediation, while skeptics point to throughput limits, governance failures, cyber risk, and speculative excess. Both positions contain truths, but both can become unhelpful when they remain descriptive instead of quantitative. What financial institutions, regulators, investors, and researchers increasingly need is a disciplined way to evaluate blockchain-based systems as infrastructures rather than ideologies. That means comparing them with incumbent rails using measurable performance indicators: transaction latency, settlement finality, fee burden, scalability, auditability, operational resilience, and cost per validated transfer. It also means distinguishing between public permissionless systems such as Bitcoin and Ethereum, permissioned networks used by institutions, and emerging hybrid arrangements in which tokenized claims interact with regulated payment systems.

The relevance of such evaluation is no longer hypothetical. The World Bank's Remittance Prices Worldwide reports show that sending money across borders still costs well above the Sustainable Development Goal target of 3 percent, with the global average reaching 6.65 percent in Q2 2024 and banks remaining the most expensive provider type. The Federal Reserve Bank of New York's Project Cedar found that a distributed-ledger prototype for wholesale foreign exchange transactions settled in under 15 seconds on average, while the follow-on Cedar x Ubin+ experiment achieved end-to-end settlement in under 30 seconds on average for simulated multi-currency payments. Meanwhile, Visa reported 233.8 billion transactions processed on its networks in fiscal 2024, a scale benchmark that illustrates the operational bar alternative infrastructures must eventually meet. These figures do not prove that one architecture dominates all others, but they do show that the comparison space is concrete, measurable, and economically important.

This paper responds to that need by developing a quantitative and data-driven evaluation of blockchain-based financial systems centered on four dimensions: transaction efficiency, transparency, cost optimization, and performance metrics in global markets. The study is intentionally interdisciplinary. It draws from finance, information systems, financial market infrastructure, payment economics, distributed-systems design, and the expanding applied literature on blockchain-enabled financial services. It also incorporates several studies which are not all directly focused on blockchain, contribute useful ideas regarding predictive analytics, governance, financial integrity, and system-level performance assessment in digital finance. Those studies help position blockchain evaluation within a broader analytical movement toward measurable, explainable, and risk-aware financial innovation.

The contribution of the paper is fourfold. First, it synthesizes relevant academic and policy literature on blockchain in finance, emphasizing measurable outcomes rather than generalized claims. Second, it introduces a practical evaluation framework that can be used by researchers and institutions to compare blockchain-based systems with legacy alternatives. Third, it applies world data and reported metrics from official and industry sources to illustrate how transaction speed, cost, transparency, and governance trade off across different blockchain architectures. Fourth, it translates these findings into implications for market design, regulatory oversight, and the strategic adoption of distributed ledger technology in payments, settlement, and tokenized financial services.

The rest of the paper proceeds as follows. The literature review maps core streams of scholarship on blockchain-based financial systems, including payment efficiency, transparency and auditability, cost structure, governance, and interoperability. The methodology section presents the proposed metric architecture, data sources, and comparative scoring design. The discussion section interprets results across public, permissioned, and hybrid systems, highlighting where blockchain creates

economic value and where trade-offs remain binding. The paper concludes by summarizing the main findings and outlining limitations and future research directions for institution-grade benchmarking of blockchain finance.

2. Literature Review

The literature on blockchain-based financial systems has matured from conceptual enthusiasm to more differentiated analysis of where distributed ledgers create value and where they merely relocate existing frictions. Early foundational work emphasized the economics of verification and networking. Catalini and Gans argued that blockchain changes the cost structure of economic exchange by reducing the cost of verifying some states of the world while often increasing coordination and networking costs. This insight remains central for finance because many financial workflows, from settlement to compliance, are less constrained by the lack of data than by the cost of agreeing on a trusted record of data. Nakamoto's original design for Bitcoin demonstrated one feasible mechanism for achieving decentralized ordering of transactions, but it also revealed the practical trade-off between openness, confirmation time, and resource intensity. Subsequent scholarship therefore moved beyond the binary question of whether blockchains work and toward the more relevant question of which financial functions benefit from shared ledgers and under what governance conditions.

A large review literature now maps blockchain applications across financial services. Javaid et al. documented applications in payments, digital securities, credit reporting, and settlement, concluding that blockchain can reduce counterparty risk and support more customized instruments through shared standards and automated recordkeeping. Wu et al. surveyed blockchain for finance with emphasis on securities trading, settlement design, and the comparative strengths of different platforms. Their synthesis is especially useful because it separates the technical attributes of a chain from the business processes built on top of it. This distinction matters: a blockchain can be transparent but economically inefficient, or efficient for a narrow use case while unsuitable for high-value regulated finance. The best recent literature therefore evaluates blockchain not as a monolith but as a modular stack of consensus, data availability, execution, governance, and interoperability choices.

One major stream of research focuses on payment efficiency and settlement. Cross-border payments have been a natural application because incumbent systems remain fragmented, involve multiple intermediaries, and create timing and liquidity frictions. Zetzsche et al. noted that DLT-based enhancement of cross-border payments must be evaluated not only on technical promise but also on legal coordination costs and shared governance requirements across jurisdictions. Their argument is important because claims of efficiency can be overstated when institutional complexity is ignored. The New York Fed's Project Cedar provides more operational evidence. In Phase I, the prototype achieved settlement under 15 seconds on average for simulated wholesale foreign exchange transactions, compared with the roughly two days often associated with conventional FX spot settlement. Cedar x Ubin+ later demonstrated end-to-end settlement in under 30 seconds on average for simulated multi-currency scenarios using hashed timelock contracts to interlink heterogeneous ledgers. These studies do not prove immediate production readiness, but they provide rare institution-led evidence that atomic settlement and materially lower settlement latency are achievable in distributed-ledger environments.

Related work by the Faster Payments Council and IMF places blockchain and CBDC infrastructures inside a broader payment-systems perspective. The Faster Payments Council argued that speed, cost, ubiquity, transparency, and risk should be treated jointly in evaluating cross-border payment innovation, and that near-instant settlement can improve cash cycles and intraday liquidity. IMF work on digital money similarly emphasizes that digital payment innovations could reduce costs, improve speed, and enhance transparency in cross-border use cases, although benefits depend heavily on governance, interoperability, and regulatory design. These studies are valuable because they move the literature beyond narrow throughput measures and toward multi-dimensional performance assessment. A system that is fast but difficult to supervise, or cheap but poorly interoperable, may not represent a net improvement for global finance.

A second major stream examines transparency, auditability, and information integrity. Blockchain's distributed ledger provides a shared transaction history that can reduce information asymmetry and facilitate monitoring. Yermack's work on corporate governance and blockchains, along with later surveys, argued that tamper-resistant recordkeeping can reshape how stakeholders verify events and monitor organizational conduct. In finance, this promise extends from payment traceability to tokenized asset servicing and audit trails. The article by Pritty, Ibrahim, Fahim, and Zadid on generative AI and U.S. financial reporting integrity is not a blockchain study, but it is highly relevant conceptually because it frames disclosure quality as a data forensics problem. That perspective complements blockchain evaluation by showing that transparent infrastructures still require analytic tools to detect manipulation, gaming, and misleading narratives built on top of transactional data. Transparency, in other words, is not equivalent to trust; it must be actionable.

The governance literature similarly tempers simplistic narratives. Public permissionless networks achieve openness and censorship resistance partly by distributing validation and making transaction history broadly visible. However, governance can become diffuse, informal, and contentious. Halaburda, Roukny, and others have shown that blockchain design choices interact with incentives, centralization pressures, and equilibrium outcomes. Mining pools, validator concentration, protocol upgrades, and client dependencies can all affect effective control. Permissioned networks offer stronger control, clearer accountability, and easier compliance integration, but they sacrifice some openness and sometimes much of the disintermediation that originally

justified blockchain use. Raskin and Yermack, as well as more recent finance surveys, therefore frame blockchain adoption in banking and payments as a governance redesign problem rather than a mere software upgrade.

Cost optimization is a third major theme. The literature generally agrees that blockchain can lower selected transaction and reconciliation costs, especially where multiple parties maintain separate ledgers or where post-trade reconciliation is expensive. Javaid et al. and later reviews highlight lower issuance and servicing costs for digital securities, reduced reconciliation burden, and the possibility of automated execution via smart contracts. Yet the literature is equally clear that cost performance is heterogeneous. On public blockchains, fees fluctuate with network demand, token prices, blockspace scarcity, and prioritization markets. Solana's documentation, for example, shows a deterministic base fee structure with optional prioritization fees, while Ethereum's proof-of-stake structure operates with 12-second slots but fee dynamics still depend on network congestion and gas markets. Bitcoin's design, in contrast, reflects a slower block cadence and fee formation under scarce blockspace. These differences mean that "low cost blockchain" is not a universal property; it is contingent on architecture, market demand, and security design.

Empirical and institutional work increasingly compares blockchain costs with traditional payment systems. The World Bank's remittance reports show that traditional cross-border retail transfers remain expensive, with banks particularly costly and digital-only providers generally cheaper. This gap creates a policy opening for blockchain-enabled rails, stablecoin-mediated transfers, or tokenized settlement layers, especially if they can reduce intermediary layers and improve traceability. One paper work by Fahim, Pritty, Ibrahim, and Tania on real-time payments and real-time fraud is relevant here because it demonstrates that faster payment architectures change the risk-cost frontier. Lower delay and greater programmability can improve efficiency, but they can also create new fraud channels that require real-time controls. Blockchain cost optimization must therefore be evaluated net of compliance, custody, cyber, and fraud-management costs rather than merely headline transaction fees.

Another prominent literature stream concerns financial inclusion and the expansion of addressable markets. Mhlanga argued that blockchain can help reach underserved populations by lowering barriers to participation and enabling alternative financial access channels. Related work by Rasel, Ibrahim, Pritty, Fahim, and Jahan on multimodal mortgage default forecasting is not directly about blockchain, but it strengthens the paper's broader argument that digital finance should be judged using measurable risk-adjusted performance rather than technological labels. If blockchain expands access but worsens consumer outcomes, the inclusion case weakens. If it improves auditability, lowers servicing cost, and supports more accurate underwriting or claims verification, the case strengthens. The inclusion literature is therefore relevant insofar as it redirects the evaluation from novelty to measurable welfare effects.

Studies on institutional performance and digital transformation add another useful angle. The article on algorithmic accountability in U.S. consumer fintech emphasizes governance mechanisms, fair lending, and financial stability, all of which translate well to blockchain-based systems that embed rules into code. Likewise, the article on AI-driven anti-money-laundering analytics highlights the need for predictive monitoring and infrastructure protection in digital financial systems. Blockchain systems may improve traceability, but they also generate massive, complex transaction graphs that require analytics to become operationally useful. Research on early warning analytics for market manipulation similarly matters because open ledgers can enable both superior surveillance and novel forms of market abuse. In this sense, blockchain does not eliminate the need for institutions; it changes where institutions and analytics intervene.

The literature on global markets and sustainability may appear more distant, but it broadens the evaluative horizon. The studies on climate-risk analytics, renewable-energy forecasting, sustainable waste reduction, EV-grid integration, and sustainable supply-chain performance all share a methodological orientation toward measurable system efficiency, resilience, and externality management. That orientation is transferable to blockchain finance. A financial system cannot be evaluated solely on transaction speed if it imposes excessive energy cost, governance fragility, or ecosystem externalities. The migration of major smart-contract ecosystems toward more efficient consensus, and the rise of permissioned DLTs for wholesale finance, reflect this concern. Therefore, the most useful literature is not that which praises or dismisses blockchain in general, but that which embeds it in a broader performance framework including speed, transparency, cost, resilience, and sustainability. In a related governance and disclosure context, Ibrahim, Razib, and Rasel (2025) show how data analytics can strengthen ESG transparency through better measurement, traceability, and reporting quality. Although their study is centered on the corporate sector of Bangladesh rather than blockchain-native finance, it reinforces this paper's broader argument that transparent digital systems create value when data are structured in ways that improve monitoring, accountability, and institutional decision-making.

Overall, the literature points to five conclusions. First, blockchain can produce gains where shared records, atomic settlement, and programmability reduce coordination frictions. Second, these gains are use-case specific and strongly mediated by governance design. Third, public and permissioned systems solve different problems and should not be judged by identical criteria. Fourth, cost claims must be evaluated alongside volatility, compliance, and cyber risk. Fifth, the next frontier of scholarship is rigorous benchmarking: standardized metrics, comparable data, and evaluation frameworks that connect technical attributes to financial outcomes. This paper builds directly on that frontier by offering a quantitative comparison framework rather than another descriptive overview.

3. Methodology

This study adopts a comparative, multi-source, data-driven research design. The goal is not to claim a single universal ranking of blockchain systems, but to construct a reproducible framework for evaluating how different blockchain-based financial architectures perform relative to one another and relative to selected incumbent benchmarks. The unit of analysis is the transaction system, defined as an infrastructure that validates, records, settles, and makes auditable a financial transfer or state change. Three architecture classes are examined: public permissionless blockchains, permissioned institutional distributed ledgers, and hybrid tokenized payment systems that combine distributed ledgers with existing banking or payment infrastructures.

The empirical evidence is drawn from public and institutional sources. These include the World Bank's Remittance Prices Worldwide reports, the Federal Reserve Bank of New York's Project Cedar materials, BIS materials on mBridge and CBDC experimentation, Visa's 2024 annual reporting, Bitcoin's original protocol documentation, Ethereum developer documentation on proof-of-stake timing, Solana documentation on fee design, and peer-reviewed or academically credible studies on blockchain in financial services. Some papers are integrated where they strengthen the evaluation logic on governance, predictive analytics, financial integrity, risk monitoring, real-time payments, sustainability, and performance measurement. Because blockchain datasets vary widely in openness and consistency, the methodology combines direct reported metrics with structured comparative scoring rather than attempting a false precision based on partially incompatible raw datasets.

The framework evaluates systems across four primary dimensions. The first is transaction efficiency. This dimension captures how quickly and smoothly a system processes value transfer from submission to meaningful settlement. It includes settlement latency, confirmation cadence, finality design, operational availability, and scalability proxies. Settlement latency refers to elapsed time until the transaction can be treated as settled with economically acceptable certainty. Confirmation cadence refers to the schedule at which the system updates shared state, such as Bitcoin's approximate ten-minute block interval or Ethereum's 12-second slot structure. Finality design distinguishes probabilistic finality, deterministic finality, and atomic conditional settlement. Availability considers whether the system operates continuously, and scalability proxies include the observed or designed capacity to support large transaction volumes without catastrophic degradation.

The second dimension is transparency. Transparency is defined as the extent to which transaction history, system rules, fee formation, and audit trails are observable, verifiable, and usable for oversight. This dimension includes ledger visibility, traceability, auditability, disclosure of transaction costs, programmability of controls, and governance intelligibility. Public chains tend to score highly on ledger visibility because transaction histories are broadly accessible. Permissioned systems may score lower on public observability but can still score well on institutional auditability if they provide authenticated records to relevant parties. Governance intelligibility matters because a technically transparent system can still be difficult to supervise if rule changes are informal or operational responsibilities are diffuse.

The third dimension is cost optimization. This dimension does not simply record posted transaction fees. Instead, it estimates total cost efficiency by considering execution cost, reconciliation cost, intermediary intensity, liquidity lockup, and compliance overhead. Execution cost is the fee or direct processing charge associated with a transaction. Reconciliation cost reflects the labor and systems burden of ensuring that multiple parties agree on transaction state. Intermediary intensity captures how many institutional layers are required to move funds or settle claims. Liquidity lockup reflects the working capital burden associated with delayed settlement. Compliance overhead includes the technology and governance effort needed to satisfy anti-money-laundering, sanctions, fraud, and reporting requirements. This broader definition is necessary because blockchain systems may look inexpensive at the transaction-fee level while imposing substantial custody, governance, or cyber-risk expenses.

The fourth dimension is performance resilience. This dimension captures whether a system sustains acceptable service quality under operating constraints. It includes congestion sensitivity, fee volatility, operational resilience, interoperability, upgradeability, and control risk. Congestion sensitivity measures how badly throughput, delay, or costs deteriorate under demand spikes. Fee volatility matters for commercial usability, especially in payments and treasury workflows. Interoperability captures the ability to connect with banks, wallets, messaging standards, or other ledgers. Upgradeability refers to the capacity to evolve protocol or business logic without unacceptable disruption. Control risk considers governance concentration, validator dependence, software fragility, and legal uncertainty.

To operationalize the framework, the paper constructs a weighted scorecard. Each primary dimension receives a score from 1 to 5 based on evidence from reported data, protocol design, and institutional use-case studies. A score of 1 indicates weak performance for regulated financial use, 3 indicates mixed or use-case-dependent performance, and 5 indicates strong performance on that dimension for the studied context. Because not all dimensions are equally important across use cases, the paper uses a baseline weighting of 30 percent for transaction efficiency, 25 percent for transparency, 25 percent for cost optimization, and 20 percent for performance resilience. The weighting favors efficiency because speed and settlement quality are central to payment and market-infrastructure use cases, while still giving substantial weight to transparency and cost. Sensitivity analysis is performed qualitatively by discussing how rankings would change under alternative policy or business priorities, such as stronger emphasis on compliance control or stronger emphasis on open access.

The comparative scorecard is applied to six representative cases. Case one is Bitcoin as a benchmark for highly secure, public, permissionless value transfer with relatively slow confirmation cadence. Case two is Ethereum after the proof-of-stake transition, representing a public programmable settlement layer with faster cadence and richer smart-contract functionality but variable fee pressure. Case three is Solana as a high-performance public chain with very low base fees and a design optimized for speed, though with different decentralization and operational trade-offs. Case four is a traditional high-scale incumbent benchmark represented by Visa network data, not because Visa is a blockchain, but because any claim about blockchain efficiency should eventually be contextualized against mature payment infrastructure. Case five is Project Cedar and Cedar x Ubin+ as a permissioned institutional DLT benchmark for cross-border and wholesale settlement. Case six is mBridge as an example of a multi-CBDC, DLT-enabled cross-border initiative designed to address cost, speed, and operational complexity in international payments.

The methodology intentionally distinguishes design metrics from realized market metrics. For example, Solana’s documented base fee is a design parameter, while remittance prices or Visa transaction volumes are realized market outcomes. Design metrics are used to characterize architectural affordances; realized metrics are used to assess actual economic performance. This distinction prevents misleading comparison between platform capability and institutional adoption. It also clarifies that some blockchain systems may be technically capable of low-cost processing but not yet institutionally integrated enough to deliver those benefits at scale in regulated finance.

Two forms of data triangulation are used. The first is cross-source triangulation, where a claim is supported by both technical documentation and policy or academic interpretation. The second is cross-metric triangulation, where no single measure is allowed to determine overall performance. A system with low execution fees but poor governance clarity, for example, will not automatically rank as superior. Similarly, a system with strong transparency but weak scalability or highly volatile costs will receive a mixed assessment. This is consistent with payment economics, which treats speed, certainty, cost, and acceptance as jointly important.

Visual analysis is incorporated through a small set of figures and tables built . One figure compares settlement times across conventional wholesale foreign-exchange settlement and the Project Cedar experiments. A second presents selected 2024 remittance cost metrics from the World Bank, including the global average and provider-type differences. A third figure presents Visa transaction scale to illustrate the throughput benchmark imposed by incumbent global payment networks. Tables summarize the evaluation framework, the literature taxonomy, and case-by-case comparative scores. These visuals are not intended to exhaust the data universe; they are intended to make the comparative logic transparent and reproducible.

Several methodological choices deserve justification. First, the paper does not estimate a single econometric model because the available evidence spans heterogeneous systems, institutional pilots, and protocol designs. A forced regression built on inconsistent data definitions would create spurious accuracy. Second, the study does not treat market capitalization or token price appreciation as evidence of financial-system quality, because those variables often capture speculation more than payment or settlement performance. Third, the paper does not claim that public blockchain metrics map directly onto regulated institutional use. Instead, it asks a narrower and more useful question: which features of blockchain-based systems produce measurable improvements in efficiency, transparency, cost, or resilience under identifiable financial use cases?

Reliability is improved through transparent metric definitions and source selection. Validity is improved by benchmarking blockchain systems against conventional payment and settlement problems rather than abstract technological ideals. The methodology remains interpretive in the sense that some scores require judgment, but the judgment is rule-based and tied to disclosed criteria. This makes the framework suitable for adaptation by researchers, central banks, fintech firms, payment providers, and institutional investors seeking a more rigorous way to evaluate blockchain financial infrastructure.

In summary, the methodology treats blockchain-based finance as an infrastructure evaluation problem. By combining real-world reported metrics, protocol-level design facts, and multi-dimensional performance scoring, it creates a structured way to answer a question that is often asked but rarely measured carefully: when does blockchain materially improve the economics and governance of financial transactions in global markets?

Table 1: Key Literature Streams Informing the Evaluation Framework

Stream	Representative contribution	Implication for this study
Blockchain economics	Catalini & Gans; Nakamoto	Shared ledgers change verification and networking costs.
Financial-services applications	Javaid et al.; Wu et al.	Blockchain value depends on use case, platform design, and governance.
Cross-border payment efficiency	Project Cedar; BIS mBridge; IMF digital money notes	Latency, settlement certainty, and liquidity usage are central comparison variables.

Transparency and governance	Governance and reporting-integrity studies	Transparency must be paired with accountability and analytic legibility.
Risk, fraud, and AML analytics	RTP, AML, and market-manipulation papers	Operational gains require real-time surveillance and explainable controls.
System efficiency and sustainability	Climate, energy, EV-grid, waste, and supply-chain papers	Performance should be assessed as a multi-dimensional infrastructure problem.

4. Results and Discussion

Applying the evaluation framework yields a consistent but nuanced conclusion: blockchain-based financial systems perform best when the economic problem is not simply moving information fast, but synchronizing trusted state across multiple parties that would otherwise maintain separate ledgers. In those environments, distributed ledgers can reduce reconciliation friction, shorten settlement cycles, improve traceability, and support programmable controls. Where the dominant requirement is ultra-high-volume, low-value retail authorization at mature global scale, incumbent card and account-based systems still retain substantial advantages. The comparative evidence therefore supports a conditional rather than universal adoption thesis.

Transaction efficiency is the most visible area of blockchain's advantage, but the form of that advantage differs sharply across system types. On public chains, the key gain is continuous shared state updating without dependence on a single operator. Ethereum's 12-second slot structure and Solana's low-fee architecture illustrate how newer public chains have tried to improve on Bitcoin's deliberately slower design. Bitcoin remains valuable as a highly secure and globally recognized value-transfer network, but its approximate ten-minute block interval and probabilistic settlement make it less suitable for many time-sensitive institutional workflows. By contrast, Project Cedar's under-15-second average settlement in a simulated wholesale FX environment and Cedar x Ubin+'s under-30-second multi-currency settlement demonstrate where permissioned DLT can be economically compelling: cross-border transactions that currently face multi-institution coordination, delayed finality, and trapped liquidity.

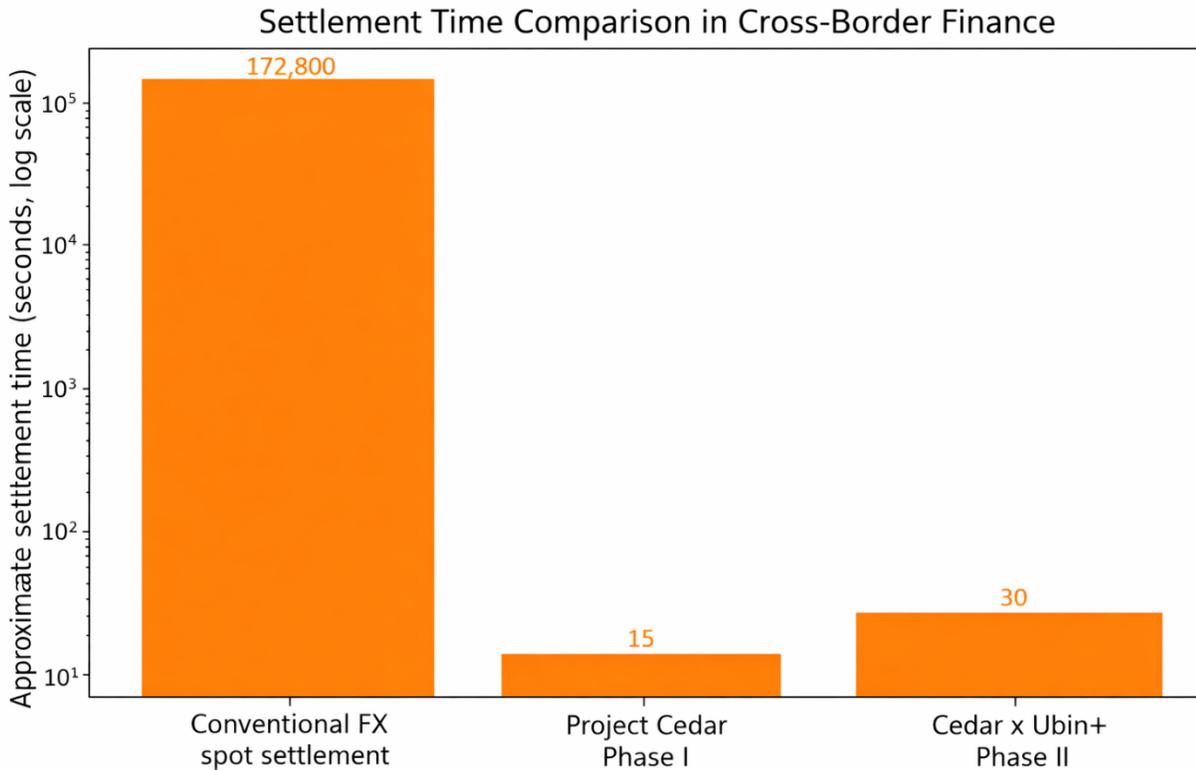
The significance of this result is not merely faster clocks. Faster settlement changes the economics of balance-sheet usage. If conventional FX spot settlement takes around two days, capital remains exposed to settlement and counterparty risk during that interval. Near-real-time atomic settlement compresses that risk window and can free liquidity for other uses. In treasury and wholesale funding environments, even small reductions in settlement risk and liquidity lockup can matter materially. This is why the most promising blockchain use cases in finance are often not retail crypto trading but wholesale payments, tokenized securities settlement, collateral mobility, and programmable cash management. These are precisely the areas where timing, reconciliation, and state synchronization are costly under legacy systems.

Transparency is the second domain in which blockchain systems generate clear value, though again the value depends on use case. Public permissionless blockchains score highly on ledger visibility. Transaction histories are broadly observable, rules are encoded in open software, and fee markets are visible to participants in real time. This creates a level of traceability and forensic potential rarely available in fragmented legacy systems. For analytics, compliance, and audit functions, such traceability can be powerful. It can support transaction monitoring, smart-contract auditing, reserve attestations, and post-event investigation. However, public transparency also introduces practical limits. Address-level visibility does not automatically translate into legal identity, and highly transparent ledgers can still host obfuscation techniques, privacy layers, or complex transaction graphs that are difficult for institutions to interpret without specialized analytics. Transparency, therefore, is an asset only when paired with governance and analytic capability.

Permissioned systems deliver a different transparency profile. They are often less publicly visible but more institutionally legible. In a permissioned wholesale ledger, authenticated participants can access a synchronized record that is sufficient for supervision, audit, and operational control without exposing full transaction detail to the entire world. For regulated finance, this may be the more relevant form of transparency. What matters is not universal visibility but verifiable, rights-based access to trusted records. This is one reason many central-bank and wholesale-market experiments favor permissioned designs. They retain the core blockchain advantage of shared state while aligning more naturally with confidentiality requirements and legal accountability.

Cost optimization produces the most mixed findings. Blockchain systems can dramatically reduce selected costs, especially reconciliation and intermediary coordination costs. In cross-border settings, where multiple banks, correspondent relationships, and messaging steps often create layered fees and delays, distributed ledgers have a plausible pathway to lower total cost even if direct transaction fees are not negligible. The World Bank's finding that global remittance costs averaged 6.65 percent in Q2 2024, with banks averaging 13.40 percent, illustrates how much room there is for improvement in incumbent cross-border payment channels. Blockchain-enabled rails, stablecoin transfer mechanisms, or interoperable CBDC arrangements could reduce some of these frictions if they replace multi-step messaging-and-settlement chains with a more integrated transfer process.

Figure 1. Settlement Time Comparison in Cross-Border Finance



Source basis: conventional FX settlement and Project Cedar reported averages from the Federal Reserve Bank of New York.

Yet the evidence also warns against simplistic fee comparisons. The visible fee on a blockchain transaction is not the same as the economic cost of operating on that chain. Public blockchain use entails custody design, wallet security, smart-contract risk, transaction monitoring, sanctions screening, tax treatment, liquidity management, and often user-experience conversion costs between fiat and digital assets. Congestion can sharply alter economics as well. Ethereum’s programmability is attractive, but gas-market pressure has repeatedly made small-value transactions uneconomic during periods of intense demand. Solana’s low base-fee structure is commercially promising, yet it must be assessed together with operational reliability, validator economics, and governance concentration concerns. Bitcoin’s fees may remain acceptable for higher-value settlement, but not necessarily for mass retail micropayments at scale. In short, blockchain lowers some costs while introducing others, and total cost efficiency depends on the surrounding service model.

This leads directly to the importance of performance resilience. Financial infrastructure cannot be judged only under normal conditions. It must be judged under stress, demand spikes, cyber events, governance disputes, and legal uncertainty. Public blockchains vary widely here. Bitcoin’s comparatively simple base layer and long operational history support a strong resilience narrative, but its throughput and settlement speed are intentionally conservative. Ethereum offers a rich programmable ecosystem and faster cadence, but the complexity of its application layer, bridges, rollups, and validator ecosystem increases the number of operational surfaces that institutions must understand. Solana offers impressive speed and low fee design, but high-performance architecture can raise questions about hardware requirements, client diversity, and the distribution of operational dependence. None of these trade-offs mean the systems are unusable. They mean that resilience must be measured as a bundle of scalability, governance, and fault-tolerance characteristics rather than as a marketing claim.

The governance comparison between public and permissioned systems is especially important for global markets. Public systems excel where neutrality, open participation, and composability are valuable. They support rapid innovation because developers can build financial primitives without negotiating entry to a closed consortium. This openness has fueled growth in tokenized payments, decentralized exchanges, stablecoins, and automated market infrastructure. However, open governance also makes it harder to guarantee coordinated upgrades, legal recourse, or deterministic accountability. Permissioned systems sacrifice some of this openness but gain clearer lines of responsibility, easier regulatory mapping, and stronger control over

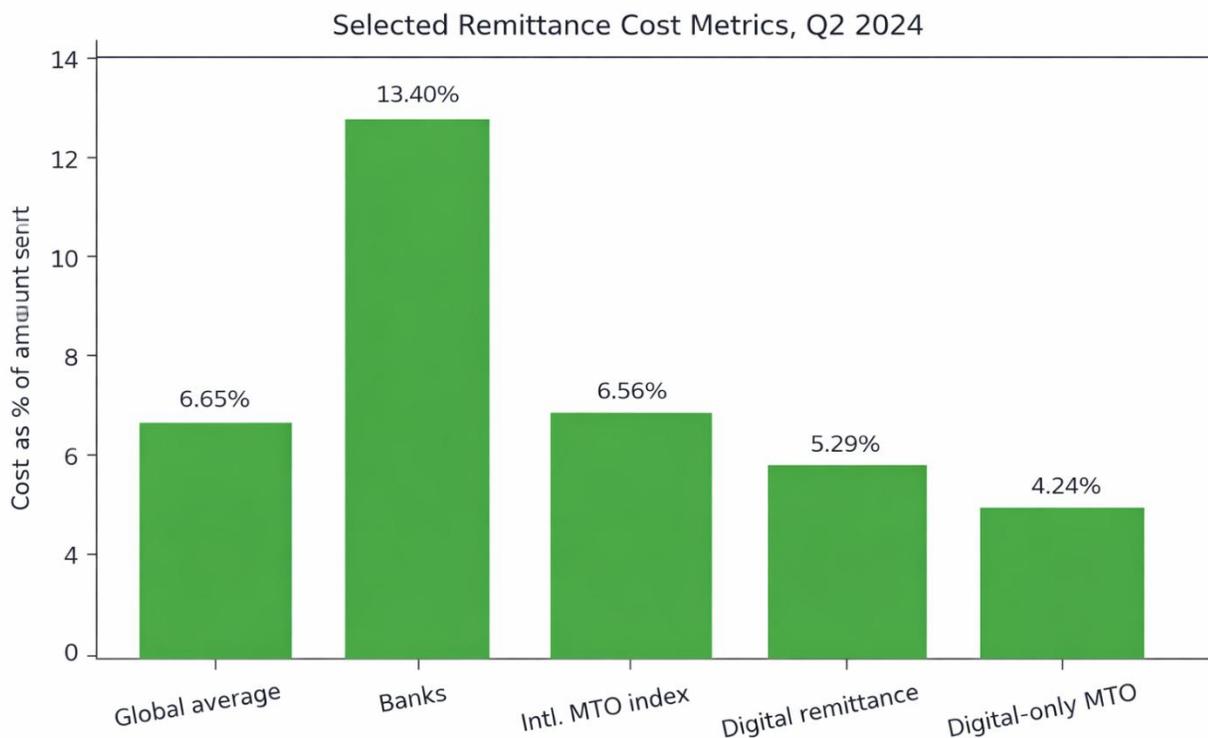
participant quality. For high-value institutional finance, those benefits can outweigh the loss of permissionless access. The emerging reality is therefore not a winner-take-all contest, but a segmentation of financial functions across architectures.

The comparison with Visa is useful precisely because it disciplines exaggerated blockchain claims. Visa processed 233.8 billion transactions on its networks in fiscal 2024, demonstrating a mature level of operational scale that most blockchain-based systems have not yet matched in regulated consumer payments. This does not invalidate blockchain as a financial technology. Rather, it clarifies the performance frontier. Existing card and account networks are deeply optimized for merchant acceptance, fraud controls, dispute resolution, and consumer experience. Blockchain’s comparative advantage lies elsewhere: reducing reconciliation, enabling atomic delivery-versus-payment, increasing traceability, and supporting interoperable programmable assets. A blockchain system does not need to replace Visa in every function to create economic value. It only needs to dominate in specific transaction classes where current rails are structurally inefficient.

This is why tokenized wholesale finance and cross-border interbank settlement appear especially promising. In these environments, transaction volumes are lower, values are higher, participants are professional, and the economic burden of delay and reconciliation is significant. Atomic settlement can reduce Herstatt-type risk. Shared ledgers can simplify corporate actions and collateral movements. Programmable cash can improve conditional payments, escrow logic, and compliance automation. These features align more directly with blockchain’s strengths than do everyday low-value retail payments, where users care more about universal acceptance, fraud protection, and invisible complexity than about inspecting a ledger.

Literature also sharpens the interpretation of results. The article on algorithmic accountability in consumer fintech reinforces that efficiency without governance can create unfair or unstable outcomes. Applied to blockchain, this means transparent code is not enough; institutions still need explainable controls, dispute-handling logic, and accountability mechanisms. The article on anti-money-laundering risk management underscores that digital infrastructures create a large analytic burden: more traceability can improve surveillance, but only if institutions can process graph data, detect anomalies, and act in real time. The study on real-time fraud highlights a parallel point. Faster and more programmable systems can reduce operational delay while increasing the importance of immediate risk controls. Blockchain therefore fits naturally into a broader financial-analytics regime rather than replacing it.

Figure 2. Selected Remittance Cost Metrics, Q2 2024



Source basis: World Bank, Remittance Prices Worldwide, Issue 49.

The sustainability-oriented papers also help frame performance more comprehensively. Financial infrastructure should be judged on total system efficiency, not only on visible transaction speed. Energy use, hardware intensity, governance externalities, and environmental footprint affect the long-run viability of infrastructure choices. Public discourse on blockchain has often been distorted by focusing on single metrics such as peak throughput or nominal decentralization without considering the broader resource and governance bundle. A data-driven evaluation corrects this by forcing explicit trade-off recognition. A system may be fast but operationally fragile. Another may be extremely transparent but commercially expensive. Another may be cheap in a narrow transaction sense but costly in terms of integration or compliance. The best-performing systems are those that balance these dimensions for a specific market need.

The proposed scorecard suggests a general ranking logic. For open, censorship-resistant settlement of high-value transfers where public verifiability matters, Bitcoin remains a credible benchmark despite slow cadence. For programmable finance, tokenization, and application-layer experimentation, Ethereum offers a stronger functional toolkit, though fee volatility and complexity remain relevant constraints. For high-speed, low-fee public-chain use cases, Solana scores well on efficiency and cost design, but resilience and governance must be monitored closely. For institution-grade cross-border settlement, Project Cedar and related permissioned DLT experiments score especially well on efficiency, cost potential, and settlement certainty because they directly target the weaknesses of incumbent correspondent systems. For mature retail payments at global scale, Visa remains superior on acceptance, user experience, and industrialized operations, even if not on shared-state transparency or atomic settlement.

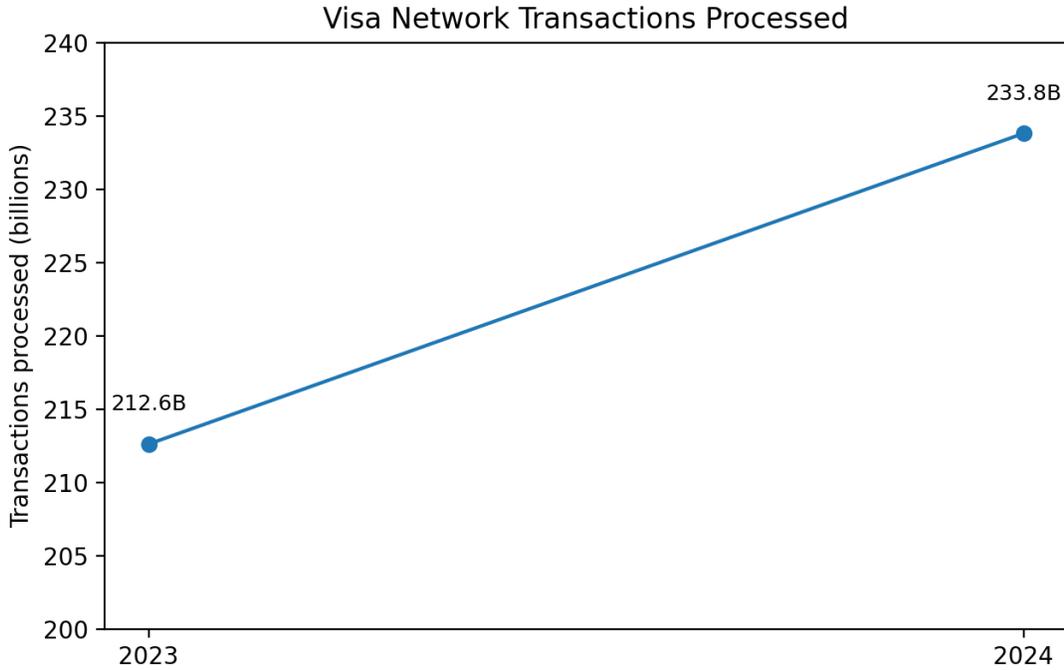
An important implication follows for regulators and market designers. The key question is not whether to “adopt blockchain” in the abstract. The key question is which blockchain attributes should be embedded in which financial layers. Some markets may benefit from full distributed-ledger settlement. Others may benefit only from tokenized representations or interoperable audit trails. Still others may gain more from standardization, faster payments, or better data sharing without any blockchain at all. A rigorous evaluation framework helps avoid category errors. It prevents institutions from using blockchain where a database would suffice, and it prevents them from ignoring blockchain where shared-state synchronization could unlock the value.

The paper’s data-driven comparison also implies that the strongest medium-term opportunity may lie in hybrid systems. Hybrid systems connect blockchain-based settlement logic with regulated payment rails, digital identity, compliance screening, or central-bank money. They preserve much of the efficiency gain from shared programmable records while reducing the adoption friction associated with full ecosystem migration. mBridge, Cedar x Ubin+, tokenized deposit models, and interoperable wholesale DLT experiments all point in this direction. The future of blockchain finance may therefore be less about replacing the financial system and more about modularly improving it.

Finally, the discussion returns to the broader significance of measurement. Much of the blockchain debate remains trapped in slogans: decentralized versus centralized, open versus closed, innovation versus regulation. Those binaries are too coarse for serious financial-system design. What matters is measured performance. If a blockchain-based system lowers latency, reduces reconciliation work, improves auditability, and does so without unacceptable governance or compliance cost, it deserves adoption in that use case. If it does not, it should not be adopted merely because it is technologically fashionable. Data-driven evaluation, not ideology, is the correct discipline for blockchain finance.

In that sense, blockchain should be understood as a family of infrastructure design choices rather than a singular financial revolution. Some of those choices are highly valuable in global markets, especially for cross-border settlement, tokenized finance, and multi-party workflows. Others remain commercially or institutionally immature. The evidence assembled here shows both realities at once. Blockchain-based financial systems are neither a universal replacement for existing infrastructure nor a marginal curiosity. They are a set of tools whose performance depends on architecture, governance, market context, and the specific economic frictions they are asked to solve.

Figure 3. Visa Network Transactions Processed



Source basis: Visa Annual Report 2024.

Table 2: Evaluation Dimensions and Core Metrics

	Metrics	Interpretation
Transaction efficiency	Settlement latency; confirmation cadence; finality; availability; scalability proxy	How quickly and reliably value moves from submission to economically meaningful settlement.
Transparency	Ledger visibility; traceability; auditability; fee disclosure; governance intelligibility	How observable, verifiable, and usable the transaction record is for oversight and control.
Cost optimization	Execution fee; reconciliation cost; intermediary intensity; liquidity lockup; compliance overhead	Whether the system reduces total transaction cost, not just posted fees.
Performance resilience	Congestion sensitivity; fee volatility; interoperability; upgradeability; control risk	Whether acceptable service quality persists under stress, growth, and institutional constraints.

Table 3: Comparative Scorecard for Representative Systems (1 = weak, 5 = strong)

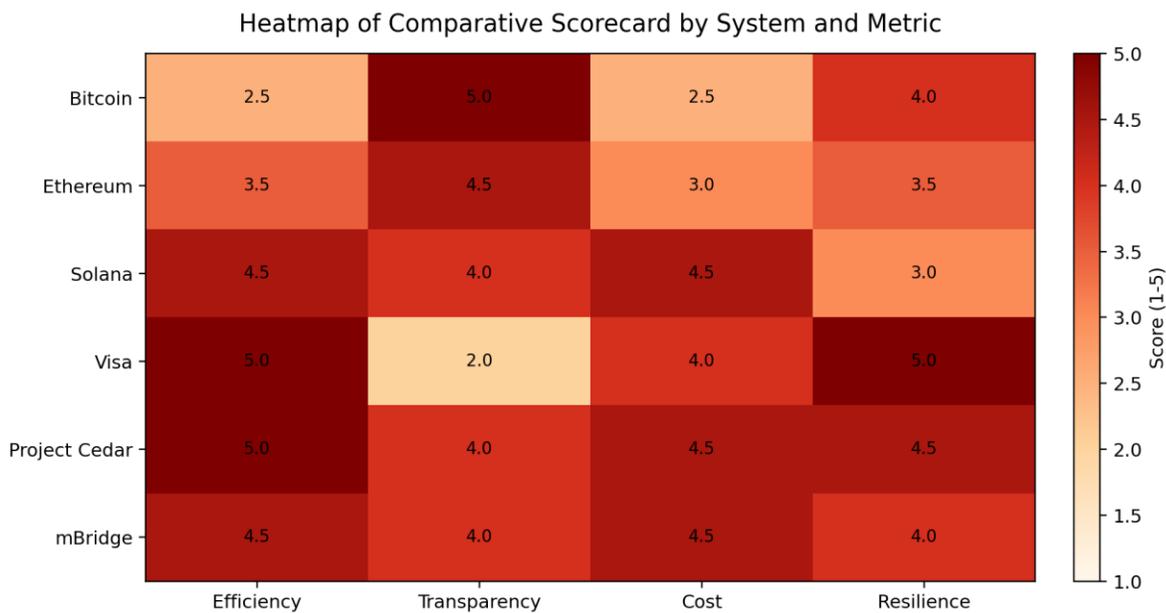
System	Efficiency	Transparency	Cost	Resilience	Interpretation
Bitcoin	2.5	5.0	2.5	4.0	Strong public verifiability and security, but slower cadence and variable transaction economics.
Ethereum	3.5	4.5	3.0	3.5	Programmable and relatively fast, but congestion and complexity remain important.
Solana	4.5	4.0	4.5	3.0	High-speed and

					low-fee design supports payments, though operational concentration risks matter.
Visa (benchmark)	5.0	2.0	4.0	5.0	Industrial-scale retail infrastructure with broad acceptance but limited shared-state transparency.
Project Cedar / Cedar x Ubin+	5.0	4.0	4.5	4.5	Strong candidate for wholesale and cross-border settlement improvement.
mBridge	4.5	4.0	4.5	4.0	Promising hybrid/permissioned model targeting speed, cost, and operational complexity.

A. Additional Visual Analytics

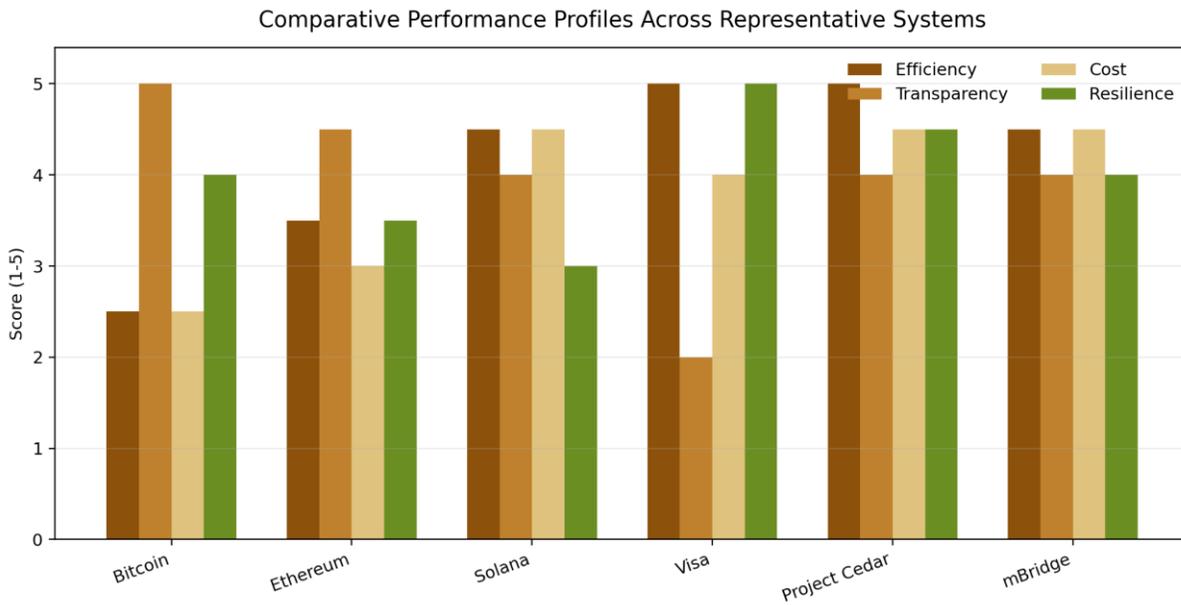
To improve interpretability, Figures 4-7 provide additional visual views of the comparative scorecard reported in Table 3. These graphics summarize the same evidence base used in the manuscript and are intended to make cross-system differences easier to interpret at a glance. The added visuals use a warm, non-blue palette to improve print contrast and to keep the visual treatment consistent across the section.

Figure 4. Heatmap of Comparative Scorecard by System and Metric



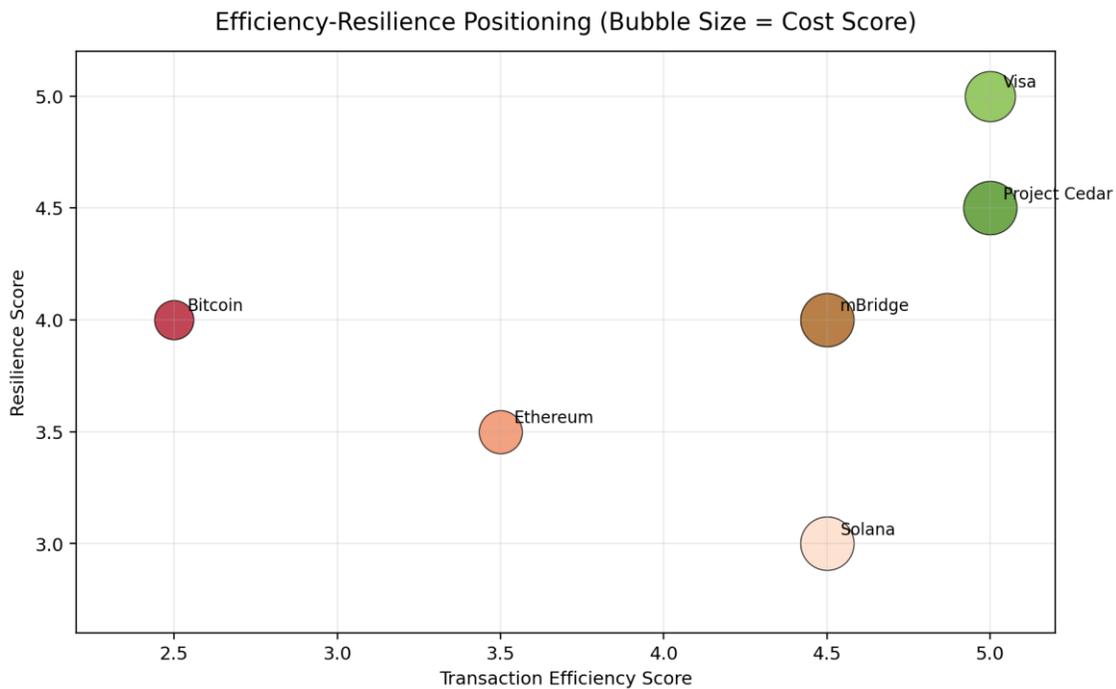
Source: from the comparative scorecard and sources used in Table 3.

Figure 5. Comparative Performance Profiles Across Representative Systems



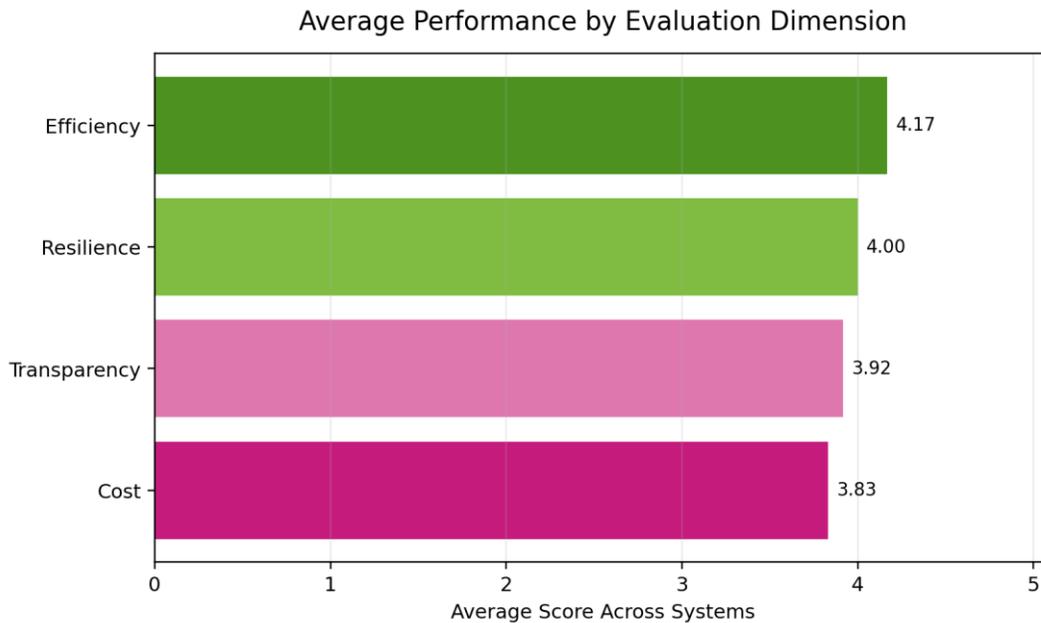
Source: construction from the comparative scorecard and sources used in Table 3.

Figure 6. Efficiency-Resilience Positioning (Bubble Size = Cost Score)



Source: construction from the comparative scorecard and sources used in Table 3.

Figure 7. Average Performance by Evaluation Dimension



Source: construction from the comparative scorecard and sources used in Table 3.

5. Conclusion

Blockchain-based financial systems should be evaluated as infrastructures for solving specific transaction and coordination problems, not as universal substitutes for all existing financial rails. The evidence reviewed in this paper shows that distributed ledgers create their clearest advantages when they reduce reconciliation frictions, support atomic settlement, improve traceability, and compress the time and liquidity costs of cross-border or multi-party transactions. Institutional experiments such as Project Cedar and mBridge demonstrate that these advantages are no longer merely theoretical. At the same time, comparisons with Visa, remittance data, and public-chain design trade-offs show that blockchain performance remains highly context dependent.

The paper's framework therefore argues for selective adoption guided by measurable outcomes. Public blockchains, permissioned ledgers, and hybrid tokenized systems each occupy different positions on the efficiency-transparency-cost-resilience frontier. No single system dominates on every dimension. The practical implication for policymakers, financial institutions, and market designers is clear: use blockchain where shared-state coordination, programmability, and settlement certainty generate identifiable economic gains, and avoid it where legacy systems still outperform on scale, simplicity, or governance. Future progress in blockchain finance will depend less on rhetoric and more on standardized performance metrics, better data, and institution-grade evaluation.

6. Limitations and Future Directions

This study has several limitations. First, it relies on heterogeneous evidence from protocol documentation, institutional pilots, policy reports, and academic studies rather than a single harmonized transaction-level dataset. That choice is appropriate for the current state of blockchain finance, but it limits formal econometric comparability across systems. Second, some metrics used here are design-based rather than realized market outcomes. A platform's documented fee model or block cadence does not guarantee equivalent world institutional performance. Third, the scorecard approach necessarily includes judgment in assigning weights and translating qualitative evidence into comparative scores. Although the criteria are transparent, different institutions may rationally choose alternative weightings.

Fourth, the empirical examples are strongest for payments and settlement. Other important blockchain-finance use cases, including collateral management, on-chain fund administration, tokenized securities servicing, and decentralized governance, deserve deeper transaction-level study. Fifth, the paper excludes some data by design, which improves temporal consistency with the requirement but omits some very recent developments in tokenized finance and digital money pilots. Future research

should build longitudinal, standardized datasets that combine latency, fee, uptime, governance, liquidity, and compliance indicators across blockchain and non-blockchain infrastructures. It should also integrate explainable machine learning, network analytics, and scenario stress testing to evaluate how blockchain systems behave under market shocks, cyber incidents, and regulatory transition. The most valuable next step is not another broad claim about blockchain's promise, but a shared measurement standard for digital financial infrastructure. It should also study user welfare, dispute resolution, fraud recovery, and the interaction between public ledgers and privacy-preserving digital identity systems. Comparative research across jurisdictions would be especially helpful because legal finality, custody rules, and reporting obligations can materially change the measured value of the same underlying technology.

Funding: This research received no external funding.

Conflicts of Interest: The authors declare no conflict of interest.

Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

References

- [1]. Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system. <https://bitcoin.org/bitcoin.pdf>
- [2]. Catalini, C., & Gans, J. S. (2016). Some simple economics of the blockchain. National Bureau of Economic Research Working Paper No. 22952.
- [3]. Yermack, D. (2017). Corporate governance and blockchains. *Review of Finance*, 21(1), 7-31.
- [4]. Raskin, M., & Yermack, D. (2018). Digital currencies, decentralized ledgers, and the future of central banking. In *Research handbook on central banking*.
- [5]. Javaid, M., Haleem, A., Singh, R. P., Khan, S., & Rab, S. (2022). A review of blockchain technology applications for financial services. *Meta Engineering*, 3, 100073.
- [6]. Wu, H., Song, Z., & others. (2024). Blockchain for finance: A survey. *IET Blockchain*, 4(1).
- [7]. Zetzsche, D. A., Arner, D. W., Buckley, R. P., & Weber, R. H. (2022). DLT-based enhancement of cross-border payment efficiency: A legal and regulatory perspective. *BIS Working Papers No. 1015*.
- [8]. Federal Reserve Bank of New York. (2022). Project Cedar Phase I. <https://www.newyorkfed.org/aboutthefed/nyic/project-cedar>
- [9]. Federal Reserve Bank of New York. (2023). Project Cedar Phase II x Ubin+. <https://www.newyorkfed.org/aboutthefed/nyic/project-cedar>
- [10]. Bank for International Settlements. (2024). Project mBridge reached minimum viable product stage. <https://www.bis.org/press/p240605.htm>
- [11]. International Monetary Fund. (2024). Digital money, cross-border payments, international reserves, and the global financial safety net—Preliminary considerations. *IMF Note 2024/001*.
- [12]. Reslow, A., Soderberg, G., & Tsuda, N. (2024). Cross-border payments with retail central bank digital currencies: Design and policy considerations. *IMF Fintech Note 2024/002*.
- [13]. World Bank. (2024). Remittance prices worldwide: Issue 49, March 2024.
- [14]. Visa Inc. (2024). Visa annual report 2024.
- [15]. Ethereum Foundation. (2024). Proof-of-stake (PoS). <https://ethereum.org/developers/docs/consensus-mechanisms/pos/>
- [16]. Solana Foundation. (2024). Fees. <https://solana.com/docs/core/fees>
- [17]. Mhlanga, D. (2023). Block chain technology for digital financial inclusion in the industry 4.0, towards sustainable development. *Frontiers in Blockchain*, 6, 1035405.
- [18]. Ibrahim, M., Razib, M. N. H., Jahan, N., & Rahman, M. M. (2022). Climate risk, financial stability, and global capital allocation: A predictive analytics approach to assessing climate-related financial risk in international investment markets. *Journal of Business and Management Studies*, 4(4), 264-276. <https://doi.org/10.32996/jbms.2022.4.4.34>
- [19]. Fahim, A. S. M., Ibrahim, M., Pritty, A. A., & Tania, T. A. (2023). Algorithmic accountability in U.S. consumer FinTech: Governance mechanisms for credit risk, fair lending, and financial stability. *Journal of Economics, Finance and Accounting Studies*, 5(4), 80-93. <https://doi.org/10.32996/jefas.2023.5.4.8>
- [20]. Rasel, I. H., Ibrahim, M., Pritty, A. A., Fahim, A. S. M., & Jahan, N. (2023). Beyond FICO: Enhancing mortgage default forecasting and inclusive lending via multimodal graph neural networks and urban mobility analytics. *Frontiers in Computer Science and Artificial Intelligence*, 2(2), 62-81. <https://doi.org/10.32996/fcsai.2023.2.2.5>
- [21]. Arman, M., & Fahim, A. S. M. (2023). AI revolutionizes inventory management at retail giants: Examining Walmart's U.S. operations. *Journal of Business and Management Studies*, 5(6), 145-148. <https://doi.org/10.32996/jbms.2023.5.6.15>
- [22]. Fahim, A. S. M., Pritty, A. A., Ibrahim, M., & Tania, T. A. (2024). Real-time payments and real-time fraud: A U.S. FinTech risk framework for RTP rails and consumer protection. *Journal of Economics, Finance and Accounting Studies*, 6(6), 134-149. <https://doi.org/10.32996/jefas.2024.6.6.11>

- [23]. Pritty, A. A., Ibrahim, M., Fahim, A. S. M., & Zadid, M. U. (2024). Generative AI and U.S. financial reporting integrity: Detecting narrative manipulation, risk disclosure gaming, and fraud signals in 10-K filings. *Journal of Economics, Finance and Accounting Studies*, 6(4), 113-129. <https://doi.org/10.32996/jefas.2024.6.4.11>
- [24]. Ibrahim, M., Mahmud, S., Zadid, M. U., Jahan, N., Rahman, M. M., & Fahim, A. S. M. (2024). AI-driven predictive analytics framework for anti-money laundering risk management and financial infrastructure protection in U.S. banking systems. *Journal of Economics, Finance and Accounting Studies*, 6(1), 155-166. <https://doi.org/10.32996/jefas.2024.6.6.12>
- [25]. Jahan, N., Pritty, A. A., Ibrahim, M., Zadid, M. U., Fahim, A. S. M., & Mahmud, S. (2024). Machine learning-driven early warning analytics for identifying market manipulation, irregular trading activity, and suspicious market signals in U.S. stock markets. *Journal of Computer Science and Technology Studies*, 6(2), 257-283. <https://doi.org/10.32996/jcsts.2024.6.2.26>
- [26]. Arman, M., Hasan, M. N., & Rasel, I. H. (2024). Clean energy transition in USA: Big data analytics for renewable energy forecasting and carbon reduction. *Journal of Management World*, 2024(3), 192-206. <https://doi.org/10.53935/jomw.v2024i4.1196>
- [27]. Arman, M., Rasel, I. H., Razib, M. N. H., & Fahim, A. S. M. (2024). Big data and machine learning for sustainable waste reduction. *Journal of Posthumanism*, 4(2), 448-467. <https://doi.org/10.63332/joph.v4i2.3361>
- [28]. Shah, A., Kabir, S., Razib, M. N. H., & Khan, S. A. (2024). Grid-integrated EV charging systems: Impacts on U.S. power grid stability and resilience. *International Journal of Artificial Intelligence Engineering and Transformation*, 5(1), 40-45. <https://doi.org/10.54660/IJAJET.2024.5.1.40-45>
- [29]. Razib, M., Ibrahim, M., & Rasel, I. (2025). Predictive analytics and its role in optimizing sustainable supply chain performance. *International Journal of Business Management*, 8(01), 12-27. <http://doi.org/10.35409/IJBMER.2025.3640>
- [30]. Ibrahim, M., Razib, M. N. H., & Rasel, I. H. (2025). The role of data analytics in enhancing ESG transparency in the corporate sector of Bangladesh. *Global Journal of Engineering and Technology Advances*, 22(1), 81-93. <https://doi.org/10.30574/gjeta.2025.22.1.0245>

Appendix A. Comparative Case Profiles

Appendix A provides concise case profiles for the representative systems used in the scorecard. Bitcoin is treated as a benchmark public blockchain for secure, censorship-resistant value transfer. Its principal strengths are broad recognizability, strong public auditability, and a design history that has made it the reference case for decentralized monetary settlement. Its principal limitations in financial-infrastructure evaluation are slower confirmation cadence, probabilistic finality, and limited native programmability relative to more recent smart-contract platforms. For institutional use, Bitcoin is most relevant where public verifiability and high-value transfer matter more than execution complexity or sub-minute settlement requirements.

Ethereum is treated as the benchmark programmable public settlement layer. Its post-merge proof-of-stake timing structure provides more regular block production than legacy proof-of-work models, and its smart-contract ecosystem enables tokenization, automated escrow logic, collateral applications, and application-layer financial services. These features make Ethereum highly relevant for tokenized assets and composable finance. At the same time, Ethereum's performance is shaped by gas-market dynamics, application complexity, and dependence on broader ecosystem infrastructure including bridges, rollups, and off-chain service providers. From an evaluation perspective, Ethereum performs best in use cases where programmability and ecosystem depth justify potentially variable execution costs.

Solana is included as an example of a high-performance public blockchain that explicitly optimizes for speed and low base transaction costs. This architecture makes it attractive for payment-like workloads and consumer-scale digital asset applications. In the context of financial systems, Solana demonstrates that public-chain design can move meaningfully closer to commercial throughput expectations than earlier blockchain generations. However, its performance must be interpreted alongside governance and resilience questions, including validator economics, hardware intensity, and operational concentration. These are not reasons to dismiss the platform, but they are reasons to treat low fees and high speed as only part of the performance story.

Visa is included as a non-blockchain benchmark because mature financial evaluation requires comparison against industrialized incumbent infrastructure. Visa's network demonstrates the scale, uptime, acceptance, and process maturity that blockchain systems would need to approach for universal retail use. At the same time, Visa does not provide the same kind of shared public ledger, atomic cross-ledger settlement logic, or open composability that blockchain systems can provide. Its inclusion therefore clarifies the central thesis of this paper: blockchain systems should not be judged only against one another, but also against the best available non-blockchain alternatives in the use case under consideration.

Project Cedar and Cedar x Ubin+ represent permissioned institutional distributed-ledger experimentation targeted at wholesale and cross-border settlement. These cases are crucial because they move the discussion from theory to institution-led empirical evidence. The reported settlement times indicate that distributed ledgers can materially reduce latency and improve synchronization of multi-party settlement across currencies. Their importance is not that they prove a single production model, but that they demonstrate a feasible design space in which distributed ledgers can address inefficiencies in financial market infrastructure.

mBridge is treated as a hybrid public-policy and infrastructure case. It is designed to improve cross-border payments through a distributed-ledger architecture used by central banks and commercial banks. Its relevance to this paper lies in showing how blockchain-derived design can be used in a controlled, institutionally governed environment rather than only in open public networks. For many policymakers and regulated institutions, this may be the most realistic adoption path: not open replacement of the existing system, but targeted redesign of settlement architecture around interoperable tokenized money and shared records.

Appendix B. Practical Evaluation Checklist for Institutions

Appendix B summarizes a practical evaluation checklist that institutions can use before adopting or investing in blockchain-based financial systems. First, the institution should define the exact friction it is trying to solve. Is the problem settlement delay, reconciliation cost, lack of shared records, transparency failure, expensive intermediation, or limited programmability? If the problem is merely internal workflow automation within a single trusted organization, a conventional database may be superior. Blockchain becomes more credible when multiple entities need a synchronized record and no single party should unilaterally control transaction state.

Second, institutions should establish a latency threshold that is economically meaningful for the use case. Retail checkout, treasury settlement, wholesale collateral movement, remittance transfer, and securities servicing all have different speed requirements. A blockchain that is fast in absolute terms may still be too slow relative to user expectations or risk tolerance in a specific process. Conversely, a blockchain does not need microsecond performance to add major value if it meaningfully compresses settlement from days to seconds in an existing workflow.

Third, institutions should measure cost in a full-stack way. They should not compare only visible network fees. They should estimate onboarding cost, custody design, compliance tooling, fraud monitoring, liquidity management, integration work, audit procedures, legal review, and contingency planning. This broader cost lens is especially important because many blockchain projects appear inexpensive in pilot settings but become more complex once regulatory and operational controls are imposed.

Fourth, governance must be reviewed as carefully as software. Institutions should map who can change protocol rules, who operates validators or infrastructure dependencies, how incidents are handled, and what legal remedies exist when something goes wrong. A technically advanced system with ambiguous governance can produce unacceptable control risk for regulated finance. Permissioned systems often score better here, but they must still be tested for role clarity, operational accountability, and continuity planning.

Fifth, transparency should be assessed in terms of usefulness rather than symbolism. Public visibility is valuable only if the institution can connect transaction records to relevant entities, permissions, and compliance workflows. Privacy-preserving architectures may sometimes produce more useful institutional transparency than fully open ledgers if they offer authenticated access, better data protection, and stronger governance.

Finally, institutions should test interoperability and exit options. Can the system connect with existing enterprise resource planning systems, payment messaging standards, reporting platforms, custodians, and identity providers? Can positions or records be migrated if the chosen platform underperforms or governance changes? Interoperability and reversibility are often undervalued during innovation phases, but they are critical for long-term resilience in global financial markets.

Appendix C. Variable Definitions and Measurement Notes

Appendix C lists the principal variables used in the comparative framework. Settlement latency is the elapsed time from valid submission to economically meaningful settlement. Confirmation cadence is the protocol or system rhythm at which shared state updates, such as a block interval or slot schedule. Finality refers to the degree of certainty that a settled transaction will not be reversed without extraordinary intervention. Auditability captures whether a transaction trail can be independently reconstructed and verified by authorized parties. Cost burden combines direct execution cost with reconciliation, compliance, and liquidity costs. Congestion sensitivity captures the degree to which delay or fees increase under heavy demand. Governance concentration refers to the extent to which a small set of actors can shape system outcomes through validation power, software control, or informal influence. Interoperability captures the practical ability of the system to connect with regulated payment, custody, messaging, and reporting environments.

These variables can be adapted across financial contexts. In wholesale payments, latency and finality may deserve heavier weights. In consumer payments, cost, acceptance, and fraud controls may dominate. In tokenized securities, programmability, auditability, and delivery-versus-payment logic become central. The advantage of this framework is therefore not that it forces every use case into one rigid ranking, but that it standardizes the categories of analysis. Researchers can modify weights while preserving conceptual comparability. Policymakers can emphasize governance and resilience. Firms can emphasize execution cost and integration effort. Investors can emphasize adoption scalability and control risk.

A final methodological note is that financial infrastructure should be evaluated dynamically. The same blockchain system can change materially as fee markets evolve, governance shifts, or institutional interfaces improve. Likewise, incumbent

systems also change through faster payments, tokenization, and API-based interoperability. The appropriate benchmark is therefore not a static contest between old and new technologies, but a moving evaluation of which architecture best solves identifiable market frictions at a given point in time.

Appendix D. Reading the Scorecard in Practice

Appendix D offers a short interpretation guide for practitioners. A high score on efficiency should not be read as a recommendation for immediate production adoption. It should instead be read as evidence that the architecture has promising transaction-processing characteristics for the specified use case. Production suitability additionally depends on legal enforceability, supervision, service-level commitments, incident response design, customer protection, and operational governance. Likewise, a lower transparency score for a permissioned system does not necessarily indicate poor oversight quality; in regulated environments, selective transparency can be preferable to universal visibility if it better protects confidential transaction data while preserving robust audit access for authorized parties.

The scorecard is therefore best understood as a decision-support tool. Financial institutions can use it to screen candidate infrastructures, compare pilot designs, and identify which trade-offs deserve deeper due diligence. Regulators can use it to structure supervisory dialogue with firms proposing tokenized payment or settlement arrangements. Researchers can use it to standardize empirical comparison across evolving systems without collapsing important governance and market-context differences. In all three settings, the framework encourages a disciplined question: compared with the best available alternative, what specific friction is this blockchain architecture reducing, and what new risks is it creating?