JFDA THE JOURNAL OF FINTECH & DIGITAL ASSETS

Forthcoming 2027

| Research Paper

The Asymmetric Attack Surface of Fiat Substitutes: A Measurement Framework and Owner-Bound Offline Custody Architecture

Alex Breton*

Authors
Alex Breton, Founder & CEO,
PrintDreams International AB, Sweden.
Email: alex.breton@printdreams.com

* Corresponding author.

Abstract
USDC and USDT can alter theft risk through remote reach, reusable exploit logic, rapid transfer, and common-mode compromise. This methodological feasibility study distinguishes attack topology from realized loss and defines exposure-normalized gross and net loss intensities under explicit comparability conditions. Public evidence does not identify a global per-dollar theft probability or a defensible cross-domain deterrence ratio. As an illustrative reported-loss calibration, 2019 Federal Bureau of Investigation data record USD 1.424 billion in reported stolen currency and USD 36.981 million recovered; relative to USD 1.7598 trillion of U.S. currency in circulation at year-end, gross reported loss intensity is 0.0809% and the recovery share is 2.60%. Binance breach reports estimate approximately USD 6.30–6.43 billion stolen in October 2026 but provide no matched USDC/USDT loss numerator and exposure denominator. Relative-vulnerability comparisons remain conditional estimands. The study specifies an owner-bound offline custody architecture in which the complete spend key is never printed, independent shares are required, and signing occurs in a verified offline environment. A framework linking attacker incentives to net loss and a control-level comparison identify conditional benefits, transferred attack paths, and residual risks without claiming empirically demonstrated superiority.

Keywords: cybercrime measurement, digital asset custody, owner-bound substrate, secret sharing, stablecoins
JEL Classification: E42, G23, K42

Article Information
Received: 25 August 2026
Revised: 13 October 2026
Accepted: 20 October 2026
Published: 03 November 2026

Citation: Breton, A. (in press). The asymmetric attack surface of fiat substitutes: A measurement framework and owner-bound offline custody architecture. *The Journal of FinTech and Digital Assets*.

1. Introduction

The fiat-backed stablecoins USDC and USDT extend monetary functions into programmable networks. Their potential benefits include rapid settlement, composability, and cross-border reach. The same characteristics change how theft may be attempted. A remote adversary can scan many systems, reuse exploit logic, compromise a shared dependency, and move assets before institutions and law-enforcement agencies coordinate. Physical cash theft generally requires proximity, access, and transport, although wholesale storage, insider access, and common security providers can also create concentrated losses.

After this manuscript was accepted, the 29 October 2026 Binance breach became the largest reported cryptocurrency theft (TRM Labs, 2026; Robinson, 2026). Its reported gross losses exceed the approximately USD 2.2 billion that each of Chainalysis and TRM Labs reported for identified hacks in 2024. This compares

nominal reported amounts, not exposure-normalized risk. The incident gives immediate practical relevance to the distinction between authorization security, correlated custody exposure, and recovery capacity. Section 4.2 presents a dated evidence snapshot.

This difference is an asymmetry in attack topology. It is not, without matched data and additional assumptions, proof that a randomly selected digital dollar has a particular annual probability of theft. Public cash-crime data typically measure reported offenses or stolen property in one jurisdiction. Blockchain-intelligence reports usually aggregate identified on-chain incidents across assets, custody models, and countries. Currency in circulation, stablecoin supply, transaction volume, and assets under custody are different economic quantities. Ratios formed from unmatched quantities may be arithmetically correct yet scientifically uninterpretable.

The distinction matters for both policy and engineering. Unsupported empirical precision can obscure a valid structural concern. Absolute technical claims can be equally misleading: a passive record has no network interface, but the provisioning system, offline signer, recovery process, and human operators remain attackable. A useful analysis must therefore separate topology, observed loss, enforcement, and control effectiveness.

This paper addresses four research questions. The first one (RQ1) is under what conditions can exposure-normalized loss intensities be compared across physical cash and USDC or USDT in a specified custody configuration? Second (RQ2), how do remote reach, automation, common-mode compromise, transfer speed, and enforcement friction enter an offender's expected payoff? Third (RQ3), which attack paths are reduced by an owner-bound passive substrate, and which remain in the end-to-end custody system? Last (RQ4, how does value partitioning affect expected loss, variance, and catastrophic-loss probability when compromise events are correlated? The paper makes two connected methodological contributions. First, it specifies measurable loss and enforcement quantities and tests the available public evidence against their identification requirements. Second, it develops a bounded reference architecture for owner-bound offline custody. The architectural invariant is that no complete spend key is present on the physical substrate or retained by the manufacturing process. An owner-controlled factor is required for authorization, and transaction signing occurs in a verified offline environment.

The analysis does not claim that all digital value is riskier than all cash, that passive storage is immune to compromise, or that a physical cryptographic factor satisfies regulatory obligations by itself. Its contribution is a falsifiable measurement framework and an engineering specification whose benefits and residual risks can be independently tested. The paper is structured as follows. Section 2 presents literature review and theoretical foundations. Section 3 explains the methodology and data. Section 4 provides evidence on loss and enforcement, followed by section 5 on owner-bound offline custody architecture, section 6 on partitioning and common-mode risk, and section 7 on threat model and comparative security. Sections 8 and 9 analyse macroeconomic, regulatory implications and limitations, respectively. Section 10 concludes the paper.

2. Literature and Theoretical Foundations

2.1 Attack topology, realized loss, and measurement

An attack surface is the set of reachable components, interfaces, dependencies, and human processes through which an adversary may attempt compromise. Realized loss is the result of attempts, controls,

detection, response, recovery, and reporting. A broad attack surface can coexist with low observed loss when controls are effective; a narrow surface can produce severe loss when value is concentrated behind a weak control.

Research on cryptocurrency security has documented that technical, economic, and operational vulnerabilities interact rather than exist as isolated software defects (Bonneau et al., 2015; Conti et al., 2018). Table 1 summarizes five dimensions that are especially relevant to monetary substitutes. Perez and Livshits (2021) distinguish detected smart-contract vulnerabilities from observed exploitation, supporting this separation of attack surface and realized loss.

These are tendencies, not categorical properties. A cash vault may hold a large concentration, while digital assets can be segregated and held offline. The relevant unit of comparison is a defined custody configuration under a defined threat class.

Table 1. Comparative attack-topology dimensions

Dimension	Physical-cash tendency	Networked digital-asset tendency	Measurement implication
Reach	Local or insider access	Potentially global and remote	State geography and custodian perimeter
Replication	New physical effort for each target	Exploit logic may be reused	Model clustered incidents and fixed costs
Concentration	Bounded by location and transport	Shared keys, bridges, or dependencies may aggregate value	Use heavy-tail and common-mode diagnostics
Transfer	Physical removal and laundering	Rapid, multi-network movement	Separate gross loss, recovery, and ultimate net loss
Observability	Police and insurer records; underreporting remains	On-chain visibility with off-chain attribution gaps	Document different capture and classification processes

2.2 Economics of offending

Becker's (1968) framework treats offending as a choice under uncertainty. For an attempt against target, let denote the probability of technical success conditional on an attempt; the probability of an enforceable sanction conditional on success; the transferable gain; the private execution and opportunity cost; the monetary equivalent of punishment; restitution, forfeiture, or asset recovery imposed on the offender; and baseline wealth. Treating the sanction amounts as fixed within each outcome, let $U(\cdot)$ denote utility of wealth and EU_j expected utility:

$$EU_j = q_j\left[(1-p_j)U(Y+G_j-c_j)+p_jU(Y+G_j-c_j-F_j-R_j)\right]+(1-q_j)U(Y-c_j) \quad (1)$$

For a risk-neutral offender, π_j denotes expected monetary payoff relative to baseline wealth; Equation (1) reduces to

$$\pi_j = q_j G_j - c_j - q_j p_j (F_j + R_j) \tag{2}$$

An economically motivated attempt is discouraged in this simplified model when $\pi_j < 0$. The condition shows why sanction probability alone is not a sufficient measure of deterrence. Technical success, target value, execution cost, penalty severity, recovery, liquidity, and risk preferences also matter. A custody control may reduce expected offending by lowering q_j, reducing the value exposed to one compromise, increasing c_j, or improving attribution and recovery. This simplified tree excludes sanctions after failed attempts and recovery without sanction; Section 6.1 separates recovery from conviction.

Network attacks may also exhibit an important scale economy. Suppose an adversary incurs fixed development cost C_0 and attacks M targets, indexed by m, with marginal costs c_m. For each target, q_m, G_m, p_m, F_m, R_m retain the definitions above; $\Pi(M)$ is total expected net payoff. Under risk neutrality,

$$\Pi(M) = -C_0 + \sum_{m=1}^{M} [q_m G_m - c_m - q_m p_m (F_m + R_m)] \tag{3}$$

Reusable exploit logic amortizes C_0 across targets. The marginal cost of an additional attempt may therefore be small, although credential acquisition, laundering, and operational security can remain costly. Physical attacks more often require target-specific access and transport, but insider networks or shared logistics can also reduce marginal cost. The model identifies a testable mechanism rather than assuming a universal ordering.

2.3 Enforcement as a pipeline

Cybercrime enforcement is not a single probability. Reporting, investigative acceptance, attribution, identification, arrest, charging, prosecution, conviction, and recovery are distinct stages with different units and censoring. A case may be cleared by exceptional means without a conviction; one event may involve several defendants; and recent cases may remain unresolved at the end of the observation window.

For a cohort of eligible events n_0, let n_k be the number reaching enforcement stage k. The observable transition rate is

$$\hat{\theta}_k = \frac{n_k}{n_{k-1}}, \qquad k = 1, \ldots, K \tag{4}$$

and the cohort probability of reaching the final recorded stage is

$$\hat{p}_{0 \to K} = \prod_{k=1}^{K} \hat{\theta}_k = \frac{n_K}{n_0} \tag{5}$$

These equalities require consistent event linkage and a sufficiently mature cohort. Right-censoring, jurisdictional transfer, multiple defendants, and missing outcomes otherwise bias the estimate. Survival or multistate models are preferable when stage dates are available.

A probability-only deterrence ratio may be defined for matched physical and digital offense cohorts,

$$DR_{c,d}^{(p)} = \frac{p_c}{p_d}, \qquad p_d > 0 \tag{6}$$

but it is descriptive, not a complete deterrence measure. A fuller expected-sanction component is $D_j = p_j(F_j + R_j)$. Because public data do not provide matched digital and cash cohorts followed through equivalent stages, this paper does not report a cross-domain point estimate for $DR_{c,d}^{(p)}$.

2.4 Offline custody and secret sharing

Offline storage can reduce the set of remotely reachable components, but it does not remove key-generation, parser, firmware, supply-chain, recovery, or coercion risk. A full-key paper wallet is particularly fragile because copying the record duplicates spending authority. Secret sharing changes that property: a single substrate can be made insufficient for authorization, provided the complementary share is independently generated or controlled and the reconstruction process is trustworthy (Shamir, 1979).

Secret sharing is not threshold signing. In a secret-sharing system the signing secret may be reconstructed, even if only briefly, on one device. A threshold-signature protocol distributes signing so that the complete key need not be assembled. The latter can reduce reconstruction exposure but introduces protocol, implementation, availability, and interoperability risks. The reference architecture in this paper uses secret sharing because it can be specified across existing transaction systems; threshold signing is an important extension.

3. Data and Methodology

3.1 Research design and scope

The empirical component is a methodological feasibility study focused on USDC and USDT. It defines estimands, assesses comparability, and provides an illustrative U.S. reported-cash-loss calibration, not a causal estimate of digitization. Table 1a distinguishes the asset classes (Bank for International Settlements, 2025b). The engineering component is an analytical reference specification, not a product certification or measured reliability study.

Table 1a. Asset taxonomy and study perimeter

Asset class	Economic distinction	Study perimeter
Physical cash	Tangible bearer currency	U.S. dollar calibration; future matched comparator
Fiat-backed stablecoins	Issuer-backed tokens linked to fiat value	USDC and USDT only; separate strata
Other cryptoassets	Includes issuerless Bitcoin and non-fiat-backed tokens	Excluded from target cohort
Tokenized commercial-bank money	Tokenized claims on a commercial bank	Excluded; distinct liability and custody model
CBDCs	Digital central-bank liabilities; designs vary	Excluded; offline design context only

The target cohort is issuer-native USDC and USDT held in externally controlled wallets, stratified by token contract, network, custody configuration, and key-compromise class. Wrapped or bridged claims, DeFi deposits, issuer-administration incidents, and other assets require separate exposure pools. Issuer terms, freezing powers, redemption eligibility, and chain controls differ (Circle, 2025; Tether, n.d.); USDC and USDT are not interchangeable observations. Future estimation should compare each stratum with physical U.S.-dollar cash held by the same institutional population over the same period.

The unit of analysis for loss should be a recorded theft event, not a complaint, transfer, wallet address, or defendant. An event enters the theft numerator only when an unauthorized transfer or legally reported taking occurred. Investment fraud, ransomware payments, market-price declines, issuer insolvency, and protocol losses without unauthorized taking must be reported separately.

3.2 Loss estimands

Let e index recorded theft events in domain j during period t. Let g_e be gross value stolen at a stated valuation timestamp and $r_e(h)$ value returned or recovered within horizon h. Gross and horizon-specific net loss are

$$G_{j,t} = \sum_{e \in \mathcal{E}_{j,t}} g_e\,, \qquad N_{j,t}(h) = \sum_{e \in \mathcal{E}_{j,t}} [g_e - r_e(h)] \tag{7}$$

Let $V_j(u)$ denote value exposed to the included event class at time u. For an observation interval of length T, time-weighted exposure is

$$E_{j,t} = \frac{1}{T}\int_t^{t+T} V_j\,(u)\,du \tag{8}$$

Gross and net exposure-normalized loss intensities are

$$\hat{L}^G_{j,t} = \frac{G_{j,t}}{E_{j,t}}, \qquad \hat{L}^N_{j,t}(h) = \frac{N_{j,t}(h)}{E_{j,t}} \tag{9}$$

Equation (9) defines value-weighted intensities, not the probability that a particular currency unit is stolen. A probability interpretation would require a closed exposure pool, a defined risk interval for every unit, homogeneous hazards or known risk weights, no repeated exposure, and complete event observation. Those conditions do not hold in the aggregate sources used here.

For matched domains a and b, a Relative Vulnerability Ratio is

$$RVR_{a,b} = \frac{\hat{L}^G_{a,t}}{\hat{L}^G_{b,t}} \tag{10}$$

The ratio supports inference only when numerator and denominator share the period, geography, asset perimeter, custody perimeter, event definition, gross or net convention, valuation rule, exposure concept, and materially comparable reporting coverage. Otherwise, it is a scenario calibration conditional on chosen inputs. No cross-domain RVR point estimate is reported because the available data fail these conditions.

3.3 Source hierarchy and metadata

Public agencies and central banks are preferred for fiat and enforcement statistics. Commercial blockchain-intelligence estimates are useful for digital-asset incident discovery and classification, but they require triangulation because event lists and taxonomies are proprietary and may change. Table 2 records the sources, coverage, required metadata, and admissible use.

The minimum event-level metadata for comparative estimation are: event identifier; loss timestamps; victim and custodian jurisdictions; asset and chain; custody type; attack vector; event-time valuation; gross loss; recovered value and horizon; insurance reimbursement; reporting source; inclusion confidence; affected accounts; time-weighted exposed value; and dated indicators for referral, attribution, identification, arrest, charging, prosecution, disposition, and recovery. Dataset version, source retrieval date, transformation code, and a cryptographic hash of each input file should be retained.

3.4 Bias and statistical uncertainty

At least four biases limit cross-domain comparison. First, underreporting differs: small cash thefts may never reach police, while large on-chain transfers are observable but may be misclassified. Second, denominators differ: currency in circulation includes notes held abroad and outside theft-prone locations, while stablecoin supply spans heterogeneous custody arrangements. Third, digital losses are heavy-tailed and clustered around common dependencies. Fourth, gross and ultimate net loss diverge as assets are frozen, returned, reimbursed, or written off. Heterogeneous loss severity (Eling & Jung, 2022) and differential underreporting (Sangari et al., 2022) reinforce the need for stratified exposure and reporting models.

Event lists should be deduplicated across vendors, issuer disclosures, court records, and law-enforcement releases. Capture–recapture methods may help assess omitted events only if dependence between lists and heterogeneous inclusion probabilities are modeled. Conventional independent Poisson intervals are unsuitable when a single infrastructure compromise produces several linked losses. With event-level observations, uncertainty should be estimated using block or cluster bootstrap procedures that preserve common-event structure. No confidence interval is reported here because neither matched microdata nor a defensible digital exposure denominator is publicly available.

Table 2. Evidence sources and analytical use

Construct	Source and dataset	Coverage and unit	Required metadata	Use in this study
Cash stolen and recovered	FBI, Crime in the United States 2019, Table 24	USD reported by 14,165 agencies; population 285.7 million	Agency coverage, offense rules, gross and recovered values	Illustrative reported-loss calibration
Currency stock	Federal Reserve, Currency in Circulation: Value, 2019	USD billions at year-end; includes notes abroad	Date, denomination, foreign-holdings caveat	Approximate stock denominator

Construct	Source and dataset	Coverage and unit	Required metadata	Use in this study
Property-crime clearance	FBI UCR, Clearances, 2019	Offenses cleared by arrest or exceptional means	Definition, category, agency coverage	Conceptual benchmark only
Bank-crime topology	FBI, Bank Crime Statistics 2023	1,362 reported violations at federally insured institutions	Violation type, location, persons identified	Attack topology; no loss denominator
Crypto hacks	Chainalysis, 2024 hacking update	USD stolen; 303 identified incidents	Incident rules, prices, recoveries, version date	Broader cryptoasset context only
Crypto hacks	TRM Labs, 2025 Crypto Crime Report analysis of 2024	USD stolen in hacks and exploits	Incident list, confidence rules, taxonomy	Broader cryptoasset cross-check only
Internet-crime complaints	FBI, 2024 IC3 Annual Report	Complaints and reported losses	Crime label, crypto descriptor, duplicate policy	Complaint burden only
Stablecoin scale and flows	IMF Working Paper WP/25/141; BIS Annual Economic Report 2025	Market stock and estimated flows	Chain coverage, price, address and geography rules	Macro-financial context
Binance breach, 29 October 2026	Binance (2026a, 2026b); Chainalysis (2026); Robinson (2026); TRM Labs (2026); FBI (2026)	Single custodian; USD 6.30–6.43 billion gross estimates	Source cutoff; token and chain; prices; freeze, seizure, and restoration status	Mechanism illustration; no probability estimate or model calibration

4. Evidence on Loss and Enforcement

4.1 Illustrative U.S. reported-loss calibration

FBI Table 24 reports USD 1,423,559,757 in stolen "currency, notes, etc." and USD 36,980,933 recovered in 2019. The table covers 14,165 agencies and an estimated population of 285,731,352 (Federal Bureau of Investigation, 2019b). The Federal Reserve reports USD 1,759.8 billion in U.S. currency in circulation at December 31, 2019 (Federal Reserve Board, 2026). Using this year-end stock as an illustrative reported-loss calibration denominator gives

$$\hat{L}^{G}_{cash,2019} = \frac{1{,}423{,}559{,}757}{1{,}759{,}800{,}000{,}000} = 0.00080893 = 0.080893\% \tag{11}$$

The reported recovery share is

$$\widehat{RR}_{cash,2019} = \frac{36{,}980{,}933}{1{,}423{,}559{,}757} = 0.025978 = 2.5978\% \tag{12}$$

The corresponding net reported loss intensity is

$$\hat{L}^{N}_{cash,2019} = \frac{1{,}423{,}559{,}757 - 36{,}980{,}933}{1{,}759{,}800{,}000{,}000} = 0.00078792 = 0.078792\% \tag{13}$$

Table 3 consolidates the source quantities and the three descriptive calculations.

Table 3. Illustrative U.S. reported currency-loss calibration, 2019

Quantity	Value	Interpretation
Reported gross currency stolen	USD 1.424 billion	Covered agencies and offenses in FBI Table 24
Reported currency recovered	USD 36.981 million	Recovery recorded within the source's reporting process
Gross reported loss intensity	0.080893%	Gross loss divided by year-end currency stock
Recovery share	2.5978%	Recovery divided by reported gross loss
Net reported loss intensity	0.078792%	Reported loss net of reported recovery, divided by stock

These are illustrative reported-loss calibrations, not theft probabilities or formal bounds on population risk. The numerator excludes unreported losses and agencies outside the table's coverage. The denominator includes currency held outside the United States and is a year-end stock rather than average value exposed to theft. The calculation does not estimate the probability that an individual banknote is stolen.

FBI Bank Crime Statistics provide complementary evidence about topology. In 2023 the FBI recorded 1,362 violations under the Federal Bank Robbery and Incidental Crimes Statute (18 U.S.C. § 2113): 1,263 robberies, 95 burglaries, and four larcenies. Fifty-four incidents involved a vault or safe, and investigators had identified 801 of 1,652 persons known to be involved at publication (Federal Bureau of Investigation, 2023). The report does not provide an aggregate stolen-value numerator and cautions that it is not a complete census. It can describe event type and concentration, but it cannot support a vault-theft loss intensity. The four larcenies refer only to this recorded category, not all U.S. cash theft or bank-related theft; no undercount multiplier is inferred.

4.2 Contextual cryptoasset hacking evidence

Chainalysis (2024) estimated approximately USD 2.2 billion stolen from crypto platforms across 303 hacking incidents during 2024 and attributed 43.8% of stolen value to private-key compromise. TRM Labs (2025) independently reported approximately USD 2.2 billion stolen in hacks and exploits and classified nearly 70% as infrastructure attacks, primarily private-key and seed-phrase compromises. Agreement in the aggregate value is useful triangulation. The attribution percentages are not directly comparable because the firms use different categories, incident sets, and confidence rules.

These estimates cover broader cryptoassets and platforms, not USDC and USDT alone. They provide context, not a stablecoin-theft numerator. Neither stablecoin supply nor total crypto market capitalization matches the reported attack perimeters. A USDC or USDT intensity requires token-specific theft events and time-weighted value exposed within the same network and custody configuration.

The FBI's 2024 Internet Crime Complaint Center report measures a different phenomenon: 859,532 complaints and USD 16.6 billion in reported losses, with a cryptocurrency descriptor attached to 149,686 complaints and USD 9.322 billion in reported loss (Federal Bureau of Investigation, 2025). The descriptor identifies a medium or tool associated with a selected crime type. It includes investment fraud and other conduct, so it must not be added to hack totals or interpreted as a count of technical compromises, arrests, or convictions.

A post-acceptance update records evidence available through 1 November 2026. Binance (2026a) reported approximately USD 6.30 billion lost in the 29 October hot-wallet breach. TRM Labs (2026) valued unauthorized outflows at approximately USD 6.36 billion across ten networks, Chainalysis (2026) estimated USD 6.40 billion, and Elliptic reported USD 6.43 billion across eleven networks (Robinson, 2026). These are separate source estimates. Coverage, address attribution, valuation timestamps, and update times must be retained. The FBI attributed the theft to North Korea's TraderTraitor actors on 31 October (Federal Bureau of Investigation, 2026). However, attribution does not establish conviction or completed recovery.

Within the study's asset scope, TRM Labs (2026) identifies 1.875 billion USDT and 823 million USDC taken on Ethereum and 612.7 million USDT on Tron. Its chronology records freezes of 971.1 million USDT at 21:01 UTC and 823 million USDC at 21:04 UTC on 29 October. These are address-specific observations. Wrapped tokens such as USDT0 and USDC obtained by converting other stolen assets require separate classification, and subsequent laundering transfers must not be counted as additional theft. Elliptic's later snapshot verified USD 2.85 billion frozen or seized across assets at approximately 05:50 UTC on 1 November, using current prices (Robinson, 2026). Subtracting this from event-time gross loss would mix valuation conventions and still not measure victim restoration.

Chainalysis (2026) describes approximately USD 325 million in stablecoins exchanged for USD 285 million in ETH over almost 20 minutes, illustrating the incentive to sacrifice proceeds for assets beyond issuer freezing powers. Binance (2026b) subsequently reported covering approximately 95% of the impacted amount using company resources as it prepared phased withdrawals on 1 November. Customer compensation, assets seized into recovery custody, and funds actually restored from the theft must be recorded separately. Neither the custodian's total holdings nor the multi-asset loss supplies the matched USDC/USDT exposure denominator required by Equation (9).

4.3 What can be inferred

Three findings are supportable. First, digital-asset theft is economically material and key-management or infrastructure compromise constitutes a substantial share in both vendor taxonomies. Second, the attack topology permits remote reach, exploit reuse, and correlated loss, which makes isolation and compartmentalization relevant design objectives. Third, the sources examined do not identify a USDC/USDT-specific per-dollar theft probability, a cash-versus-digital RVR, or a cross-domain deterrence ratio.

The absence of a valid point estimate does not imply equal risk. It means that the magnitude and uncertainty of the difference remain unidentified. Policy should distinguish an observed mechanism from an estimated population parameter.

4.4 The enforcement evidence gap

FBI UCR clearance statistics define clearance by arrest or exceptional means, not conviction (Federal Bureau of Investigation, 2019a). Bank Crime Statistics report people identified in investigations, not disposition. IC3 reports complaints and losses, not a complete enforcement cohort. Blockchain-intelligence reports focus on incident attribution, and fund flows rather than the full criminal-justice process.

A credible comparison would begin with matched event cohorts and follow each event through the stages in Equations (4) and (5). The analysis would define the jurisdictional rule for multi-country events, keep event-level and defendant-level tables separate, and use stage dates to account for unresolved cases. Penalty severity and asset recovery would then enter Equation (2) rather than being collapsed into a conviction percentage. Until such data exist, “deterrence deficit” is best used as a hypothesis about the combined effects of low attribution, cross-border friction, rapid transfer, limited recovery, and scalable attacks.

5. Owner-Bound Offline Custody Architecture

5.1 Security objective and invariant

The architecture is intended for assets or authorities that are used infrequently and justify reduced online availability in exchange for a smaller remote attack surface. The application considered here is segregated USDC or USDT treasury custody. It is not intended to replace every operational wallet.

The mandatory invariant is:

No complete spend key exists on the physical substrate or in manufacturing logs, and authorization requires an owner-controlled factor r45independent of the substrate.

The physical record is therefore not a bearer instrument for the full balance. Copying or stealing it should not confer unilateral spending authority. This property depends on genuine independence: the complementary share must not be generated, retained, backed up, or recoverable by the same manufacturing or custody domain.

The system must also preserve availability. Share loss, operator incapacity, succession, and substrate damage require a documented recovery design. Threshold selection, duplicate storage, revocation, and disaster recovery are governance choices, not properties supplied automatically by cryptography.

5.2 Record and protocol binding

Each physical record should carry a versioned, authenticated envelope containing only public parameters and one secret share. A minimum envelope includes protocol and suite identifiers, network and asset identifiers, threshold and share index, public address or public-key commitment, derivation-path information when applicable, encoding and error-detection parameters, record serial number, creation epoch,

recovery-policy identifier, and authenticated provenance metadata. Personal identity data should not be printed unless a documented legal and privacy basis requires it.

Secret confidentiality does not provide record integrity. The envelope should therefore include a digital signature or message-authentication mechanism whose verification key is controlled outside the printer. Error-detection codes address accidental corruption but not malicious substitution. Unknown versions, duplicate share indices, malformed fields, inconsistent commitments, and unsupported algorithms must fail closed.

Before funding, an independent verification device should decode the record, combine the required owner factor in a controlled environment, derive the public address, and compare it with the printed commitment and the intended destination. This ceremony is necessary because a malicious printer can substitute an attacker-controlled address while producing a visually plausible record.

5.3 Two-of-two XOR sharing

For a uniformly distributed 256-bit secret K, select an independent uniform share

$$S_1 \xleftarrow{\$} \{0,1\}^{256}, \qquad S_2 = K \oplus S_1 \tag{14}$$

Reconstruction is

$$K = S_1 \oplus S_2 \tag{15}$$

Either share alone is information-theoretically independent of K when S_1 is uniform, secret, and used once. XOR is appropriate only for a strict two-of-two design. Reusing a share across keys, deriving both shares within an untrusted printer, or retaining a copy in logs defeats the security claim. The share format must specify bit ordering, length, encoding, checks, and protocol version.

5.4 Threshold Shamir sharing

For a t-of-n scheme over a finite field $\mathbb{F}$, encode the secret K as a field element and sample independent coefficients $a_1, \ldots, a_{t-1}$. Define

$$f(x) = K + \sum_{j=1}^{t-1} a_j x^j, \qquad a_j \xleftarrow{\$} \mathbb{F} \tag{16}$$

Each share is (x_i, y_i), where $x_i \neq 0$, all indices are distinct, and $y_i = f(x_i)$. Any t valid shares recover

$$K = f(0) = \sum_{i=1}^{t} y_i \prod_{\substack{\ell=1 \\ \ell \neq i}}^{t} \frac{x_\ell}{x_\ell - x_i} \tag{17}$$

All operations in Equation (17) occur in the selected field. In characteristic two, subtraction equals addition, but implementations should retain the field definition rather than substitute integer or bitwise arithmetic informally.

A profile using $\mathrm{GF}(2^{256})$ must name a vetted irreducible degree-256 polynomial, define polynomial-basis encoding and byte order, specify element and index validation, and publish known-answer and malformed-input tests. No polynomial should be selected merely because it appears syntactically plausible. Implementers may instead use a well-reviewed prime field with an unambiguous secret encoding. The chosen suite must be versioned and independently reviewed.

Shamir sharing provides confidentiality against fewer than t shares, not authenticity. Robust reconstruction requires authenticated envelopes, unique indices, commitment verification, duplicate detection, and explicit error handling. For high-value deployments, verifiable secret sharing or a threshold-signature design may be preferable.

5.5 Key generation, encryption, and signatures

Key material should be generated from an approved deterministic random bit generator seeded by documented entropy sources, with health tests and failure handling consistent with NIST SP 800-90B (National Institute of Standards and Technology, 2018). Key lifecycle and destruction procedures should follow a documented security strength and retention policy (National Institute of Standards and Technology, 2020). The ceremony should combine independent entropy where feasible and retain non-secret evidence that the procedure and software version were followed. Entropy values, complete keys, and shares must not enter telemetry, crash dumps, printer queues, swap, or general-purpose logs.

Passphrase-based encryption can add a factor but should not replace sharing or access control. BIP 38 is a Bitcoin-specific legacy construction based on scrypt and AES (Bitcoin Improvement Proposals, 2012); it is not a general encryption profile for stablecoins or arbitrary chains. A contemporary design should specify a versioned authenticated-encryption suite, memory-hard key derivation, salt, work factors, encoding, and migration rules. Passphrase entropy and recovery procedures must be tested because encryption cannot compensate for a guessable secret.

Signature rules are chain-specific and should be evaluated against the applicable signature standard (National Institute of Standards and Technology, 2023). Ed25519 verification must enforce canonical encodings and the $S < L$ condition specified by RFC 8032 (Josefsson & Liusvaara, 2017). Bitcoin's BIP 66 constrains DER encoding (Bitcoin Improvement Proposals, 2015), while low-s normalization follows the applicable transaction-policy and consensus rules. Ethereum's EIP-2 applies its own low-s rule (Buterin, 2015). These controls reduce malleability or ambiguity; they do not make transactions secret or prove that a human intended the displayed transfer.

5.6 Air-gapped ingestion and signing

The online coordinator constructs an unsigned transaction or signing request containing only public data. It displays the request through a constrained optical format. The offline device decodes it with a bounded, fail-closed parser, independently derives the transaction identifier or canonical preimage, and presents human-readable destination, asset, amount, fee, network, nonce, and policy effects on a trusted display. Authorization proceeds only after policy and operator checks.

TRM Labs (2026) reports that the Ethereum component used 50 pre-signed EIP-7702 authorizations with consecutive nonces to delegate a Binance wallet to attacker-controlled code. EIP-7702 authorizes

persistent account behavior, so checking a displayed transfer amount alone cannot establish the scope of consent (Buterin et al., 2024). For this reference profile, delegation requests, including batches spanning multiple nonces, should be rejected by default. Any supported exception requires a separately approved policy, verified delegate code, explicit review of ongoing authority and revocation, and tests against unauthorized requests. Air-gapping alone cannot make a malicious authorization safe.

The required shares are introduced inside the controlled offline environment. The signer reconstructs or accesses the key for the shortest feasible interval, signs under the chain-specific profile, erases secret-bearing memory, and emits only the signed public transaction and a minimal non-secret receipt. The online system decodes the output again, compares all fields with the authorized request, and broadcasts it.

An optical channel is an interface control, not proof of end-to-end isolation. QR decoders, cameras, displays, firmware, and peripherals can contain vulnerabilities or covert channels. Input size, character set, schema, image complexity, and the number of accepted messages should be bounded. Output should be deterministic or strictly allowlisted so that a malicious signer cannot exfiltrate shares through unused fields or visual variation. Signed transactions are public artifacts, not zero-knowledge proofs.

Secret zeroization should be implemented with platform-specific primitives resistant to compiler removal and tested across normal completion, error paths, power loss, and crash recovery. Some caches, managed runtimes, storage controllers, or secure elements may prevent complete erasure. The security assessment must state those residuals rather than treating a zeroization call as proof.

6. Partitioning and Common-Mode Risk

Partitioning limits value exposed to one compromise only when compromise events are sufficiently independent. Let total value W be divided into N equal compartments. Let $X_i \in \{0,1\}$ indicate loss of compartment i, with $\Pr(X_i = 1) = p$. Total loss is

$$L = \frac{W}{N}\sum_{i=1}^{N} X_i\,, \qquad \mathbb{E}[L] = pW \tag{18}$$

If per-compartment compromise probability remains p, partitioning does not change expected loss. Under pairwise correlation ρ,

$$\mathrm{Var}(L) = W^2 p(1-p)\left[\frac{1}{N} + \rho\frac{N-1}{N}\right] \tag{19}$$

When $\rho = 0$, variance falls with N. When $\rho = 1$, partitioning provides no variance benefit. The design objective must therefore reduce shared dependencies as well as nominal balance size.

An explicit common-mode model makes the limitation clearer. Let a common event compromise every compartment with probability q. Conditional on no common event, assume independent compartment compromise probability p. Then

$$\Pr(L = W) = q + (1-q)p^N \tag{20}$$

Increasing N reduces only the second term. Shared firmware, one provisioning ceremony, one recovery administrator, one entropy source, or one physical location can dominate the residual risk.

A practical portfolio objective should include expected loss, tail risk, operating cost, and common-mode controls:

$$\min_N \mathbb{E}[L_N] + \lambda\,\mathrm{CVaR}_\alpha(L_N) + C(N) \tag{21}$$

subject to recovery-time, quorum, geographic-diversity, and policy constraints. The parameter λ expresses institutional aversion to tail loss, and $C(N)$ includes materials, ceremonies, monitoring, testing, and recovery exercises. Estimating the objective requires incident and operational data; Equation (21) is a decision framework, not an empirical optimum.

6.1 Linking attacker incentives and net loss

Let z denote custody controls and w_i the fixed value in compartment i. For full-compartment theft, set $G_i = w_i$ in Equation (3) and compare candidate target sets with the zero payoff from not attacking. Controls alter fixed development cost, marginal cost, technical success, and sanction/recovery terms. Target selection therefore affects compromise probability; holding it constant as compartment size changes is an assumption, not a result.

Over a fixed period with no replenishment or repeated counting, let $X_i \in \{0,1\}$ indicate compartment loss and $u_i(z) = Pr(X_i = 1|z)$ its probability, jointly determined by attack selection and success. Let $r_i(h) \in [0,1]$ be the recovered fraction by horizon h, and $\overline{r}_i(h,z) = E[r_i(h)|X_i = 1, z]$, set to zero for zero-probability losses. Net loss L_h and its expectation are

$$L_h = \sum_{i=1}^{N} w_i X_i [1 - r_i(h)], \quad E[L_h|z] = \sum_{i=1}^{N} w_i u_i(z)[1 - \overline{r}_i(h,z)] \tag{21a}$$

Writing for the net loss fraction, dependence enters throughYi=Xi[1-ri(h)]

$$Var(L_h|z) = \sum_{i=1}^{N}\sum_{j=1}^{N} w_i w_j Cov(Y_i, Y_j|z) \tag{21b}$$

No independence between compromise and recovery is assumed. Recovery may occur without conviction; freezing USDC or USDT is not recovery until value is restored to the victim. Correlation changes variance and tail loss directly, but mean loss only through marginal compromise or recovery terms. Equal balances, no recovery, and the common probability/correlation assumptions recover Equations (18)–(19); common-mode events enter through the joint distribution. Equation (21) can therefore use net loss alongside control costs. These are analytical identities, not estimated effects.

The Binance case illustrates these relationships without calibrating them. Binance (2026a) reported additional breach points and no evidence of private-key leakage. That statement does not establish that leakage was impossible. Cross-chain unauthorized signing is consistent with shared control dependencies, but the complete attack path remained unresolved at the snapshot date. One event cannot identify the common-event probability in Equation (20), pairwise correlation in Equation (19), or the fixed and marginal costs in Equation

(3). Moreover, repeated sweeps of newly arriving deposits violate the fixed, unreplenished exposure assumption of Equation (21a); applying it requires explicitly defined exposure intervals. Freezing can constrain realizable offender proceeds but expected deterrence and victim recovery require separate observations.

7. Threat Model and Comparative Security

7.1 Threat classes

Table 4 states the control objective and residual risk for each principal threat class. Elimination is boundary-specific: a passive substrate has no network interface, so direct network exploitation of that component is absent. This does not eliminate an end-to-end route to unauthorized spending. Attack effort may shift to provisioning, optical parsing, key reconstruction, recovery, or operator approval. Requiring an independent share blocks substrate-only spending under the stated assumptions, but transfers reliance to share independence and the signer. “Zero remote exploitability” and “unhackable” are therefore inappropriate system descriptions.

Table 4. Threat-control-residual matrix

Threat	Principal control	Residual risk
Remote malware against the passive record	No networked component on the substrate	Coordinator or offline signer may be compromised
Theft or copying of the substrate	Independent owner factor; no complete key printed	Coercion or simultaneous factor theft
Malicious printer or provisioning station	Owner entropy, signed metadata, independent pre-funding verification	Supply-chain collusion or compromised verifier
Entropy failure or key reuse	Approved generator, health tests, independent entropy	Correlated or undetected entropy defects
Parser exploit	Canonical schema, strict bounds, fail-closed parsing, fuzzing	Unknown parser or display vulnerability
Transaction substitution	Trusted display, policy engine, independent output comparison	Human approval error or display-path compromise
Side-channel or fault attack	Constant-time software, shielding, fault detection, testing	Capable laboratory adversary
Firmware rollback	Signed measured boot and monotonic version policy	Vendor-key compromise or immutable boot defect
Optical exfiltration	Output allowlist and deterministic encoding	Subliminal channels or malicious peripherals
Share loss or damage	Threshold recovery, geographic separation, drills	Common-mode disaster or hidden dependency
Coercion	Value limits, time delays, multi-person control, duress procedure	Cryptography cannot eliminate physical coercion
Supply-chain substitution	Chain of custody, provenance, sampling, teardown	Sophisticated component or material substitution

7.2 Comparison with custody modes

Table 5 is a control-level analytical comparison of specified custody configurations, not evidence of product or architecture superiority. Each criterion identifies an attack path or operational dependency; entries are conditional on the stated controls and are not measured failure rates.

The substrate-only copying property is an architectural condition, not a demonstrated reduction in total loss. Hardware-wallet and threshold-signature configurations may provide comparable or stronger controls. Testing should verify whether copied artifacts or backups permit unilateral spending, malformed requests alter signed intent, and provisioning or recovery compromises cross compartment boundaries. Issuer, contract, and network failures remain outside this storage-control comparison. Independent testing is required before claiming comparative effectiveness.

Table 5. Control-level analytical comparison of custody configurations

Attack-path criterion	Hot wallet	Hardware wallet	Full-key paper wallet	Owner-bound substrate
Remote extraction during storage	Online host or signing service	Host interface; device isolation required	No paper interface; generation/import remain	No substrate interface; provisioner/signer remain
Does copying the stored artifact suffice?	If sufficient credentials are exposed	Only if key or recovery material is exposed	Yes, for an unencrypted full key	No, if the other share is independent
Malicious input or firmware path	Online software stack	Host request, device and firmware	Generation, printing and import	Provisioning and offline parser/signer
Transaction availability	High	Moderate	Low	Low to moderate
Recovery bypass of authorization	Account reset or backup policy	Seed backup and recovery path	Each readable copy holds full key	Threshold or owner-factor reset policy
Coercion at authorization	Limited resistance	Limited resistance	Limited resistance	Policy dependent; not eliminated
Shared-domain compromise	Shared host, service or key	Shared firmware or backup	Shared generator or storage	Shared provisioner, signer or recovery

7.3 Hardware and environmental qualification

No secure-element vendor is normative. Candidate devices should be evaluated for secure boot, code authenticity, anti-rollback, key isolation, entropy behavior, side-channel and fault resistance, debug-port lifecycle, vulnerability disclosure, update policy, and configuration-specific certification. Certification is evidence about a defined target and configuration, not proof of system security.

Durability is a property of the complete record stack: substrate, pigment or conductor, adhesive, laminate, optical geometry, encoding redundancy, and storage environment. Qualification should define sample size, production batch, environmental profile, acceptance threshold, failure definition, confidence interval, and any acceleration model.

A test program should include baseline and aged optical grading under ISO/IEC 15415; ultraviolet exposure under an applicable ISO 4892-3 method; water absorption or immersion testing appropriate to the material system, including ASTM D570 where applicable; tensile testing such as ASTM D882 for thin plastic sheeting; thermal and humidity cycling under an appropriate IEC 60068 profile; and chemical, abrasion, fold, puncture, salt-fog, and fire exposure when required by the use case. Every ageing condition should conclude with optical decoding, envelope authentication, and cryptographic reconstruction tests. Material names alone do not establish moisture immunity, ultraviolet protection, or multi-decade readability.

8. Macroeconomic and Regulatory Implications

8.1 USDC and USDT custody implications

The IMF documents large international stablecoin flows and the difficulty of assigning pseudonymous activity to geography (Reuter, 2025). BIS analysis places stablecoins within a broader transformation of the monetary system while identifying governance, monetary, and operational concerns (Bank for International Settlements, 2025a, 2025b). These sources establish scale and cross-border relevance; they do not supply a stablecoin theft rate.

Beyond holder custody, reserve-control keys, mint and burn authorities, bridge administrators, upgrade keys, and recovery credentials are potential concentration points. Compromise can create correlated losses or undermine confidence beyond the directly stolen amount. Owner-bound offline controls may therefore serve as one defense layer for rarely used authorities and contingency reserves. They do not address reserve credit quality, redemption liquidity, governance capture, smart-contract defects, oracle manipulation, consensus failure, or legal claims on reserves.

The October 2026 Binance breach highlights why the preventive authorization controls discussed above must be considered alongside containment and loss absorption. Issuer freezing can interrupt disposal of stolen USDC or USDT, while custodian reserves can restore customer access before stolen assets are recovered. The ability of one custodian to absorb losses does not establish sector-wide resilience. The Binance breach therefore motivates evaluation of independent authorization for rarely used treasury compartments alongside operational liquidity and recovery arrangements. It does not establish that the proposed owner-bound architecture would have prevented the attack.

The macro-financial relevance arises from common-mode dependency and recovery capacity, not from the physical medium alone. A sovereign-grade control must have explicit assurance, redundancy, succession, and recovery objectives; independent components and operators; tested restoration; and legal records that remain effective if a chain, issuer, or communications network is unavailable.

8.2 Offline resilience

As design context outside the USDC/USDT empirical scope, BIS Project Polaris treats offline CBDC functionality as a system-level problem (Bank for International Settlements, 2023a, 2023b). An owner-bound substrate can support offline custody, backup, or authorization continuity during a network disruption. It does not provide offline payment finality. Offline settlement also requires double-spend controls, value synchronization, device limits, revocation, privacy rules, fraud allocation, and later reconciliation.

The defensible institutional claim is therefore offline authorization resilience for defined keys or compartments. Deployment requires dual control, segregation of duties, geographic redundancy, tested succession, recovery-time objectives, and contingency procedures for unavailable redemption or settlement channels.

8.3 AML/CFT and payment transparency

FATF Recommendation 15 and associated guidance apply anti-money-laundering and counter-terrorist-financing obligations to covered virtual-asset service providers. Payment-transparency rules require obligated entities to obtain, retain, and transmit specified originator and beneficiary information for covered transfers (Financial Action Task Force, 2021, 2025a, 2025b). A physical cryptographic factor does not itself perform customer due diligence, sanctions screening, transaction monitoring, or Travel Rule messaging.

The architecture can support a compliant operating model when it is integrated with regulated controls: customer due diligence at onboarding; binding of a non-secret instrument identifier to a protected off-chain record; screening at funding, signing, and redemption gates; transaction and value limits; Travel Rule communication between obligated entities; retained custody and authorization logs; suspicious-activity escalation; and privacy controls for personal information. Self-hosted wallet interactions remain subject to jurisdiction-specific risk-based measures.

Offline custody can reduce the availability of real-time monitoring, so some controls must move to funding, signing policy, and redemption. Owner binding is compatible with AML/CFT objectives but is neither a safe harbor nor evidence of compliance by itself.

9. Limitations and Research Agenda

The principal empirical limitation is comparability. The sources examined supply no USDC/USDT event-level theft numerator matched to average value exposed and followed through a comparable enforcement pipeline. Commercial blockchain-intelligence datasets are useful but proprietary, updateable, and sensitive to classification rules. Cash data are jurisdiction-specific and incomplete. The study therefore does not estimate a causal effect of digitization on theft or a population RVR. The Binance update is a dated observational case. Vendor totals remain revisable, the final root cause and ultimate recovery were unresolved in the cited disclosures, and no causal control-effectiveness claim follows.

The U.S. cash calculation uses a year-end currency stock rather than a time-weighted custody exposure. It includes notes held abroad, while its numerator reflects participating U.S. agencies. It is included to demonstrate transparent calculation and caveat discipline, not as a benchmark for a digital ratio.

The owner-bound architecture is a reference specification, not an evaluated product. Its security depends on implementation, entropy, parser behavior, signing rules, operator training, material quality, recovery policy, and independence among shares and control domains. No failure probability is estimated. Usability burdens may create workarounds that reverse the intended security benefit.

Five empirical priorities follow:

First, preregister separate issuer-native USDC and USDT cohorts by network, custody configuration, and attack class; exclude fraud, insolvency, market loss, and wrapped or bridged claims.

Second, publish a versioned, deduplicated event dataset with event-time valuation, recovery horizons, custody type, and confidence labels while preserving disagreements among sources.

Third, construct daily time-weighted exposure denominators by asset, custody configuration, and attack class; estimate intensities with clustered uncertainty.

Fourth, link incidents to enforcement records and estimate reporting-to-disposition transition hazards with right-censoring and jurisdiction controls.

Fifth, implement the reference profile, publish test vectors and reproducible code, and commission independent cryptographic, hardware, red-team, usability, coercion, and accelerated-ageing studies.

Further work should evaluate threshold-signature deployments that avoid full-key reconstruction, formally verified transaction parsers, privacy-preserving compliance attestations, post-quantum migration plans based on standardized algorithms, and portfolio models that jointly represent cyber compromise, physical disaster, operational error, liquidity stress, and common-mode dependency.

10. Conclusion

USDC and USDT custody can expose value to remote, reusable, and common-mode attack paths that differ materially from the local and transport-constrained risks of most physical cash. Public evidence establishes economically significant digital-asset theft and recurring key-management compromise, but it does not identify a global per-dollar theft probability or a valid cash-versus-digital multiplier. Such a comparison requires matched event definitions, exposure denominators, periods, jurisdictions, custody perimeters, valuation rules, recovery conventions, and reporting coverage. The October 2026 Binance breach gives these measurement and authorization questions immediate practical relevance while leaving the requirements for valid comparative inference unchanged.

An owner-bound passive substrate addresses a defined part of this risk. Prohibiting a complete spend key on the record and requiring an independent owner factor avoids the bearer property of a full-key paper wallet. Passive storage removes a direct network interface from the stored share, and offline transaction verification can constrain remote interaction. These benefits depend on verified entropy, correct secret sharing, authenticated records, pre-funding verification, canonical transaction display, controlled public-data input and output, secure recovery, environmental qualification, and explicit treatment of supply-chain, side-channel, common-mode, and coercion risks.

The policy implication is bounded but consequential. Institutions that hold concentrated, rarely used digital authorities should evaluate offline, independently controlled authorization against explicit control criteria when their threat model includes remote or common-mode compromise. The precise implementation should be proportionate to value, availability needs, and regulatory obligations. The architecture does not confer absolute immunity; it converts several implicit trust assumptions into explicit, testable controls.

Data Availability Statement

This study uses publicly accessible aggregate sources listed in the references. It does not use confidential, personal, or human-subject data. Replication materials should include source snapshots where licensing permits, machine-readable metadata, exact transformations, numerical precision and rounding rules, recovery cutoffs, dataset hashes, and cryptographic test vectors. Replication materials for the original numerical analysis are publicly available at:

GitHub repository: https://github.com/printdreams/cryptoeconomic-theft-asymmetry

DOI Badge: https://doi.org/10.5281/zenodo.22057611

SHA-256 checksum of the ZIP-release file:
6D11CD49924EE3818DB3C5E595E894A3D9A5358E64104EC30DD53795C2BDAF4D

Key public sources include Federal Reserve, FBI, and Chainalysis (citations in References here below). The repository contains raw input CSVs (where permitted), processing scripts, and a Jupyter notebook to reproduce the original numerical tables and figures. The late-breaking Binance case is a narrative update from the dated sources cited in Section 4.2; it is not represented as a new dataset or computation in the archived replication release.

The notebook for Equations (1) – (21) can be accessed using an online viewer tool for Jupyter notebooks:

https://nbviewer.org/github/printdreams/cryptoeconomic-theft-asymmetry/blob/main/reproduce_report_figures.ipynb

Funding

This research was funded by the author himself and has received no external funding.

Acknowledgements

No additional acknowledgements.

Declaration of the Use of Generative AI

During preparation of this work, the author used Gemini 3.7 Flash for the purpose of finding bibliography references and data sources, and for suggesting adequate mathematical models to address the research topics. The author also used 2026 Copilot model for formatting, mathematical notation typesetting, structural organization, and stylistic editing. Finally, the author also used OpenAI ChatGPT (GPT-5) for methodological critique, source triage, mathematical consistency checking, structural editing, and language refinement. All core research topics, theoretical formulations, and empirical interpretations were performed and verified independently by the author. The author reviewed, edited, and verified the text, data, equations, and citations and accepts full responsibility for the manuscript.

Conflicts of Interest

The author is Founder, Chief Executive Officer, and majority owner of PrintDreams International AB, applicant for patent application WO 2025/237814 A1, "Method and printer unit for printing a physical note for accessing a digital asset." This commercial and intellectual-property interest is directly relevant to the architecture discussed and is disclosed for transparency. No external funder influenced the study design, source selection, analysis, manuscript preparation, or decision to submit.

References

Bank for International Settlements. (2023a). Project Polaris: Handbook for offline payments with CBDC. https://www.bis.org/publ/othp64.htm

Bank for International Settlements. (2023b). Project Polaris: High-level design guide for offline payments. https://www.bis.org/publ/othp79.htm

Bank for International Settlements. (2025a). Advancing in tandem: Results of the 2024 BIS survey on central bank digital currencies and crypto (BIS Papers No. 159). https://www.bis.org/publ/bppdf/bispap159.pdf

Bank for International Settlements. (2025b). The next-generation monetary and financial system. In Annual Economic Report 2025 (Chapter III). https://www.bis.org/publ/arpdf/ar2025e3.htm

Becker, G. S. (1968). Crime and punishment: An economic approach. Journal of Political Economy, 76(2), 169–217. https://doi.org/10.1086/259394

Binance. (2026a, October 30). *Security incident update #2 (2026-10-30).* https://www.binance.com/en/support/announcement/detail/d3306d9dca5243548c813b605e1d1751

Binance. (2026b, November 1). *Security incident | Recovery update #4: Withdrawals resume in phases (2026-11-01).* https://www.binance.com/en/support/announcement/detail/73911fb819b240f88b88256bdd318eda

Bitcoin Improvement Proposals. (2012). BIP 38: Passphrase-protected private key. https://github.com/bitcoin/bips/blob/master/bip-0038.mediawiki

Bitcoin Improvement Proposals. (2015). BIP 66: Strict DER signatures. https://github.com/bitcoin/bips/blob/master/bip-0066.mediawiki

Bonneau, J., Miller, A., Clark, J., Narayanan, A., Kroll, J. A., & Felten, E. W. (2015). SoK: Research perspectives and challenges for Bitcoin and cryptocurrencies. In 2015 IEEE Symposium on Security and Privacy (pp. 104–121). IEEE. https://doi.org/10.1109/SP.2015.14

Buterin, V. (2015). EIP-2: Homestead hard-fork changes. Ethereum Improvement Proposals. https://eips.ethereum.org/EIPS/eip-2

Buterin, V., Wilson, S., Dietrichs, A., & lightclient. (2024, May 7). *EIP-7702: Set code for EOAs.* Ethereum Improvement Proposals. https://eips.ethereum.org/EIPS/eip-7702

Chainalysis. (2024, December 19). $2.2 billion stolen from crypto platforms in 2024, but hacked volumes stagnate toward year-end as DPRK slows activity post-July. https://www.chainalysis.com/blog/crypto-hacking-stolen-funds-2025/

Chainalysis. (2026, October 30). *The $6.4 billion Binance hack: Tracing the North Korea link.* https://www.chainalysis.com/blog/binance-exploit-latest-findings-and-response/

Circle. (2025, December 12). USDC terms. https://www.circle.com/legal/usdc-terms

Conti, M., Kumar, E. S., Lal, C., & Ruj, S. (2018). A survey on security and privacy issues of Bitcoin. IEEE Communications Surveys & Tutorials, 20(4), 3416–3452. https://doi.org/10.1109/COMST.2018.2842460

Eling, M., & Jung, K. (2022). Heterogeneity in cyber loss severity and its impact on cyber risk measurement. Risk Management, 24, 273–297. https://doi.org/10.1057/s41283-022-00095-w

Federal Bureau of Investigation. (2019a). Crime in the United States 2019: Clearances. https://ucr.fbi.gov/crime-in-the-u.s/2019/crime-in-the-u.s.-2019/topic-pages/clearances

Federal Bureau of Investigation. (2019b). Crime in the United States 2019: Table 24, property stolen and recovered by type and value. https://ucr.fbi.gov/crime-in-the-u.s/2019/crime-in-the-u.s.-2019/topic-pages/tables/table-24

Federal Bureau of Investigation. (2023). Bank Crime Statistics: Federally insured financial institutions, January 1–December 31, 2023. https://www.fbi.gov/file-repository/reports-and-publications/bank-crime-statistics-2023-091724.pdf

Federal Bureau of Investigation. (2025). 2024 IC3 Annual Report. https://www.ic3.gov/AnnualReport/Reports/2024_IC3Report.pdf

Federal Bureau of Investigation. (2026, October 31). *North Korea responsible for $6 billion+ Binance attack.* https://www.fbi.gov/investigate/cyber/alerts/2026/north-korea-responsible-for-6-billion-binance-attack

Federal Reserve Board. (2026). Currency in circulation: Value. https://www.federalreserve.gov/paymentsystems/coin_currcircvalue.htm

Financial Action Task Force. (2021). Updated guidance for a risk-based approach to virtual assets and virtual asset service providers. https://www.fatf-gafi.org/en/publications/Fatfrecommendations/Guidance-rba-virtual-assets-2021.html

Financial Action Task Force. (2025a). FATF updates Standards on Recommendation 16 on payment transparency.https://www.fatf-gafi.org/en/publications/Fatfrecommendations/update-Recommendation-16-payment-transparency-june-2025.html

Financial Action Task Force. (2025b). Targeted update on implementation of the FATF Standards on virtual assets and VASPs. https://www.fatf-gafi.org/en/publications/Fatfrecommendations/targeted-update-virtual-assets-vasps-2025.html

Josefsson, S., & Liusvaara, I. (2017). Edwards-curve digital signature algorithm (EdDSA) (RFC 8032). RFC Editor. https://www.rfc-editor.org/rfc/rfc8032.html

National Institute of Standards and Technology. (2018). Recommendation for the entropy sources used for random bit generation (NIST SP 800-90B). https://doi.org/10.6028/NIST.SP.800-90B

National Institute of Standards and Technology. (2020). Recommendation for key management: Part 1—General (NIST SP 800-57 Part 1, Rev. 5). https://doi.org/10.6028/NIST.SP.800-57pt1r5

National Institute of Standards and Technology. (2023). Digital Signature Standard (DSS) (FIPS 186-5). https://doi.org/10.6028/NIST.FIPS.186-5

Perez, D., & Livshits, B. (2021). Smart contract vulnerabilities: Vulnerable does not imply exploited. In 30th USENIX Security Symposium (pp. 1325–1341). USENIX Association. https://www.usenix.org/conference/usenixsecurity21/presentation/perez

Reuter, M. (2025). Decrypting crypto: How to estimate international stablecoin flows (IMF Working Paper WP/25/141). International Monetary Fund. https://www.imf.org/-/media/Files/Publications/WP/2025/English/wpiea2025141-source-pdf.ashx

Robinson, T. (2026, November 1). *The largest theft in history: North Korea behind the $6.4 billion Binance hack.* Elliptic. https://www.elliptic.co/insights/binance-largest-hack-in-history/

Sangari, S., Dallal, E., & Whitman, M. (2022). Modeling under-reporting in cyber incidents. Risks, 10(11), Article 200. https://doi.org/10.3390/risks10110200

Shamir, A. (1979). How to share a secret. Communications of the ACM, 22(11), 612–613. https://doi.org/10.1145/359168.359176

Tether. (n.d.). Legal terms. Retrieved October 11, 2026, from https://tether.to/en/legal/

TRM Labs. (2025, March 17). Category deep-dive: $2.2 billion was stolen in crypto-related hacks in 2024. https://www.trmlabs.com/resources/blog/category-deep-dive-2-2-billion-was-stolen-in-crypto-related-hacks-in-2024

TRM Labs. (2026, October 29). *North Korea steals USD 6.36 billion from Binance in the largest hack in history.* https://www.trmlabs.com/resources/blog/north-korea-steals-usd-6-36-billion-from-binance-in-the-largest-hack-in-history

Appendix A. Minimum Empirical Protocol

A.1 Inclusion and valuation

An event enters the theft numerator only when an unauthorized transfer or legally reported taking occurred. Each value must use a documented event-time price and currency-conversion source. Recoveries require a stated horizon. Market loss, fraud, ransomware, insolvency, and unauthorized technical transfer should remain distinct categories unless category-specific results are also shown.

A.2 Exposure

Exposure must correspond to the attack class. For custodian-key compromise, use average assets controlled by the included key-management perimeter. For smart-contract exploitation, use average value locked in the included contracts. For self-custody drainers, use value in reachable wallets or an explicitly modeled proxy. Stablecoin supply is not automatically the exposure denominator for every digital-asset incident. For USDC/USDT, record the issuer-native contract and network separately; the proposed key-compromise cohort excludes bridged claims and DeFi exposure.

A.3 Enforcement

Store event-level and person-level records separately. Record civil, criminal, administrative, sanctions, and recovery outcomes independently. Use event and disposition dates to address censoring; unresolved recent cases must not be treated as enforcement failures.