



Regulatory-Grade Evidence for AI in FinTech: A Control-Mapped Framework and Blockchain-Anchored Architecture for Audit Replay

Ian T. Staley^{1,*}

¹Independent Researcher, USA

Abstract: Financial institutions increasingly deploy artificial intelligence to support high-stakes operational decisions in anti-money laundering investigations and transaction dispute resolution. However, many deployments fail during audits and regulatory reviews not because models underperform but because decision processes cannot be reliably reconstructed, defended, or proven untampered over time. This paper addresses that gap by proposing a regulatory-grade evidence framework for artificial intelligence (AI) in financial technology (FinTech) that translates governance and documentation expectations into concrete, system-level evidence requirements. Using a document-driven control-mapping methodology, this study synthesizes requirements from the National Institute of Standards and Technology (NIST) Artificial Intelligence Risk Management Framework and the European Union Artificial Intelligence Act into an evidence control taxonomy spanning decision traceability, evidence integrity, audit replay, and retention governance. Two reference architectures are evaluated: a centralized AI governance and logging stack, and the same stack augmented with selective blockchain anchoring to provide tamper-evident integrity guarantees. The analysis identifies evidentiary coverage gaps, architectural trade-offs, and bounded conditions under which blockchain anchoring provides material incremental compliance value. This research contributes an actionable evidence control taxonomy, a controls-to-artifacts traceability matrix, and a Minimum Viable Evidence Layer (MVEL) architecture that enables audit-ready “decision replay” for regulated AI workflows, positioning AI governance in finance as an evidence system design problem.

Keywords: AI governance, regulatory-grade evidence, FinTech compliance, audit replay, Minimum Viable Evidence Layer (MVEL)

1. Introduction

Artificial intelligence (AI) is increasingly embedded in regulated financial technology (FinTech) workflows, including anti-money laundering (AML) investigations, transaction dispute resolution, and credit decision support. In these contexts, AI systems influence outcomes with legal, financial, and reputational consequences, placing them under regulatory scrutiny, supervisory review, and post-hoc audit. As a result, the primary challenge facing enterprise AI adoption in FinTech extends beyond model accuracy to the ability to reconstruct, justify, and defend AI-assisted decisions over time [1].

Recent governance frameworks emphasize that trustworthy AI requires robust organizational and technical controls encompassing documentation, traceability, accountability, and lifecycle oversight. The National Institute of Standards and Technology (NIST) Artificial Intelligence Risk Management Framework articulates these expectations through governance, mapping, measurement, and management functions that must be operationalized across AI systems [1]. In parallel, the European Union Artificial Intelligence Act establishes binding obligations for many high-risk AI systems, including requirements for record-keeping, transparency, post-market monitoring, and human oversight [2]. Together, these frameworks shift regulatory focus from model behavior alone toward the preservation and defensibility of decision evidence.

Despite these developments, many FinTech organizations continue to rely on fragmented evidence practices, such as centralized logs, model registries, and monitoring dashboards, that were designed for operational observability rather than long-term evidentiary integrity. Over time, model updates, data drift, system migrations, and human interventions can erode provenance and complicate the reconstruction of past decisions. This gap between governance expectations and operational evidence capabilities has emerged as a critical barrier to scaling AI in regulated financial environments [3, 4].

To address this gap, this study proposes a regulatory-grade evidence framework for AI-assisted FinTech decisioning. The framework translates governance and regulatory expectations into a concrete set of evidence controls and artifacts and evaluates architectural approaches for preserving evidentiary integrity over time. In doing so, this study adopts a deliberately bounded view of blockchain technology, treating it not as a platform for AI computation but as a mechanism for anchoring selected decision evidence in a tamper-evident and verifiable manner when such guarantees are operationally justified [5, 6].

Although AI governance frameworks increasingly mandate documentation, traceability, and accountability, FinTech organizations lack a clear operational model for implementing regulatory-grade evidence systems. Existing governance tooling prioritizes real-time monitoring and lifecycle management but often fails to ensure that decision evidence remains immutable, complete, and defensible across extended retention horizons. Moreover, while blockchain is frequently proposed as a solution to AI governance challenges, its application

*Corresponding author: Ian T. Staley, Independent Researcher, USA. Email: ian.t.staley@gmail.com

in FinTech contexts is often poorly scoped, introducing architectural complexity without a clear mapping to regulatory requirements [5, 6].

Against this backdrop, this study advances an evidence-centric approach to AI governance in FinTech. It synthesizes regulatory and governance requirements into a structured evidence control taxonomy, evaluates centralized and blockchain-anchored reference architectures, and delineates bounded conditions under which blockchain anchoring provides material incremental compliance value. In doing so, this paper reframes AI governance in finance as an evidence system design problem and provides practical guidance for achieving audit-ready “decision replay” under emerging regulatory regimes [3, 4]. Shifting AI governance from model-centric controls to evidence sufficiency and audit replay fills a critical void between regulatory demands and the technical reality of financial AI systems.

2. Literature Review

Recent scholarship on artificial intelligence governance highlights a growing disconnect between high-level governance principles and the operational mechanisms required to demonstrate accountability in practice. Systematic reviews emphasize documentation, traceability, human oversight, and lifecycle management as core elements of trustworthy AI, particularly in regulated domains such as financial services [3, 4]. However, these studies consistently note that governance failures often arise not from deficient ethical intent or model performance but from insufficient evidence practices that prevent organizations from reconstructing and defending past decisions. This perspective aligns with emerging scholarship that frames AI governance as the convergence of risk management, IT governance, and accountability infrastructures across socio-technical systems [7].

In FinTech and banking contexts, this challenge is amplified by longstanding regulatory expectations for record-keeping, auditability, and supervisory review. Studies examining AI adoption in financial services identify decision provenance and evidentiary durability as recurring compliance pain points, particularly as systems evolve through model updates, data drift, and organizational change [8–10]. As AI systems increasingly support AML investigations, transaction dispute resolution, and credit decisioning, the evidentiary burden associated with AI-assisted decisions approaches that of traditional financial records.

A growing body of literature explores the integration of blockchain technologies with AI systems as a means of enhancing integrity, provenance, and accountability. Surveys and conceptual studies argue that blockchain can provide tamper-evident logging, non-repudiation, and version tracking for AI-generated artifacts [5, 6, 11]. Applied studies further suggest potential efficiency and cost benefits when AI and blockchain are jointly deployed in financial systems [12]. However, much of this literature remains generalized, often proposing broad integration architectures without clearly identifying which regulatory requirements are addressed or which evidence artifacts benefit most from anchoring. Recent studies also emphasize blockchain’s role in improving AI transparency, explainability governance, and secure decision provenance in financial and decentralized finance systems [11, 13–15]. Parallel research on AI agent architectures highlights the growing importance of agent accountability, version provenance, and system-level governance in financial decision automation [16–19].

More recent work adopts a more critical and bounded perspective, emphasizing that blockchain should be evaluated as a governance instrument rather than a computational substrate for AI. These studies argue that blockchain’s value lies in selectively strengthening evidentiary integrity and auditability while cautioning against unselective integration that introduces complexity without proportional compliance benefit [7, 20, 21]. Collectively, the literature

points to the need for an evidence-centric framework that translates governance expectations into concrete, verifiable system artifacts that are capable of supporting regulatory review and audit replay.

Table 1 organizes these obligations into a four-domain evidence control taxonomy—decision traceability, evidence integrity, audit replay, and retention and access governance—comprising twelve controls (C1–C12). Table 2 then identifies the minimum evidence artifacts required to operationalize each control in a replay-capable AI governance environment.

Table 1
Regulatory-grade evidence control taxonomy for AI-assisted FinTech systems

Control domain	Control ID	Control description
Decision traceability	C1	Preserve a complete lineage linking each decision to input data sources, derived features, model identifiers, and execution context.
	C2	Preserve decision rationale artifacts (e.g., explanation outputs, rules triggered, and confidence scores) associated with each decision.
	C3	Preserve records of human actions, including approvals, overrides, escalations, and exceptions.
Evidence integrity	C4	Ensure decision evidence is tamper-evident, such that any modification or deletion can be reliably detected.
	C5	Ensure non-repudiation of key decision artifacts through cryptographic signing or equivalent mechanisms.
	C6	Preserve reliable ordering and timestamping of evidence events across systems.
Audit replay	C7	Enable reconstruction of the full decision context, including data, model, policy, and governance state.
	C8	Preserve version provenance for models, data schemas, prompts, and policies applicable at decision time.
	C9	Support export of evidence in usable, regulator-acceptable formats within reasonable timeframes.
Retention & access governance	C10	Retain decision evidence in accordance with regulatory and organizational retention requirements.
	C11	Enforce access controls and segregation of duties for evidence systems, with logged access events.
	C12	Monitor evidence completeness and integrity, detecting gaps, anomalies, or control failures.

2.1. Theoretical framework

This research is grounded in an evidence-centric interpretation of AI governance, drawing on risk management theory, accountability theory, and socio-technical systems perspectives. Contemporary AI governance frameworks conceptualize trustworthy AI as the outcome of organizational controls, technical safeguards, and oversight mechanisms operating across the AI lifecycle [3, 4]. Within this framing, accountability is realized not only through transparency or

Table 2
Mapping of evidence controls to required artifacts

Control ID	Evidence control	Minimum required artifact(s)
C1	Decision lineage	Input data reference, feature/derived-data reference, model identifier, decision ID, timestamp
C2	Rationale capture	Explanation output, rule trigger, confidence score, prompt/parameter snapshot where applicable
C3	Human accountability	Approver/analyst identity, action taken, override/escalation record, timestamp, justification note
C4	Evidence integrity	Cryptographic hash, digital signature or attestation, immutable event log entry
C5	Non-repudiation	Signed approval, identity-bound attestation, provenance record linking actor to action
C6	Ordering and timestamping	Sequenced event record, trusted timestamp, append-only log or anchored commitment
C7	Audit replay readiness	Reconstructable evidence bundle containing lineage, rationale, human actions, and decision outcome
C8	Version provenance	Model version, policy/rule version, prompt version if applicable, schema/version metadata, approval state
C9	Evidence bundling	Bundle ID, artifact manifest, decision packet structure, linkage across related records
C10	Retention governance	Retention schedule, storage class, access policy, deletion/disposal rule
C11	Access-controlled retrieval	Role-based access log, retrieval authorization record, query/audit access event
C12	Supervisory exportability	Exportable audit package, machine-readable metadata, regulator-review file set

explainability but also through the ability to demonstrate, after the fact, that decisions were made in accordance with defined policies, controls, and regulatory requirements.

The NIST Artificial Intelligence Risk Management Framework provides a foundational risk-based structure by articulating governance, mapping, measurement, and management functions that must be operationalized throughout AI system deployment [1]. These functions implicitly require the preservation of decision evidence sufficient to support traceability, monitoring, and post-deployment evaluation. Similarly, the European Union Artificial Intelligence Act formalizes accountability expectations for high-risk AI systems through explicit obligations related to record-keeping, technical documentation, human oversight, and post-market monitoring [2]. Together, these frameworks position evidence preservation as a central mechanism through which AI accountability is operationalized in regulated environments.

From a theoretical standpoint, this study conceptualizes regulatory compliance as a problem of evidence sufficiency over time. Rather than treating governance as a static set of policies or controls, the framework emphasizes the dynamic preservation of decision artifacts that enable audit replay, defined as the reconstruction and validation of a past decision using trustworthy and verifiable evidence.

This perspective aligns with socio-technical governance scholarship that frames accountability as emerging from the interaction between technical architectures, organizational processes, and regulatory oversight [21, 22].

Blockchain technology is incorporated into this framework in a deliberately bounded manner. Consistent with recent critiques of overgeneralized AI-blockchain integration, blockchain is treated as a specialized integrity mechanism capable of strengthening non-repudiation, ordering, and tamper resistance for selected evidence artifacts, rather than as a universal solution for AI governance [5, 6, 20]. This positioning enables a proportional evaluation of blockchain’s contribution relative to centralized governance mechanisms and provides a theoretical basis for identifying conditions under which blockchain anchoring yields material incremental compliance value.

3. Research Methodology

This research adopts a design-oriented qualitative research methodology to examine how regulatory-grade evidence can be operationalized for AI-assisted decision systems in regulated FinTech environments. Rather than relying on empirical surveys, interviews, or institutional case studies, this research focuses on the systematic analysis of regulatory texts, governance frameworks, and peer-reviewed scholarship to derive implementable evidence controls and architectural design patterns. In this study, audit replay refers to the ability to reconstruct and validate a historical AI-assisted decision using preserved, trustworthy evidence artifacts that capture the decision context, system state, and human oversight applied at the time the decision was made.

This methodological choice is appropriate given the regulatory sensitivity of AI governance in financial services and the limited accessibility of production-level AI systems subject to supervisory oversight. Prior AI governance research similarly emphasizes conceptual rigor, reproducibility, and architectural sufficiency over self-reported adoption data when examining compliance-oriented system design [3, 4]. The approach aligns with design science traditions in information systems and governance research, where normative requirements are translated into operational system constructs.

The methodology proceeds through four sequential analytical stages: (1) identification of evidence-relevant governance and regulatory requirements, (2) extraction and normalization of regulatory-grade evidence controls, (3) development of a structured evidence control taxonomy, and (4) comparative evaluation of alternative system architectures against the taxonomy. Together, these stages support a defensible assessment of how AI governance expectations can be satisfied through system design rather than policy declarations alone.

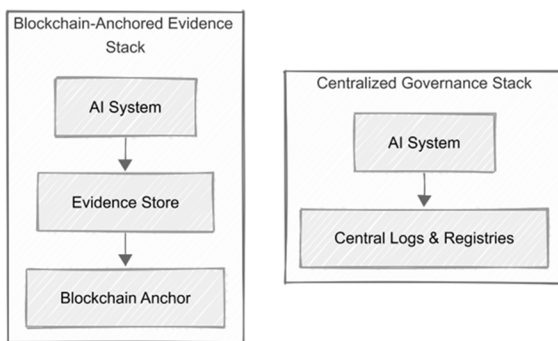
To improve methodological transparency and reproducibility, the control-mapping process followed a structured normalization logic. First, regulatory and governance texts were reviewed for clauses that imposed or implied evidentiary obligations relevant to AI-assisted decision systems, including documentation, lineage preservation, human oversight, logging, retention, integrity, and post-hoc reviewability. Second, each clause was translated into a functionally stated evidence requirement using technology-neutral language so that overlapping obligations across frameworks could be harmonized without duplicating controls. Third, normalized requirements were grouped by common evidentiary purpose and consolidated into the final control taxonomy. The resulting control IDs therefore represent synthesized evidence obligations derived from multiple normative sources rather than one-to-one restatements of individual clauses. Because this study is document-driven and single-author in design, the framework is intended to provide analytical transparency rather than statistical inter-rater validation. Future work may test the taxonomy through multi-reviewer coding and empirical supervisory simulations.

Blockchain is treated in this framework as a bounded governance mechanism for strengthening the integrity and non-repudiation of selected evidence artifacts, rather than as a general platform for AI computation or decision execution.

Figure 1 contrasts the two reference architectures used in the comparative analysis. The centralized architecture represents prevailing enterprise practice, in which logs, model records, approval workflows, and storage controls are managed within conventional governance systems. The blockchain-anchored architecture preserves the same operational stack but adds selective integrity anchoring for high-value evidence artifacts such as approvals, version provenance, and evidence-bundle commitments. The figure is intended to show that blockchain is introduced as a narrowly scoped integrity layer rather than as a replacement for core AI governance infrastructure.

Figure 1

Centralized vs. blockchain-anchored evidence architectures



3.1. Research design

The research design is document-driven and analytical in nature. Authoritative regulatory and governance sources are treated as primary normative inputs, while peer-reviewed academic literature provides contextual grounding and theoretical support. The central premise of the design is that regulatory accountability for AI systems is ultimately assessed through the availability, integrity, and reconstructability of decision evidence over time.

Core regulatory sources include the NIST Artificial Intelligence Risk Management Framework (AI RMF 1.0) and the European Union Artificial Intelligence Act, both of which articulate enforceable or de facto expectations related to documentation, traceability, logging, and post-deployment monitoring for AI systems deployed in high-risk contexts [1, 2]. These frameworks are increasingly referenced by financial regulators and supervisory bodies and therefore serve as appropriate anchors for deriving evidence requirements.

The source-selection logic was purposive rather than exhaustive. The NIST AI RMF and the European Union AI Act were selected because they are widely cited, operationally influential, and especially relevant to regulated, high-risk AI use cases involving documentation, accountability, and post-deployment oversight. Peer-reviewed literature was then used to contextualize these obligations within broader scholarship on AI governance, auditability, accountability, and blockchain-enabled integrity mechanisms. This bounded selection supports this study’s objective of deriving a practical evidence framework for regulated FinTech settings without claiming universal coverage of all international regulatory regimes.

The research design does not seek to measure adoption levels or organizational attitudes toward AI governance. Instead, it evaluates architectural sufficiency by asking whether a given system design can support audit replay, defined as the reconstruction and validation of a historical AI-assisted decision using preserved and verifiable evidence. This framing reflects supervisory priorities in AML investigations, transaction disputes, and credit oversight, where post-hoc defensibility often outweighs real-time model performance [3, 9].

To support comparative analysis, two reference architectures are defined: a centralized AI governance stack reflecting prevailing enterprise practice and an augmented architecture incorporating selective blockchain anchoring for evidentiary integrity. These architectures are evaluated qualitatively against a common evidence control taxonomy using sufficiency criteria rather than performance benchmarks. The comparative outcome of this evaluation, expressed as coverage levels (low/medium/high) for each control under the centralized and blockchain-anchored architectures, is summarized in Table 3. This design enables proportional assessment of blockchain’s incremental governance value without presupposing its necessity across all AI system components [5, 6].

Table 3
Control coverage matrix

Control ID	Evidence control	Centralized stack (A)	Blockchain-anchored (B)	Analytical commentary
C1	Decision lineage	Medium	High	Anchoring prevents silent modification of lineage records
C2	Rationale capture	Medium	Medium	Explanation quality unchanged by blockchain
C3	Human accountability	Medium	High	Signed attestations strengthen approval traceability
C4	Tamper evidence	Low–medium	High	Immutable commitments detect post-hoc edits
C5	Non-repudiation	Medium	High	Cryptographic signatures enforce accountability
C6	Time integrity	Medium	High	Ordered commitments strengthen event sequencing
C7	Audit replay	Medium	High*	*Dependent on completeness of captured artifacts
C8	Version provenance	Medium	High	Prevents overwriting or retroactive version changes
C9	Evidence exportability	High	High	Process-driven, not ledger-dependent
C10	Retention posture	High	High	Policy-based; blockchain supports proof-of-existence
C11	Access control	High	High	Remains IAM-driven in both architectures
C12	Monitoring & exceptions	Medium	Medium	Operational control independent of ledger

Note: *For control C7 (audit replay), the “high” coverage rating under the blockchain-anchored architecture is conditional upon the completeness of the captured evidence artifacts. Blockchain anchoring strengthens the integrity and non-repudiation of anchored commitments but does not compensate for missing lineage, rationale, or human-action records.

To translate the evidence control taxonomy into an implementable system design, this study introduces a Minimum Viable Evidence Layer (MVEL): the smallest set of architectural capabilities required to enable audit replay (i.e., reconstruction and validation of a historical AI-assisted decision) in regulated FinTech workflows. MVEL operationalizes governance expectations by ensuring that each decision generates a durable “evidence bundle” containing decision lineage, rationale outputs, and human-in-the-loop actions, preserved with verifiable integrity over the applicable retention horizon [1, 2]. Consistent with risk-based AI governance, MVEL emphasizes evidence sufficiency rather than explainability alone, recognizing that supervisory review typically demands reconstructable decision context, version provenance, and defensible record-keeping for high-impact workflows such as AML investigations and transaction dispute resolution [3, 4].

MVEL comprises six baseline capabilities aligned to the control domains in Table 1 and the artifact mapping in Table 2: (1) evidence capture (C1–C3) at decision time across data inputs, model identifiers, and human actions; (2) evidence normalization and bundling to produce consistent, replayable records (C7–C9); (3) versioned governance state for models, prompts/policies, schemas, and approvals (C8); (4) evidence store and retention governance enforcing access control and retention posture (C10–C12); (5) integrity services providing cryptographic signing, ordering, and tamper-evidence for integrity-critical artifacts (C4–C6); and (6) audit replay interface enabling timely export of regulator-ready evidence packages (C7–C9) [1, 2]. Where audit contexts are adversarial or retention horizons are long, selective blockchain anchoring can be applied to integrity-critical commitments (e.g., approvals, attestations, and version provenance) to strengthen non-repudiation and post-hoc defensibility without relocating AI computation onto a ledger [5, 6, 23].

Figure 2 illustrates MVEL as a control-aligned evidence subsystem layered onto conventional AI governance stacks. The

architecture begins at decision time, where inputs, model identifiers, rationale outputs, and human interventions are captured and normalized into a replayable evidence bundle. That bundle is then stored under retention and access-control rules designed to support later supervisory retrieval. Integrity services generate hashes, timestamps, signatures, or comparable commitments over selected artifacts so that critical records can be verified at a later date. In the blockchain-anchored variant, only selected commitments are anchored to provide tamper-evident integrity guarantees for long-lived, governance-critical artifacts such as approvals, model-version provenance, and finalized evidence manifests, while operational processing remains off-chain [5, 6]. Table 4 illustrates how these MVEL capabilities map to representative high-risk FinTech workflows—AML investigations, transaction disputes, and optional credit review—by identifying the high-risk decisions, critical controls, evidence artifacts warranting anchoring, and likely audit replay triggers for each workflow. This figure therefore depicts MVEL not as a new decision engine but as the minimum evidence infrastructure required to support audit-ready replay in regulated FinTech environments.

4. Conclusion

Overall, this study examined how artificial intelligence-assisted FinTech decision systems can satisfy emerging regulatory expectations for traceability, accountability, and auditability under frameworks such as the NIST Artificial Intelligence Risk Management Framework and the European Union Artificial Intelligence Act [1, 2]. Through a document-driven control-mapping methodology and comparative architectural evaluation, the research demonstrated that regulatory-grade AI governance in finance is fundamentally an evidence system design problem rather than a model performance challenge.

This study introduced a regulatory-grade evidence control taxonomy that translates abstract governance requirements into implementable system controls centered on audit replay. By evaluating centralized AI governance architectures against blockchain-anchored evidence architectures, the analysis showed that while centralized systems remain necessary for operational oversight, they exhibit structural limitations in preserving long-lived, tamper-evident decision evidence. Selective blockchain anchoring provides material incremental value for integrity-critical artifacts such as approvals, version provenance, and decision attestations while offering limited benefit for real-time monitoring or explainability generation [5, 6, 23].

This study is subject to several limitations. First, the framework is derived from document-driven analysis of regulatory texts and peer-reviewed literature rather than empirical deployment within a production financial institution. As such, the findings evaluate architectural sufficiency for regulatory-grade evidence rather than operational performance or organizational adoption outcomes. Second, while the proposed control taxonomy and MVEL architecture are grounded in widely referenced governance frameworks, they do not assert legal sufficiency or substitute for jurisdiction-specific regulatory interpretation.

An additional limitation concerns jurisdictional scalability and regulatory change. Although the framework is grounded in widely

Figure 2
MVEL architecture

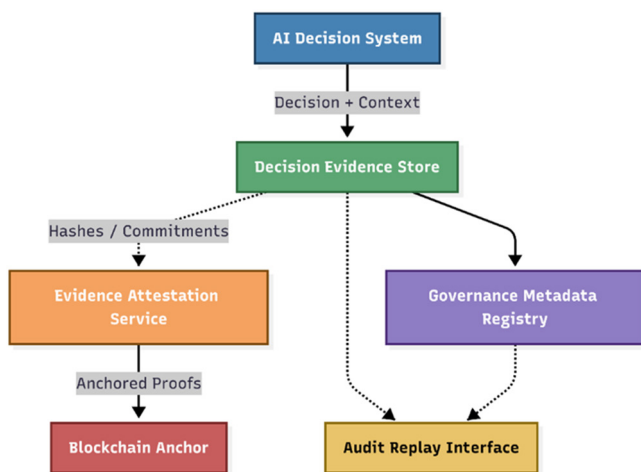


Table 4
Evidence requirements by FinTech workflow

Workflow	High-risk decisions	Critical controls	Evidence requiring anchoring	Audit replay trigger
AML investigations	Alert disposition, escalation, SAR recommendation	C1, C3, C4, C5, C7, C8	Approvals, lineage, version provenance	Regulatory exam, enforcement inquiry
Transaction disputes	Dispute acceptance or denial	C1, C4, C5, C6, C7	Evidence bundles, timestamps	Legal dispute, customer appeal
Credit review (optional)	Override of model output	C1, C2, C3, C8	Override rationale, model version	Adverse action challenge

referenced governance sources, regulatory expectations for AI documentation, retention, explainability, and model oversight continue to evolve across jurisdictions. As a result, the control taxonomy should be interpreted as a baseline evidence framework rather than a static or globally exhaustive compliance checklist. Institutions adopting the model may need to extend or recalibrate specific controls to reflect local supervisory expectations, sector-specific obligations, or future amendments to AI governance regimes.

This study also does not empirically quantify the latency, throughput, or cost implications of selective blockchain anchoring. That omission is important because high-frequency FinTech environments may be sensitive to even modest additional coordination or integrity verification overhead. For that reason, the framework advances a deliberately bounded design: blockchain anchoring is recommended only for integrity-critical artifacts with longer retention horizons or elevated audit risk, while primary decision processing, storage, and operational governance remain within conventional enterprise systems. Future implementation studies should evaluate these trade-offs under realistic workload conditions.

Future research may extend this work through empirical validation in live or simulated regulatory examinations, longitudinal analysis of evidence durability under system evolution, or comparative assessment of evidence architectures across jurisdictions. Additional work could also explore how MVEL interacts with emerging AI agent systems, model orchestration frameworks, and automated compliance tooling in complex financial environments. In particular, non-deterministic or agentic workflows may require expanded evidence capture for intermediate tool calls, orchestration state, delegated actions, and dynamic prompt chains to preserve replayability and supervisory defensibility.

Recommendations

From a design and governance perspective, institutions may use the proposed framework as an implementation baseline for building replay-capable evidence systems and then selectively add blockchain anchoring only for integrity-critical artifacts where audit risk and retention horizons justify the incremental governance overhead [5, 6, 23]. On the basis of these findings, this study recommends that FinTech organizations adopt evidence-first design principles when deploying AI in regulated workflows. Rather than treating compliance as a documentation exercise, institutions should explicitly architect systems to preserve decision evidence with sufficient integrity, provenance, and longevity to support supervisory review, investigations, and dispute resolution. Blockchain anchoring should be applied proportionally and selectively, aligned with workflow risk profiles and regulatory exposure.

For regulators and supervisors, the framework provides a practical lens for evaluating AI governance claims. Oversight efforts may benefit from focusing on whether institutions can demonstrate audit replay through preserved and verifiable evidence artifacts, rather than relying solely on transparency or explainability assertions. This approach aligns with the increasing emphasis on post-market monitoring, documentation, and accountability under risk-based AI governance regimes [1, 2].

Ethical Statement

This study does not contain any studies with human or animal subjects performed by the author.

Conflicts of Interest

The author declares that he has no conflicts of interest to this work.

Data Availability Statement

Data sharing is not applicable to this article as no new data were created or analyzed in this study.

Author Contribution Statement

Ian T. Staley: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Resources, Data curation, Writing – original draft, Writing – review & editing, Visualization, Project administration.

References

- [1] Tabassi, E. (2023). *Artificial intelligence risk management framework (AI RMF 1.0) (NIST AI 100-1)*. National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.AI.100-1>
- [2] European Parliament & Council of the European Union. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>
- [3] Batool, A., Zowghi, D., & Bano, M. (2025). AI governance: A systematic literature review. *AI and Ethics*, 5, 3265–3279. <https://doi.org/10.1007/s43681-024-00653-w>
- [4] Papagiannidis, E., Mikalef, P., & Conboy, K. (2025). Responsible artificial intelligence governance: A review and research framework. *Journal of Strategic Information Systems*, 34(2), 101885. <https://doi.org/10.1016/j.jsis.2024.101885>
- [5] Ifedayo, A. E., Olugbade, D., & Hamid, S. (2025). Integrating artificial intelligence with blockchain: A literature review on opportunities, challenges, and applications. *Blockchain, Artificial Intelligence, and Future Research*, 1(1), 52–69. <https://doi.org/10.70211/bafr.v1i1.179>
- [6] Karim, M. M., Van, D. H., Khan, S., Qu, Q., & Kholodov, Y. (2025). AI agents meet blockchain: A survey on secure and scalable collaboration for multi-agents. *Future Internet*, 17(2), 57. <https://doi.org/10.3390/fi17020057>
- [7] Batte, B. (2025). Toward an integrated future: The convergence of AI, IT governance, blockchain, NLP, and deep learning. *IT Governance, Blockchain, NLP, and Deep Learning* (June 26, 2025). <https://doi.org/10.2139/ssrn.5325590>
- [8] Addula, S. R., Meduri, K., Nadella, G. S., & Gonaygunta, H. (2024). AI and blockchain in finance: Opportunities and challenges for the banking sector. *International Journal of Advanced Research in Computer and Communication Engineering*, 13(2), 184–190. <https://doi.org/10.17148/IJARCC.2024.13231>
- [9] Cheng, X., Du, A. M., Yan, C., & Goodell, J. W. (2025). Internal business process governance and external regulation: How does AI technology empower financial performance? *International Review of Financial Analysis*, 99, 103927. <https://doi.org/10.1016/j.irfa.2025.103927>
- [10] Fahlevi, M., Moeljadi, M., Aisjah, S., & Djazuli, A. (2023). Corporate governance in the digital age: A comprehensive review of blockchain, AI, and big data impacts, opportunities, and challenges. *E3S Web of Conferences*, 448, 02056. <https://doi.org/10.1051/e3sconf/202344802056>
- [11] Rane, N., Choudhary, S., & Rane, J. (2023). Blockchain and artificial intelligence (AI) integration for revolutionizing security and transparency in finance. Available at SSRN 4644253. <http://dx.doi.org/10.2139/ssrn.4644253>

- [12] Naeem, M., Khan, A., Rehman, A., Farooq, S., Mehboob, A., Abdali, A. S., & Ahmad, B. (2025). Does artificial intelligence with blockchain reduce the costs of the financial sector?. In I. Shah & N. Jhanjhi (Eds.), *Generative AI for web engineering models* (pp. 147–160). IGI Global Scientific Publishing. <https://doi.org/10.4018/979-8-3693-3703-5.ch007>
- [13] Akther, A., Arobee, A., Adnan, A. A., Auyon, O., Islam, A. S. M., & Akter, F. (2025). *Blockchain as a platform for artificial intelligence (AI) transparency*. arXiv. <https://doi.org/10.48550/arXiv.2503.08699>
- [14] Bihani, D., Ubamadu, B. C., & Daraojimba, A. I. (2025). Integrating blockchain with AI: A data-driven model for secure, scalable decentralized finance (DeFi) systems. *World Journal of Innovation and Modern Technology*, 9(5), 58–85. <https://doi.org/10.56201/wjimt.v9.no5.2025.pg58.85>
- [15] Chahar, S., Kaur, K., Kaswan, K. S., & Dhattewal, J. S. (2025). Explainable AI in blockchain system for decentralized governance. In *2025 International Conference on Pervasive Computational Technologies*, 725–729. <https://doi.org/10.1109/ICPCT64145.2025.10941600>
- [16] Hettiarachchi, I. (2025). The rise of generative AI agents in finance: Operational disruption and strategic evolution. *International Journal of Engineering Technology Research & Management*, 9(4), 447. <https://ijetrm.com/issues/files/Apr-2025-26-1745643054-APR65.pdf>
- [17] Joshi, S. (2025). Advancing innovation in financial stability: A comprehensive review of AI agent frameworks, challenges and applications. *World Journal of Advanced Engineering Technology and Sciences*, 14(2), 117–126.
- [18] Joshi, S. (2024). A literature review of Gen AI agents in financial applications: Models and implementations. <https://ssrn.com/abstract=5133985>
- [19] Rizinski, M., & Trajanov, D. (2025). AI agents in finance and FinTech: A scientific review of agent-based systems, applications, and future horizons. *Computers, Materials and Continua*, 86(1), 1–34. <http://dx.doi.org/10.32604/cmc.2025.069678>
- [20] Alzoubi, M. M. (2025). Investigating the synergy of blockchain and AI: Enhancing security, efficiency, and transparency. *Journal of Cyber Security Technology*, 9(3), 227–255.
- [21] Asif, R., Hassan, S. R., & Parr, G. (2023). Integrating a blockchain-based governance framework for responsible AI. *Future Internet*, 15(3), 97. <https://doi.org/10.3390/fi15030097>
- [22] Coeckelbergh, M. (2024). Artificial intelligence, the common good, and the democratic deficit in AI governance. *AI and Ethics*, 5, 1491–1497. <https://doi.org/10.1007/s43681-024-00492-9>
- [23] Saeed, U. F. (2025). Can artificial intelligence and blockchain condition governance mechanisms to restrict earnings management?. *SN Business & Economics*, 5(12), 206. <https://doi.org/10.1007/s43546-025-00973-x>

How to Cite: Staley, I. T. (2026). Regulatory-Grade Evidence for AI in FinTech: A Control-Mapped Framework and Blockchain-Anchored Architecture for Audit Replay. *Journal of Computational Law and Legal Technology*. <https://doi.org/10.47852/bonviewJCLLT62029281>