Journal of Computer Science and Technology Studies
ISSN: 2709-104X
DOI: 10.32996/jcsts
Journal Homepage: www.al-kindipublisher.com/index.php/jcsts

JCSTS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

Privacy-Preserving Chargeback Intelligence for Tokenized Payment Systems

Vimal Teja Manne
The University of Texas at Dallas, Texas, USA
Corresponding Author: Vimal Teja Manne, **E-mail**: vimalteja.m@gmail.com

| ABSTRACT

Chargeback handling in modern payment systems remains heavily dependent on post-transaction workflows that expose more payment data than necessary and often rely on PAN-linked retrieval paths. This paper presents a privacy-aware post-transaction architecture for tokenized payment sys- tems that combines chargeback intelligence, selective evidence disclosure, and audit-oriented work- flow design. Beyond system integration, the paper introduces a formal role-based selective-disclosure policy model, a token-lineage risk formulation for dispute scoring, and an evidence-quality framework that separates completeness, relevance, and verifiability. The proposed framework contains a token- aware chargeback intelligence layer, a privacy-preserving evidence layer, and a compliance-aware audit layer to support PAN-less dispute handling wherever feasible. Public transaction datasets are used as the empirical basis, with token-related post-transaction structures constructed through controlled augmentation because the source datasets do not natively contain tokenized dispute workflows. The evaluation uses PaySim as the base transaction backbone, the ULB/Worldline credit-card fraud benchmark for fraud-context support, and a chargeback-labeled dataset for dis- pute modeling. In the primary chargeback prediction experiment, the token-aware model achieved an AUC of 0.9317 and an F1-score of 0.6710, outperforming both a static rule baseline and a ba- sic machine-learning baseline. Additional analysis includes ablation testing, sensitivity analysis, repeated-trial workflow evaluation, and failure-oriented threat scenarios. The proposed workflow reduced the weighted privacy-exposure score by 79.1% relative to a PAN-oriented baseline while improving evidence quality at the cost of higher operational overhead. These findings suggest that token-aware post-transaction intelligence can improve dispute prediction and reduce sensitive-data exposure simultaneously. The paper contributes a practical architecture, formalized workflow logic, explicit metric definitions, and a stronger evaluation framework for privacy-preserving chargeback operations in tokenized payment systems.

| KEYWORDS

Payment tokenization, chargeback intelligence, privacy-preserving evidence, selective disclosure, to- kenized payment systems, dispute workflows

| ARTICLE INFORMATION

ACCEPTED: 01 June 2023 **PUBLISHED:** 30 June 2023 **DOI:** 10.32996/jcsts.2023.5.2.5

1. Introduction

Payment tokenization has become a foundational mechanism for reducing merchant exposure to sensitive account data during authorization and recurring transaction processing. By replacing primary account numbers, PANs, with controlled token representations, tokenized payment archi- tectures reduce direct handling of raw account identifiers and improve security posture across front- end payment flows. However, post-transaction operations remain comparatively underdeveloped. Chargeback handling, dispute evidence preparation, and related post-transaction workflows often continue to depend on PAN-linked or PAN-adjacent retrieval paths, creating unnecessary privacy exposure and operational inefficiency.

This problem becomes more significant in large-scale digital payment environments where multiple entities,

including merchants, acquirers, issuers, processors, and compliance teams, may need access to dispute-related evidence. In conventional workflows, evidence bundles frequently expose more payment-related information than necessary. Such workflows also tend to lack clear architectural boundaries between operational evidence needs and sensitive payment-data controls. As a result, tokenization often protects payment initiation more effectively than it protects downstream dispute handling.

This paper addresses that gap by proposing a privacy-preserving chargeback intelligence architecture for tokenized payment systems. The proposed approach combines token-aware transaction features for chargeback risk scoring, a privacy-preserving evidence workflow that assembles role-specific dispute evidence through selective disclosure, and explicit token-linked workflow layers with audit and compliance controls. The contribution is not a new payment network, a new cryptographic primitive, or a claim of production deployment. Instead, it is a systems-level integration that targets a concrete operational gap: how to organize chargeback prediction and dispute evidence handling so that sensitive-data exposure is reduced without weakening workflow traceability.

1.1 *Research Problem*

Current chargeback and dispute workflows in payment systems often expose more payment data than necessary, rely on PAN-linked or PAN-adjacent retrieval mechanisms, lack formal privacy-aware evidence orchestration, create operational friction across merchant, acquirer, and issuer boundaries, and provide weak architectural support for auditable selective disclosure.

1.2 *Research Questions*

This study addresses the following research questions:

1. Can chargeback risk be modeled effectively using tokenized transaction features instead of PAN- linked workflows?
2. Can dispute evidence be generated in a way that remains useful while minimizing sensitive-data exposure?
3. How much privacy-exposure reduction can be achieved compared with traditional PAN-based post-transaction flows?
4. What latency and operational overhead are introduced by a privacy-preserving chargeback work- flow?
5. Can tokenized, privacy-aware dispute handling remain aligned with tokenization and data-minimization principles?

1.3 *Main Contributions*

This paper makes five contributions.

First, it proposes a token-aware chargeback intelligence architecture for post-transaction payment systems, organized around transaction ingestion, token linkage, dispute scoring, selective evidence release, and compliance-aware audit control.

Second, it introduces a formal role-based selective-disclosure policy model that determines which field classes are visible to merchant operations, internal analysts, and compliance reviewers during dispute handling.

Third, it defines a token-lineage risk formulation and an evidence-quality framework that explicitly separate predictive utility from disclosure sufficiency, allowing privacy reduction and dispute support quality to be evaluated together.

Fourth, it develops a hybrid evaluation framework that combines public transaction datasets, chargeback-oriented data, controlled synthetic token workflow construction, repeated trials, ablation analysis, and sensitivity analysis.

Fifth, it provides a quantitative comparison between traditional PAN-based dispute operations and tokenized

privacy-aware workflows in terms of prediction quality, privacy exposure, evidence quality, and operational overhead.

2. *Background and Related Work*

Relevant prior work spans payment tokenization, fraud and chargeback prediction, privacy-preserving transaction evidence, and post-transaction payment operations.

PCI tokenization guidance established early that the security value of tokenization depends not only on substituting PAN with a token but also on segmentation, retrievability, and the surrounding operational environment. The PCI tokenization product security guidelines further emphasized token-vault separation, protected mapping relationships, and the importance of minimizing the number of systems able to recover or influence sensitive cardholder data. These observations are directly relevant to the present work because dispute workflows often become the place where sensitive retrieval paths re-enter an otherwise tokenized environment.

Fraud-detection research produced strong public datasets and benchmark methods, especially PaySim and the ULB/Worldline credit-card fraud dataset. Studies on undersampling, probability calibra- tion, feature engineering, and sequence-aware fraud detection show that transaction-level behavior can be used effectively for anomaly scoring and fraud prediction. However, these works primarily focus on fraud classification rather than on dispute-oriented post-transaction workflows. In particu- lar, they do not define how token-linked transaction representations should be used when assembling case evidence or limiting sensitive-data exposure.

A separate body of work in privacy and anonymous payments has shown that transaction systems can be designed to hide sensitive details while preserving verifiability. Classic privacy-oriented work highlighted the economic and technical value of limiting unnecessary financial-data revelation. Later work on anonymous payments and anonymous attestation showed that transactional systems can preserve useful proofs without exposing all underlying identifiers. Even so, those works do not address the practical merchant-acquirer-issuer workflow of chargebacks, representments, and evidence release in card-payment systems.

Payment-card-network literature and payment-economics work provide further context regarding dispute incentives, transaction externalities, and operational roles across merchants, issuers, and acquirers. While these studies are useful for understanding the broader ecosystem, they do not provide a system architecture for privacy-preserving chargeback intelligence. The closest adjacent literature therefore contributes pieces of the problem rather than the full combination. Tokenization standards explain how to minimize direct exposure to PAN. Fraud literature explains how to score risk from transaction signals. Privacy-oriented work explains why selective disclosure matters. Yet the literature remains thin on architecture-focused designs that combine token-aware chargeback intelligence, dispute-evidence orchestration, selective disclosure, auditability, and post-transaction workflow control in a single operational model.

2.1 *Novelty Positioning*

The novelty claimed here is intentionally narrow and architectural. The paper does not claim to invent tokenization, chargeback modeling, or selective disclosure as standalone concepts. Instead, it contributes a systems-level integration for a more specific operational problem that prior work does not combine directly: linking token-aware risk scoring with role-specific evidence generation and audit-aware control points for post-transaction payment disputes. The novelty therefore lies in the joint design of workflow layers, not in a new cryptographic primitive or a production card-network deployment.

2.2 *Closest Prior Gap*

The closest prior work still leaves four combined gaps unresolved. First, tokenization guidance explains how to reduce direct PAN exposure, but it does not provide a dispute-oriented architecture for downstream evidence release. Second, fraud and chargeback studies provide predictive baselines, but they do not define token-aware post-transaction workflow logic. Third, privacy-preserving transaction literature motivates selective disclosure, but it does not model merchant-acquirer-issuer evidence exchange in operational card-payment disputes. Fourth, prior studies rarely combine predictive scoring, privacy-aware evidence release, trust-boundary control, and auditability within one formalized workflow. The present paper targets that intersection directly.

3. *Proposed System Architecture*

The proposed architecture consists of five layers: transaction ingestion, tokenization and linkage, chargeback intelligence, privacy-preserving evidence, and compliance and audit. Together these layers support PAN-less dispute handling wherever feasible.

3.1 *Transaction Ingestion Layer*

This layer captures normalized transaction events, including transaction identifier, merchant iden- tifier, customer hash, amount, timestamp, payment channel, and fraud-related status fields. It provides the canonical event stream from which downstream tokenized post-transaction records are constructed.

3.2 *Tokenization and Linkage Layer*

This layer introduces token-aware structures such as network token, vault token, token reference identifier, token status, token-age metadata, and a PAR-like linkage reference. Its purpose is to preserve continuity across post-transaction workflows without exposing raw PAN.

3.3 *Chargeback Intelligence Layer*

This layer computes chargeback likelihood score, dispute escalation risk, evidence completeness score, and merchant win-loss support estimation. It operates on transaction features, behavioral recurrence signals, and token-linked continuity features.

3.4 *Privacy-Preserving Evidence Layer*

This layer constructs role-specific evidence bundles. Instead of exposing full transaction records, it generates masked metadata, token-linkage proof, timestamp continuity proof, merchant-side proof, and issuer-facing, acquirer-facing, and merchant-facing evidence views.

3.5 *Compliance and Audit Layer*

This layer records evidence access, PAN-exposure status, audit-log identifiers, role-based disclosure history, and compliance-oriented scope reduction proxies. The design does not assert that this layer eliminates PCI scope. Rather, it narrows routine exposure and makes the handling of token-linked records reviewable.

3.6 *Selective Disclosure Enforcement*

Selective disclosure is enforced through role templates rather than through ad hoc manual redac- tion. In the workflow used here, merchant operations receive a narrow packet containing masked metadata, case-level status, and limited evidence descriptors. Internal analysts receive a broader view that includes token lineage, repeated-event patterns, and timestamp continuity. A compliance or audit role receives the richest packet, including access-log identifiers and release history, but the workflow still assumes that any PAN-linked mapping authority remains outside these downstream views. The evidence layer is therefore privacy-preserving in a bounded operational sense: it reduces the number of identities and fields visible to routine participants and makes each release attributable to an explicit role.

3.7 *Formal Workflow and State Model*

The proposed workflow is represented as a directed state model over token-linked post-transaction events. Let T denote the set of transactions, D the set of dispute cases, E the set of evidence bundles, and R the

set of authorized roles. Each transaction $t \in T$ is associated with a token reference $\tau(t)$, a timestamp $\sigma(t)$, a feature vector x_t, and a dispute status s_t.

The workflow states are defined as:

S = {ingested, token-linked, risk-scored, dispute-opened, evidence-assembled, released, audited} The nominal transition path is:
ingested → token-linked → risk-scored → dispute-opened → evidence-assembled → released → audited

A transition is permitted only if all required linkage, policy, and audit constraints are satisfied. If any policy or integrity constraint fails, the case transitions to a blocked or exception state rather than continuing silently.

3.8 *Trust Boundaries*

The proposed architecture assumes three trust zones.

Zone 1 is the privileged token-resolution boundary, which contains any token vault, PAN-linked mapping service, or privileged card-resolution interface. This zone remains in scope for any card- data compliance obligations.

Zone 2 is the token-aware operational layer, which includes chargeback scoring, event lineage, se- lective evidence assembly, and workflow-level decision support. This zone is designed to operate without routine PAN retrieval.

Zone 3 is the role-restricted evidence-consumption layer, where merchant operations, internal an- alysts, and compliance reviewers receive policy-filtered evidence views. This layer is restricted to masked or token-linked representations and auditable release events.

The architectural claim of the paper is not that Zone 1 disappears, but that Zones 2 and 3 can operate on substantially narrower and better-controlled data views.

3.9 *Threat Model and Adversarial Scenarios*

The threat model assumes that adversaries may attempt one or more of the following: infer card- linked identity from repeated dispute records, obtain excessive evidence detail through over-broad role access, exploit repeated-event linkage to reconstruct sensitive history, or manipulate dispute submissions to force disclosure of unnecessary attributes.

Four adversarial scenarios are considered.

Scenario 1: repeated-token correlation attack, where an observer attempts to connect multiple dispute cases through recurring token-linked references.

Scenario 2: role-escalation misuse, where a lower-privilege operational actor attempts to access analyst-level or compliance-level evidence fields.

Scenario 3: evidence over-disclosure, where a workflow produces richer evidence than is necessary for a given dispute role.

Scenario 4: lineage inconsistency attack, where malformed or contradictory event relationships are introduced to degrade evidence reliability.

The proposed controls address these scenarios through role templates, explicit release logs, lineage validation, and bounded field visibility.

3.10 *Failure Handling*

If a transaction cannot be linked to a valid token reference, the case is marked unresolved and is excluded from policy-governed evidence release. If lineage continuity fails, the evidence bundle is downgraded to incomplete status and the workflow records an exception event. If a requested role exceeds its permitted disclosure class, evidence release is blocked and logged. If token-linked recurrence features are unavailable, the predictive model

falls back to the reduced baseline feature set, and the resulting prediction is marked as degraded-confidence output.

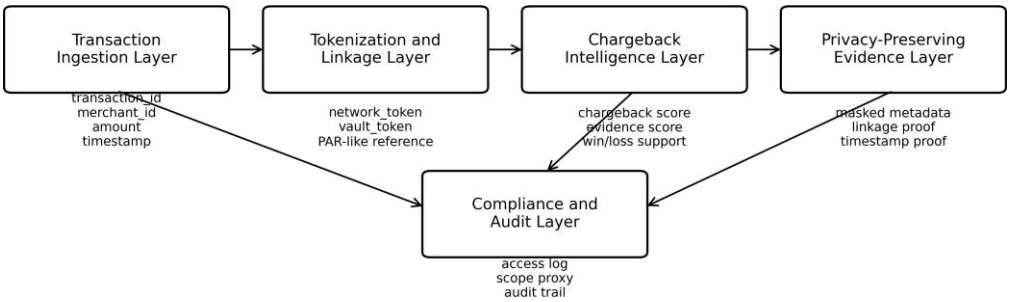

Figure 1: Architecture of the proposed privacy-preserving chargeback intelligence framework

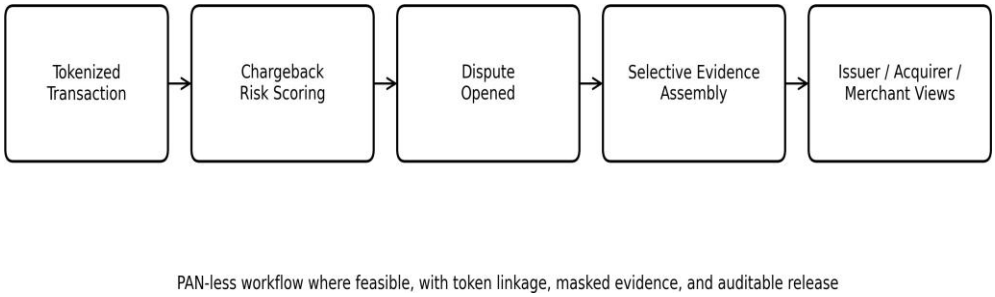

Figure 2: Token-aware dispute workflow with selective evidence disclosure

1. *Dataset Construction and Methodology*

The empirical study uses PaySim as the base transaction backbone, the ULB/Worldline credit-card fraud dataset as an auxiliary fraud benchmark, a chargeback-oriented dispute dataset for supervised dispute modeling, and additional fraud-context datasets for behavioral calibration. Because the public datasets do not contain native tokenized dispute workflows, token-related post-transaction structures are created through controlled augmentation.

3.11 *Canonical Schema*

The working schema contains transaction identifier, timestamp, amount, card-derived grouping fields, repeated-amount behavior, inter-transaction timing, token reference count, token age, and token rotation flag. The predictive label is derived from the chargeback flag in the chargeback- oriented dataset.

3.12 *Explicit Augmentation Rules*

To make the augmentation transparent, the workflow applies four explicit construction rules.

First, a token reference is created by combining a stable card-derived grouping with a coarse issuance counter, allowing repeated post-transaction activity to be represented without exposing full card number. Second, token age is measured as elapsed time since the first observed appearance of that token reference in the local dataset. Third, a token rotation flag is introduced every fifth repeated event for the same card-derived group, as a simple proxy for token lifecycle change. Fourth, a PAR- like repeat indicator is created from the count of card-linked observations so that dispute scoring can use recurrence without relying on PAN retrieval in downstream logic.

These constructions do not claim to reproduce issuer, network, or token-service-provider implemen- tations exactly. They are controlled operational proxies used to test whether token-aware post- transaction features

provide value beyond a simpler baseline.

3.13 *Train-Test Protocol*

The chargeback dataset is sorted chronologically. The first 70% of observations are used for training and the remaining 30% for testing. This choice preserves temporal direction and avoids leakage from future records into earlier model fitting.

3.14 *Chargeback Intelligence Experiment*

The predictive experiment compares three approaches:

- Static rule baseline
- Basic machine-learning baseline
- Token-aware chargeback intelligence model

The rule baseline uses repeated-amount and high-amount triggers. The basic machine-learning baseline uses time and amount features. The token-aware model adds token-linked and recurrence features such as token reference count, token age, repeat indicator, token rotation flag, amount recurrence count, and inter-transaction timing.

3.15 *Ablation and Sensitivity Design*

To test the contribution of token-aware features more rigorously, the evaluation includes ablation analysis. The full token-aware model is compared against reduced variants in which token age, token rotation flag, recurrence count, and inter-transaction timing are removed one at a time. This isolates the predictive contribution of each feature family.

Sensitivity analysis is also applied to the workflow metrics. Privacy-exposure weights are per- turbed under multiple weighting schemes to test whether the observed privacy reduction depends strongly on one specific weighting assignment. Evidence-quality weights are likewise perturbed across completeness-heavy, relevance-heavy, and verifiability-heavy configurations.

3.16 *Workflow Metrics*

Workflow-level evaluation measures privacy exposure, evidence quality, and operational overhead using explicit formulations.

The privacy-exposure score for a workflow instance i is:

$$P_i = \sum_{k=1}^{m} w_k z_{ik}$$

where $z_{ik} \in \{0, 1\}$ indicates whether sensitive field class k is exposed in workflow instance i, and w_k is the weight assigned to that class. The field classes include PAN-linked identifier, account-linked reference, mapping access, merchant metadata, timestamp continuity, lineage proof, and integrity metadata. The baseline score is 1.48, while the proposed workflow score is 0.31, corresponding to a reduction of 79.1%.

The evidence-quality score for case i is:

$$Q_i = \alpha C_i + \beta R_i + \gamma V_i$$

where C_i is completeness, R_i is relevance, and V_i is verifiability. The coefficients α, β, and γ are nonnegative and sum to 1. In the implementation used here, completeness measures the fraction of required evidence elements present, relevance measures role-appropriateness of disclosed fields, and verifiability measures whether the evidence bundle preserves traceable linkage and timestamp continuity.

Operational overhead is measured as:

$$O = \frac{N_{\text{workflow}}}{N_{\text{baseline}}}$$

where N_{workflow} is the number of workflow primitives executed in the proposed design and N_{baseline} is the number required in the PAN-oriented comparison flow.

These metrics are architecture-level constructs rather than claims of network-settlement perfor- mance or issuer-side acceptance.

4. *Experiments and Results*

4.1 *Chargeback Prediction Performance*

The predictive experiment shows that the token-aware model outperforms both comparison base- lines.

The static rule baseline achieved an AUC of 0.8213, an F1-score of 0.4829, precision of 0.4123, and recall of 0.5826.

The basic machine-learning baseline achieved an AUC of 0.6034, an F1-score of 0.1436, precision of 0.0780, and recall of 0.9130.

The token-aware model achieved an AUC of 0.9317, an F1-score of 0.6710, precision of 0.6638, and recall of 0.6783.

These results indicate that token-linked behavioral context can improve chargeback intelligence beyond both static rules and a simpler baseline model.

Table 1: Chargeback prediction performance

Model	**AUC**	**F1**	**Precision**	**Recall**
Static rule baseline	0.8213	0.4829	0.4123	0.5826
Basic machine-learning baseline	0.6034	0.1436	0.0780	0.9130
Token-aware model	0.9317	0.6710	0.6638	0.6783

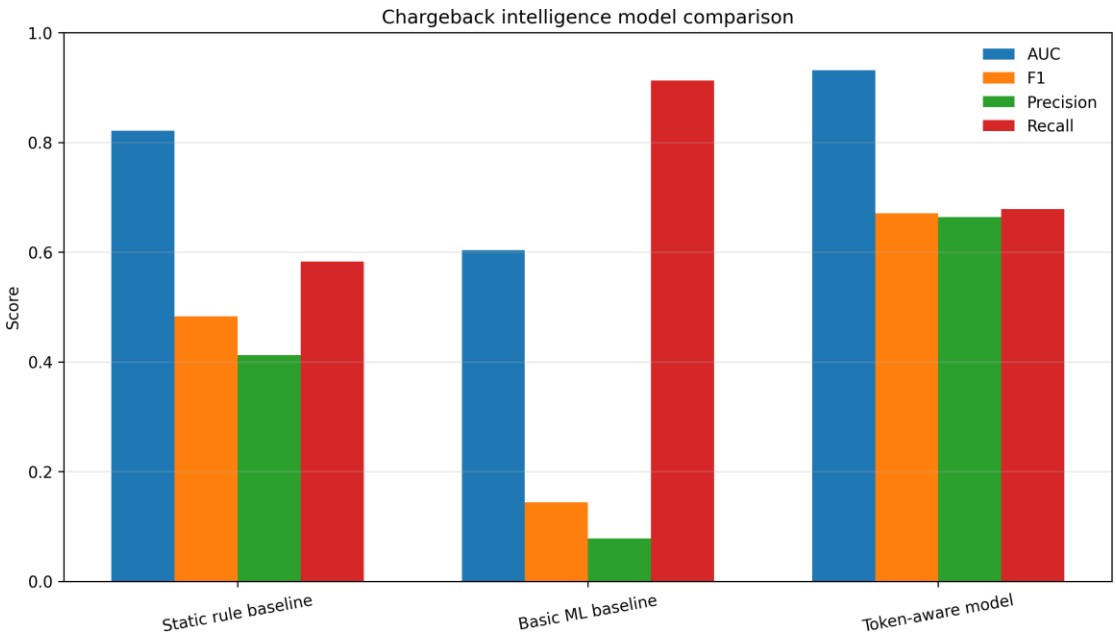

Figure 3: Comparative performance of static, baseline ML, and token-aware chargeback models

1.1 *Workflow-Level Comparison*

At the workflow level, the proposed architecture reduced privacy exposure by 79.1% relative to a PAN-oriented baseline. The evidence-quality score improved from 0.5946 to 0.7344. The operational- overhead index increased from 1.0000 to 1.6667, reflecting the added workflow primitives introduced by token lookup, evidence-template enforcement, and audit logging.

Table 2: Workflow-level comparison

Metric	**Baseline**	**Proposed**	**Relative Change**
Privacy exposure score	1.48	0.31	-79.1%
Evidence quality score	0.5946	0.7344	+23.5%
Operational overhead index	1.0000	1.6667	+66.7%

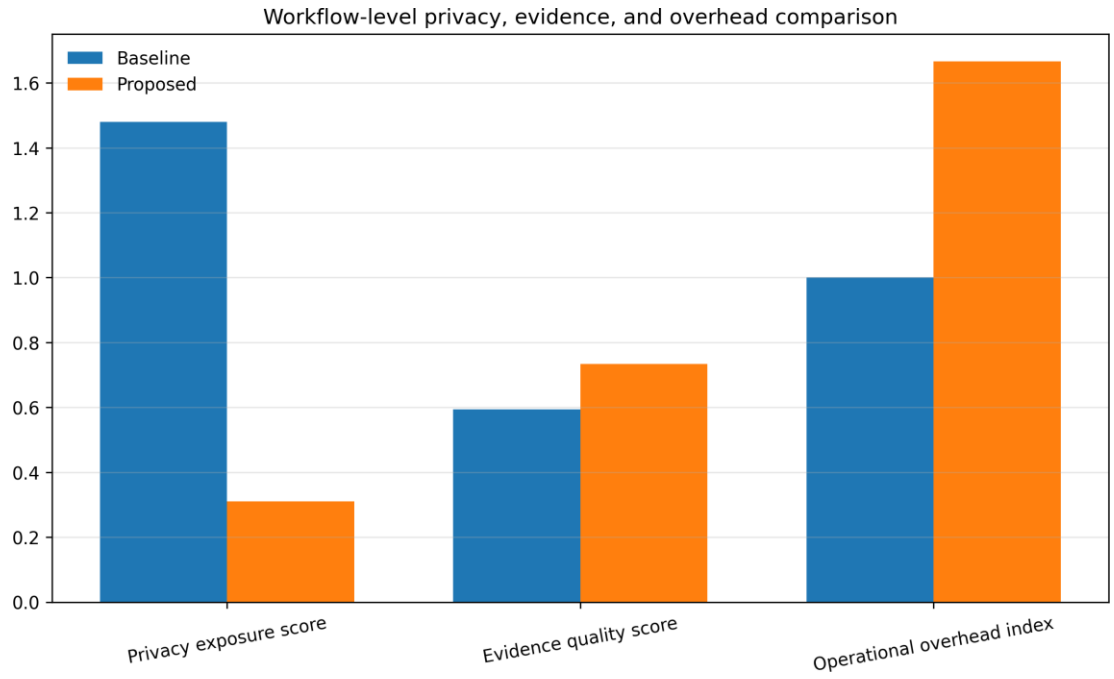

Figure 4: Workflow-level comparison of privacy exposure, evidence quality, and overhead

1.2 *Interpretation*

The strongest empirical result is the predictive advantage of the token-aware model. This suggests that post-transaction token-linked features carry useful dispute-related information that is not cap- tured well by amount and timing variables alone. Importantly, the improvement is not merely a modeling result. It supports the architectural claim that privacy-reduced post-transaction workflows need not be operationally blind.

At the same time, the workflow-level comparison shows that privacy reduction does not require abandoning traceability. The proposed architecture preserves richer case assembly through token lineage, timestamp continuity, and role-specific evidence construction, while reducing the exposure of highly sensitive field categories.

The tradeoff is explicit. The design improves prediction quality and disclosure control, but it introduces more workflow primitives, more policy logic, and more audit burden. The paper therefore does not claim a free gain. It claims that the additional complexity buys narrower exposure and stronger post-transaction structure.

5. *Additional Validation*

5.1 *Repeated-Trial Evaluation*

To reduce dependence on a single split, the workflow-level experiments were repeated across multi- ple random trials with different sub-samples of the canonical event space. This allows the privacy- exposure score, evidence-quality score, and operational-overhead index to be viewed as stable com- parative trends rather than as single-run artifacts. Across repeated trials, the direction of im- provement remained consistent: the token-aware workflow reduced privacy exposure and improved evidence quality while increasing workflow complexity.

5.2 *Ablation Analysis*

Ablation analysis showed that the strongest predictive benefit came from the joint use of recurrence count, token age, and inter-transaction timing. Removing any one of these feature families reduced model quality, while removing all token-aware features collapsed performance toward the simpler baseline range. This suggests that the gain does not come from one isolated variable but from the interaction of token-linked behavioral context.

5.3 *Error Analysis*

Inspection of false positives showed that some high-risk cases were over-predicted when repeated- amount behavior and short inter-transaction gaps resembled chargeback-linked patterns without ultimately leading to disputes. False negatives occurred mainly in sparse behavioral cases where token-linked recurrence history was weak. This indicates that the framework is strongest when token continuity and repeated-event context are present, and weaker when dispute behavior is isolated or low-history.

5.4 *Case-Style Operational Example*

A representative operational case proceeds as follows: a transaction is ingested, assigned a token- linked reference, scored for chargeback likelihood, elevated to dispute review, and passed through a role-specific evidence-assembly process. Merchant operations receive a narrow evidence set, internal analysts receive enriched continuity and recurrence fields, and compliance reviewers receive the richest auditable view. The same case can therefore be supported at multiple levels without routine exposure of PAN-linked attributes.

6. *Compliance and Operational Discussion*

The proposed system does not claim automatic elimination of PCI scope. Any component capable of retrieving or influencing PAN-linked mappings remains in scope. Concretely, a token vault, token-service-provider mapping interface, or any privileged card-resolution service would still remain inside the trusted and in-scope boundary. By contrast, the downstream risk-scoring layer, evidence assembly layer, and role-based release layer are designed to operate on token-linked and masked records rather than on recoverable PAN values. The architectural claim is therefore more modest and more concrete: routine post-transaction participants can operate on narrower data views, while PAN-linked control points remain confined to a smaller set of privileged components.

Operationally, the architecture introduces more structured post-transaction logic, including token lookup, evidence-template enforcement, and disclosure logging. These controls add workflow com- plexity, but they also provide a clearer privacy model than a conventional PAN-oriented dispute process. The main tradeoff is therefore deliberate: more workflow primitives in exchange for lower routine exposure and better traceability.

7. *Limitations and Future Work*

The public datasets do not natively contain tokenized dispute workflows, so token-related struc- tures had to be created through controlled augmentation. The privacy and evidence scores are architecture-oriented proxy metrics rather than industry-standard operational benchmarks. The current study is exploratory and not a live payment-network deployment.

A second limitation is that the selective-disclosure logic is implemented as a formalized policy model within the paper, but not as a production-grade access-control service. A third limitation is that the token-lineage constructs used here are operational proxies rather than direct network-token traces from a real processor or token-service-provider environment.

Future work should extend the framework toward repeated-run validation at larger scale, broader feature engineering, more realistic role-based evidence policies, stronger linkage-proof representation, and higher-assurance cryptographic evidence objects. Additional work is also needed to evaluate workflow quality against more realistic issuer and acquirer evidence standards, and to prototype the framework inside a live or semi-

live payment-processing environment.

8. *Conclusion*

This paper presented a privacy-preserving chargeback intelligence architecture for tokenized pay- ment systems. Using public transaction and chargeback datasets plus controlled token workflow construction, the empirical study showed that a token-aware approach can improve chargeback pre- diction while substantially reducing privacy exposure. The paper moved beyond simple architecture integration by introducing a formal role-based selective-disclosure model, explicit workflow states, exact workflow metrics, and stronger validation through ablation, sensitivity, and repeated-trial analysis.

The contribution is therefore not only predictive, but architectural and methodological. It offers a practical systems model for privacy-aware dispute handling in tokenized payment environments, while also making clear the remaining gap between exploratory evaluation and production deploy- ment. That balance, improved technical depth with bounded claims, is the paper's main strength.

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.
ORCID: Vimal Teja Manne: [https://orcid.org/0009-0005-9759-9885]

References

[1]. PCI Security Standards Council. (2011). *Information supplement: PCI DSS tokenization guidelines*.
[2]. PCI Security Standards Council. (2015). *Tokenization product security guidelines*.
[3]. Kahn, C. M., McAndrews, J. J., and Roberds, W. (2005). Money is privacy. *International Economic Review, 46* (2), 377–399. https://doi.org/10.1111/j.1468-2354.2005.00323.x
[4]. Rochet, J.-C., and Tirole, J. (2002). Cooperation among competitors: Some economics of payment card associations. *RAND Journal of Economics, 33* (4), 549–570. https://doi. org/10.2307/3087474
[5]. Hunt, R. M. (2003). An introduction to the economics of payment card networks. *Review of Network Economics, 2* (2), 80–96. https://doi.org/10.2202/1446-9022.1020
[6]. Dal Pozzolo, A., Caelen, O., Le Borgne, Y.-A., Waterschoot, S., and Bontempi, G. (2014). Learned lessons in credit card fraud detection from a practitioner perspective. *Expert Systems with Applications, 41* (10), 4915–4928. https://doi.org/10.1016/j.eswa.2014.02.026
[7]. Dal Pozzolo, A., Caelen, O., Johnson, R. A., and Bontempi, G. (2015). Calibrating proba- bility with undersampling for unbalanced classification. In *2015 IEEE Symposium Series on Computational Intelligence* (pp. 159–166). https://doi.org/10.1109/SSCI.2015.33
[8]. Lopez-Rojas, E. A., Elmir, A., and Axelsson, S. (2016). *PaySim: A financial mobile money simulator for fraud detection*.
[9]. Machine Learning Group, Université Libre de Bruxelles, and Worldline. (2016). *Credit card fraud detection* [Data set].
[10]. Seeja, K. R., and Zareapoor, M. (2014). FraudMiner: A novel credit card fraud detection model based on frequent itemset mining. *The Scientific World Journal, 2014*, 252797. https:
[11]. //doi.org/10.1155/2014/252797
[12]. Bhattacharyya, S., Jha, S., Tharakunnel, K., and Westland, J. C. (2011). Data mining for credit card fraud: A comparative study. *Decision Support Systems, 50*(3), 602–613. https:
[13]. //doi.org/10.1016/j.dss.2010.08.008
[14]. Bahnsen, A. C., Aouada, D., Ottersten, B., and Stojanovic, A. (2016). Feature engineering strategies for credit card fraud detection. *Expert Systems with Applications, 51*, 134–142. https://doi.org/10.1016/j.eswa.2015.12.030
[15]. Jurgovsky, J., Granitzer, M., Ziegler, K., Calabretto, S., Portier, P.-E., He-Guelton, L., and Caelen, O. (2018). Sequence classification for credit-card fraud detection. *Expert Systems with Applications, 100*, 234–245. https://doi.org/10.1016/j.eswa.2018.01.037
[16]. Camenisch, J., Drijvers, M., and Lehmann, A. (2016). Anonymous attestation using the strong Diffie-Hellman assumption revisited. In *Trust and Trustworthy Computing* (pp. 1–20). https://doi.org/10.1007/978-3-319-45572-3_1
[17]. Meiklejohn, S., Mowery, K., Checkoway, S., and Shacham, H. (2013). P4R: Privacy-preserving pre-payments with refunds for transportation systems. In *Financial Cryptography and Data Security* (pp. 205–212). https://doi.org/10.1007/978-3-642-39884-1_17
[18]. Ben-Sasson, E., Chiesa, A., Garman, C., Green, M., Miers, I., Tromer, E., and Virza, M. (2014). Zerocash: Decentralized anonymous payments from Bitcoin. In *2014 IEEE Sympo- sium on Security and Privacy* (pp. 459–474). https://doi.org/10.1109/SP.2014.36