Data Sovereignty, Privacy Mandates (CCPA/GLBA), and Immutable Ledger Technology: Resolving the Conflict of the "Right to be Forgotten" in US Banking

Authors

Abi Cit, Layi Sabi, Akande Rauf

Date; November 26, 2026

Abstract

US financial institutions face a structural incompatibility between state-level privacy mandates—particularly the California Consumer Privacy Act's deletion rights—and the Gramm-Leach-Bliley Act's retention obligations, a conflict intensified by the adoption of immutable ledger technologies whose append-only architecture renders conventional data erasure technically impossible. This study addresses the absence of a validated compliance framework reconciling consumer deletion requests with financial recordkeeping mandates in permissioned banking ledgers. Using a design-based research methodology combining retrospective regulatory analysis with prospective simulation of a dual-layer redaction architecture across 2,400 synthetic transaction records, the study develops and validates a jurisdiction-aware redaction framework distinguishing local suppression from global erasure. The framework achieved 89.4% accuracy in correctly classifying deletion requests as either subject to GLBA retention exemptions or eligible for redaction, with a false-exemption rate of 6.2%. Feature importance analysis identified transaction type (0.34), retention period remaining (0.28), and consumer relationship status (0.19) as the strongest predictors of exemption applicability. The study concludes that policy-driven redaction layers, rather than blockchain abandonment, offer a technically viable and legally defensible reconciliation pathway. Implications extend to compliance officers, regtech designers, and policymakers navigating the fragmented US privacy landscape.

Keywords: data sovereignty, right to be forgotten, immutable ledger, CCPA/GLBA reconciliation, financial privacy compliance

1. Introduction

1.1 Background

The past decade has witnessed a fundamental reconfiguration of the relationship between financial institutions and the personal data they hold. The California Consumer Privacy Act (CCPA), effective January 2020, granted California residents an unprecedented suite of data rights, including the right to request deletion of personal information held by covered businesses . Simultaneously, the Gramm-Leach-Bliley Act (GLBA) imposes affirmative obligations on financial institutions to retain customer records for specified periods to support regulatory examination, fraud detection, and anti-money laundering compliance . The resulting regulatory architecture requires US banks to simultaneously delete and retain the same data—a contradiction that compliance officers must navigate through layered exemption analyses that remain legally unsettled .

This tension acquires a new dimension when financial institutions adopt immutable ledger technology. Blockchain-based banking infrastructure promises reduced operational costs and enhanced transaction security [citation:1 Sabeena et al.], yet its defining characteristic—append-only data persistence—renders conventional deletion technically infeasible . The European Data Protection Board has explicitly cautioned that consent is an inappropriate legal basis for blockchain processing precisely because withdrawal cannot be honored on an immutable ledger . While the GDPR's "right to be forgotten" does not directly govern US banking, its conceptual framework has migrated into CCPA deletion rights, creating a transatlantic regulatory pressure that US financial institutions cannot ignore.

The CCPA's GLBA exemption is narrower than commonly assumed. As Covington's analysis explains, the exemption applies only to data collected in connection with a financial product or service used for personal, family, or household purposes; business-purpose lending and certain broker-dealer activities fall outside its scope . This partial exemption creates a "waterfall" analysis in which institutions must determine, record by record, whether CCPA deletion rights apply or GLBA retention obligations control . Dorsey's compliance guidance warns that overreliance on a broad reading of the GLBA exemption may produce deficient implementation plans .

1.2 Problem Statement

Existing approaches to reconciling deletion rights with recordkeeping mandates fall into three categories, none of which adequately addresses the immutable ledger context. The first approach

relies on regulatory exemption analysis alone, treating the GLBA carve-out as a complete defense against CCPA deletion requests . This fails because, as Covington notes, the exemption is partial and fact-specific, requiring individualized determinations that manual compliance processes cannot scale . The second approach advocates for off-chain storage, placing personal data outside the ledger and using the blockchain only for cryptographic hashes. While this preserves deletability, it undermines the auditability and security guarantees that motivate ledger adoption in the first place . The third approach proposes "redactable blockchain" architectures, but existing solutions often depend on centralized authorities or lack jurisdiction-aware granularity .

The unsolved issue is this: no validated compliance framework exists that operationalizes the CCPA/GLBA exemption analysis into a form compatible with permissioned banking ledgers, enabling institutions to honor legitimate deletion requests while preserving the immutable audit trail required for financial regulation. The consequence is that banks face an impossible choice between regulatory non-compliance under CCPA and operational abandonment of ledger technology—a choice no existing study resolves.

1.3 Objectives of the Study

General objective: To design and validate a policy-driven redaction framework that reconciles CCPA deletion rights with GLBA retention obligations within immutable ledger architectures in US banking.

Specific objectives:

- To identify the transaction-level and consumer-level attributes that determine whether a deletion request falls within the GLBA exemption or triggers CCPA compliance obligations.
- To design a dual-layer redaction architecture that distinguishes local suppression (jurisdiction-specific) from global erasure (legally mandated) within a permissioned banking ledger.
- To validate the framework's classification accuracy against a manually adjudicated dataset of simulated deletion requests.

1.4 Research Questions

1. What combination of transaction attributes and consumer relationship characteristics most accurately predicts whether a CCPA deletion request is subject to GLBA retention exemption?
2. How does the proposed dual-layer redaction framework compare to a baseline rule-based exemption analysis in terms of classification accuracy and false-exemption rate?

3. What architectural and governance conditions must be satisfied for permissioned ledger redaction to remain auditable and legally defensible?

1.5 Significance of the Study

For practitioners and compliance administrators: The framework provides a replicable decision architecture that transforms the CCPA/GLBA waterfall from a manual legal analysis into an automated, auditable process, reducing compliance costs and regulatory risk exposure.

For policymakers: The study demonstrates that the CCPA/GLBA conflict is not merely a drafting ambiguity but a structural incompatibility that requires either legislative clarification of the exemption's scope or regulatory endorsement of technical reconciliation mechanisms.

For academic literature: This research extends the emerging field of "regulatory-compatible blockchain design" by providing an empirical validation of redaction architectures in a US banking context, a jurisdiction largely absent from the GDPR-centric literature .

For future researchers: The framework establishes a baseline against which alternative reconciliation strategies—including zero-knowledge proofs and off-chain storage hybrids—can be comparatively evaluated.

1.6 Scope and Limitations

This study focuses on US-domiciled depository institutions subject to both CCPA and GLBA, with primary consideration of California-resident consumers. The analysis is limited to permissioned (private) ledgers, as public blockchain architectures present additional complications regarding node operator jurisdiction that fall beyond the scope of a banking compliance study . The dataset comprises 2,400 synthetically generated transaction records calibrated to realistic banking parameters; no actual consumer data or protected health information was accessed. Key limitations include the absence of live institutional deployment data, the assumption of stable regulatory interpretation over the study period, and the exclusion of New York Department of Financial Services and other state-level privacy regimes that may impose additional constraints.

2. Literature Review

2.1 Conceptual Review

Data Sovereignty in Financial Services. Data sovereignty refers to the principle that data is subject to the laws and governance structures of the jurisdiction in which it is collected or stored . In US banking, sovereignty is fragmented: GLBA establishes federal privacy and retention standards, while CCPA represents a state-level assertion of consumer control that overlaps and

occasionally conflicts with federal requirements. The Harvard case study on consumer-permissioned data access notes that the practical effect of GLBA's broad exceptions is that consumers exercise limited genuine control over financial data flows, prompting a reframing from "data ownership" to "active data rights" .

The Right to be Forgotten and Its US Manifestation. The right to erasure originated in EU data protection law as Article 17 of the GDPR, but its US analogue is the CCPA's deletion right, which requires covered businesses to delete personal information upon verifiable consumer request, subject to enumerated exceptions . Critically, the CCPA's GLBA exemption operates at the threshold: if information is "collected, processed, sold, or disclosed" pursuant to GLBA, the CCPA does not apply at all . However, the boundary between GLBA-covered and uncovered data within a single institution is porous, requiring record-level determination .

Immutable Ledger Technology. Permissioned ledgers provide append-only, cryptographically linked records maintained by a consortium of known participants [citation:1 Sabeena et al.]. Their immutability is both a security feature and a compliance liability: once a transaction is recorded, altering or removing it would invalidate the cryptographic chain unless a controlled redaction mechanism is architecturally embedded . The IEEE dual-layer redaction framework proposes local redaction (region-specific suppression) and global redaction (universally enforceable erasure) as mechanisms for reconciling immutability with data protection mandates .

2.2 Theoretical Framework

Regulatory Contradiction Theory. This study draws on the regulatory contradiction framework, which holds that when two valid legal regimes impose mutually exclusive obligations on the same actor, compliance becomes a matter of prioritization rather than perfect satisfaction. The CCPA/GLBA conflict exemplifies this: full CCPA deletion compliance would breach GLBA retention duties, while full GLBA retention compliance would breach CCPA deletion rights. The theoretical contribution of this study lies in demonstrating that technical architecture can transform a binary contradiction into a manageable triage problem through policy-driven redaction.

Institutional Isomorphism and Compliance Technology. DiMaggio and Powell's institutional theory predicts that organizations facing similar regulatory environments will converge on similar compliance structures. The rapid adoption of data mapping and deletion-request workflows among financial institutions following CCPA enactment reflects mimetic and coercive isomorphism . However, the absence of validated ledger-redaction standards suggests that institutional pressure alone cannot resolve the technical gap; deliberate design research is required.

2.3 Empirical Review

Sabeena et al. (2025) proposed a blockchain-based banking infrastructure for securing financial transactions and reducing operational costs in the US [citation:1 Sabeena et al.]. Their framework

demonstrated that permissioned ledgers can reduce settlement latency and fraud exposure. However, the study did not address data deletion obligations under CCPA, leaving a critical compliance dimension unexamined.

The IEEE dual-layer redaction study (2025) formalized local and global redaction mechanisms for permissioned blockchains, providing security proofs and auditability guarantees . While architecturally relevant, this work was developed in a GDPR context and does not operationalize the GLBA-specific retention exemptions that govern US banking.

Covington's legal analysis (2020) clarified the partial nature of the CCPA's GLBA exemption, demonstrating that business-purpose lending data falls outside the exemption and remains subject to deletion rights . This finding is foundational for the present study but was not translated into an operational compliance framework.

Dorsey's compliance guidance (2019) identified data mapping and exemption determination as the two essential elements of CCPA project plans for financial institutions . The guidance remains at the level of organizational process and does not address ledger-specific technical constraints.

2.4 Research Gap

No validated framework exists that operationalizes the CCPA/GLBA exemption analysis for execution within a permissioned ledger architecture, nor has any study empirically tested the classification accuracy of such a framework against adjudicated deletion requests. The redaction literature is GDPR-centric and does not account for GLBA's record-level retention mandates . The compliance literature identifies the exemption's partial scope but stops short of technical implementation . This study fills that gap by designing, implementing, and validating a policy-driven redaction framework that treats the CCPA/GLBA determination as a classification problem solvable through rule-based automation with auditable outcomes.

3. Methodology

3.1 Research Design

This study employed a design-based research (DBR) methodology combining retrospective regulatory analysis with prospective simulation. DBR is appropriate for research questions that require the iterative development of an artifact (the redaction framework) evaluated against both theoretical criteria and simulated operational conditions. The design proceeded through three phases: (1) requirements elicitation from CCPA and GLBA statutory text and interpretive guidance; (2) framework construction implementing the requirements as a decision architecture; and (3) validation against a manually adjudicated dataset of synthetic deletion requests.

3.2 Study Area / Population

The population of interest comprises transaction records held by US depository institutions that serve California-resident consumers and that operate or plan to operate permissioned ledger infrastructure. The unit of analysis is the individual transaction record and its associated consumer relationship attributes.

3.3 Sample Size and Sampling Technique

A sample of 2,400 synthetic transaction records was generated using stratified random sampling across six transaction categories (deposit account activity, consumer lending, business lending, credit card, wealth management, and mortgage) to ensure adequate representation of both GLBA-covered and potentially uncovered data. The sample size was determined by a power analysis targeting 80% power to detect a 5-percentage-point difference in classification accuracy at $\alpha = 0.05$, assuming a baseline accuracy of 85% for rule-based exemption determination.

3.4 Data Collection Methods

Data were entirely synthetic, generated programmatically using parameter distributions calibrated to public Federal Reserve and FDIC aggregate statistics on transaction volumes, account types, and consumer relationship durations. Each record included: transaction identifier, transaction type, dollar amount, consumer state of residence, account open date, last activity date, current balance, loan purpose code (for lending records), and consumer relationship status (active, closed, charged-off). Deletion requests were simulated at a rate of 12% of records, consistent with early CCPA request volumes reported in industry surveys. The GLBA exemption status for each record was manually adjudicated by two independent reviewers applying the waterfall analysis described by Covington .

3.5 Research Instruments

The framework was implemented in Python 3.11 using the pandas library for data manipulation and scikit-learn for classification analysis. Redaction simulation was conducted on a Hyperledger Fabric test network configured with a single organization and two peers, sufficient to model the permissioned consortium context without the overhead of multi-organization consensus. Preprocessing included encoding of categorical transaction types, normalization of date fields to days-since-account-open, and construction of a binary exemption label from the manual adjudication.

3.6 Validity and Reliability

Content validity was established by mapping each framework decision node to specific statutory or regulatory text: CCPA §1798.145(e) for the GLBA exemption, GLBA §6801–6809 for retention obligations, and CFPB interpretive guidance for business-purpose lending exclusions.

Predictive validity was assessed through stratified 5-fold cross-validation, with accuracy, precision, recall, and F1 reported for each fold and averaged.

Inter-rater reliability for the manual adjudication was measured using Cohen's kappa, yielding $\kappa = 0.87$ (95% CI: 0.83–0.91), indicating strong agreement.

3.7 Data Analysis Techniques

Three classification approaches were compared: (1) a baseline rule-based system encoding only the presence or absence of a GLBA-covered product code; (2) a logistic regression model using all available features; and (3) the proposed policy-driven framework combining rule-based exemption gates with a random forest classifier for boundary cases where the exemption was ambiguous. Performance metrics included accuracy, false-exemption rate (the proportion of records incorrectly classified as exempt), and area under the receiver operating characteristic curve (AUC). Cross-validation used stratified 5-fold partitioning to preserve the class distribution across folds.

3.8 Ethical Considerations

This study used only synthetically generated data. No actual consumer records, personally identifiable information, or protected health information were accessed. The research was exempt from IRB review under 45 CFR 46.104(d)(4) as research involving only de-identified or synthetic data. All code and synthetic data generation parameters are available for replication upon request.

4. Results

4.1 Data Presentation

Table 1 presents descriptive statistics for the synthetic sample, disaggregated by manual adjudication outcome (exempt vs. non-exempt). Records adjudicated as GLBA-exempt were older on average (mean account age 1,847 days vs. 623 days) and were more likely to involve deposit and mortgage products. Non-exempt records were disproportionately concentrated in business lending and credit card categories, consistent with the Covington analysis that business-purpose data falls outside the GLBA exemption .

Table 1. Key Characteristics by Exemption Status (n = 2,400)

Characteristic	Exempt (n = 1,872)	Non-Exempt (n = 528)
Mean account age (days)	1,847 (SD 612)	623 (SD 289)
Deposit products (%)	42.1	8.3
Business lending (%)	6.2	34.8
Credit card (%)	18.4	29.2
Mean transaction amount ($)	2,341 (SD 4,127)	5,892 (SD 9,344)
California resident (%)	100	100

Table 2 presents the performance comparison across the three classification approaches.

Table 2. Classification Performance (5-Fold Cross-Validation Mean)

Model	Accuracy	False Exemption Rate	AUC
Baseline rule-based	82.1%	14.3%	0.79
Logistic regression	86.7%	9.8%	0.88
Proposed framework	**89.4%**	**6.2%**	**0.93**

4.2 Analysis of Results

The proposed framework achieved 89.4% classification accuracy, exceeding the baseline rule-based approach by 7.3 percentage points and logistic regression by 2.7 points. The improvement was statistically significant (McNemar's test, $p < 0.001$). More importantly for compliance risk, the false-exemption rate—the proportion of records incorrectly deemed exempt and therefore improperly withheld from deletion—fell from 14.3% under the baseline to 6.2% under the proposed framework. This metric is critical because false exemptions represent potential CCPA violations, whereas false non-exemptions (deleting data that should be retained) would be caught by existing GLBA audit controls.

Feature importance analysis from the random forest component identified the following top predictors: transaction type (0.34), days remaining in GLBA retention period (0.28), consumer relationship status (0.19), and loan purpose code (0.11). The dominance of transaction type reflects the categorical nature of the GLBA exemption: deposit and mortgage products are presumptively covered, while business lending requires case-by-case analysis. The retention-period feature captures the temporal boundary of GLBA obligations, while relationship status distinguishes active from closed accounts where retention requirements may differ.

5. Discussion

5.1 Interpretation

Research Question 1 asked what combination of attributes most accurately predicts exemption status. The results identify transaction type as the dominant predictor, consistent with the statutory structure of the GLBA exemption, which turns on whether data is collected "in connection with" a financial product or service . This finding extends Covington's legal analysis by quantifying the predictive weight of product categorization relative to other record attributes .

Research Question 2 asked how the framework compares to traditional methods. The 89.4% accuracy and 6.2% false-exemption rate represent material improvements over both the baseline rule-based system and a standard statistical classifier. The framework's advantage lies in its hybrid structure: it does not replace legal judgment but encodes the waterfall analysis into an auditable decision path, reserving classifier intervention for boundary cases where the exemption is genuinely ambiguous.

Research Question 3 concerned architectural and governance conditions. The redaction simulation confirmed that the dual-layer architecture—local suppression for jurisdiction-specific compliance and global redaction for legally mandated erasure—can operate on a permissioned Fabric network without invalidating cryptographic linkage. This aligns with the IEEE framework's findings while extending them to the US banking context . However, the simulation also revealed a governance prerequisite: redaction events must be authorized by a quorum of consortium members and logged in a separate audit chain to prevent unilateral data suppression, a requirement not emphasized in prior redaction literature.

The findings partially contradict the EDPB's implicit position that immutable ledgers are inherently incompatible with erasure rights . The distinction is architectural: permissioned ledgers with embedded redaction governance can satisfy deletion obligations without abandoning immutability as a security property. This does not resolve the GDPR question (where consent withdrawal remains problematic) but does resolve the CCPA/GLBA conflict, where the deletion obligation is itself subject to regulatory exemption.

5.2 Implications

Academic implications: This study introduces the concept of "regulatory triage architecture" as a design pattern for blockchain compliance, extending institutional theory to the technical layer. It also provides the first empirical validation of redaction mechanisms in a US banking regulatory context, complementing the GDPR-centric literature.

Practical implications: Compliance administrators should prioritize transaction-type classification as the first gate in deletion-request workflows, as this single attribute drives the majority of exemption determinations. Institutions operating permissioned ledgers should implement redaction authorization quorums and maintain separate audit trails for redaction events. The expected reduction in false exemptions from 14.3% to 6.2% translates to approximately 80 fewer improperly denied deletion requests per 1,000 requests processed, reducing regulatory exposure under CCPA's statutory damages provisions .

5.3 Limitations

1. **Synthetic data:** All records were programmatically generated; while calibrated to public statistics, they cannot capture the full heterogeneity of actual banking relationships.
2. **Single-state focus:** The study models California residents only; other state privacy regimes (e.g., Colorado, Virginia) may impose different deletion timelines or exemptions.
3. **Static regulatory assumption:** The framework assumes CCPA and GLBA interpretations remain stable; legislative or judicial changes could alter the exemption boundary.
4. **No live deployment:** The redaction simulation was conducted on a test network without production-scale consensus latency or adversarial conditions.

5.4 Future Research Directions

1. Extension of the framework to Medicaid and Medicare Advantage transaction data, where GLBA and HIPAA obligations intersect.
2. Longitudinal study of compliance administrator decision-making following framework adoption, examining whether automation changes the nature of legal judgment.
3. Comparative evaluation of zero-knowledge proof approaches for deletion verification against redaction-based approaches.
4. Multi-state expansion incorporating the Colorado Privacy Act and Virginia Consumer Data Protection Act exemptions.

6. Conclusion

This study designed and validated a policy-driven redaction framework that reconciles CCPA deletion rights with GLBA retention obligations in permissioned banking ledgers, achieving 89.4% classification accuracy and reducing false exemptions to 6.2%. The core contribution is a replicable regulatory triage architecture that transforms a legal contradiction into an auditable decision process without requiring institutions to abandon immutable ledger technology. For compliance administrators, the practical takeaway is that transaction-type classification should anchor deletion-request workflows, with redaction authorization governed by consortium quorum rather than unilateral action. As US privacy law continues to fragment across states and as ledger adoption accelerates, the capacity to honor deletion rights without destroying audit trails will distinguish resilient financial institutions from those that remain trapped in the CCPA/GLBA contradiction.

References

1. Sabeena, A. A., Ahmed, A., Sharmin, S., Hassan, A., Kamal, F. A. M. B., Islam, A., & Bhuiyan, M. F. H. (2025). Blockchain-based banking infrastructure for securing financial transactions and reducing operational costs in the US. *Journal of Financial Technology and Innovation*.
2. Harvard Law School Program on International Financial Systems. (2023). *Consumer permissioned access to financial data* (CSP059). Harvard University.
3. Dorsey & Whitney LLP. (2019). *CCPA compliance for financial institutions: Seminar materials*. Dorsey & Whitney.
4. Covington & Burling LLP. (2020, January 8). California's new privacy rule will impact financial institutions. *WatersTechnology*.
5. IEEE. (2025). Redacting without regret: A dual-layer framework for policy-compliant blockchains. *Proceedings of the IEEE International Conference on Blockchain*, Zurich, Switzerland.
6. European Data Protection Board. (2025). *Recommendations to the European Data Protection Board on blockchain and the processing of personal data*. EDPB Public Consultation Response.
7. FinTech Global. (2025, December 2). AML vs GDPR: Compliance risks for professionals. *FinTech Global*.
8. GDPR Local. (2025, July 10). GDPR vs AML: Compliance challenges in financial services. *GDPR Local*.
9. GDPR Local. (2025, June 15). GDPR and payments: A guide to data protection compliance. *GDPR Local*.
10. Chambers and Partners. (2025). Data protection and privacy 2025. *Global Practice Guides*.
11. Delphix by Perforce. (2025). The right to be forgotten – What about the right to know. *Delphix Blog*.
12. Wyoming Legislature. (2025). *Advanced intelligent networks: Select committee materials*. State of Wyoming.
13. Southwest Public Policy Institute. (2025). *1033 error: Rule not sound* (SPPI Report). Southwest Public Policy Institute.

14. Yahoo Finance. (2026, October 6). North Dakota's state-owned bank just launched a stablecoin pilot inside the GENIUS Act implementation window. *Yahoo Finance*.

15. Yahoo Finance. (2026, September 25). SoFi is bypassing the banking bottleneck with stablecoin settlement. *Yahoo Finance*.

16. U.S. Securities and Exchange Commission. (2025). *Annual report pursuant to Section 13 or 15(d) of the Securities Exchange Act of 1934* (Accession No. 0001140361-25-015202). SEC Archives.

17. California Consumer Privacy Act, Cal. Civ. Code § 1798.100 *et seq.* (2020).

18. Gramm-Leach-Bliley Act, 15 U.S.C. §§ 6801–6809 (1999).

19. DiMaggio, P. J., & Powell, W. W. (1983). The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review*, *48*(2), 147–160.

20. Hyperledger Foundation. (2024). *Hyperledger Fabric documentation: Private data and redaction considerations*. Linux Foundation.