AMERICAN JOURNAL OF AI CYBER COMPUTING MANAGEMENT
PEER REVIEWED, REFERRED & INDEXED JOURNAL
E-ISSN:3069-0102
VOL.5, NO. 1(2025)

On-Chain Reputation and Anomaly Scoring:

A Cross-EVM Risk Engine for Sybil Detection, Wash-Trading Identification, and Marketplace Trust Operationalisation

Ibtihajul Islam
Independent Researcher
ibtihajul@gmail.com

Dr. Kashif Saleem
Associate Professor

Abstract

Decentralised marketing platforms that operate across multiple EVM-compatible blockchains face a distinctive risk surface: the same adversarial patterns observed in decentralised finance (sybil networks, wash-trading rings, bridge-based provenance laundering) migrate directly into influencer campaign ecosystems, where inflated wallet activity translates into inflated audience claims and fraudulent escrow eligibility. Existing blockchain analytics tools provide raw attribution data but stop short of operationalising their signals into real-time marketplace controls. This paper presents the Cross-EVM Reputation and Anomaly Scoring Engine (CRASE), a distributed analytics architecture that normalises on-chain telemetry across Ethereum, Polygon, Arbitrum, Base, and Solana, computes composite wallet reputation scores, and feeds actionable risk verdicts into both discovery ranking and escrow gating logic. CRASE employs a four-component scoring model: transaction graph topology analysis for sybil cluster detection, temporal pattern analysis for wash-trading identification, cross-chain provenance tracing for bridge-based obfuscation detection, and wallet-age and activity normalisation for cold-start equity. Evaluation against a labelled dataset of 120,000 wallets drawn from publicly documented DeFi exploit post-mortems and influencer fraud reports yields a sybil detection F1 of 0.88 and a wash-trading detection F1 of 0.83, with score computation latency below 850 ms at the 99th percentile. Critically, CRASE outputs are consumed directly by marketplace discovery and pre-escrow settlement controls, closing the gap between blockchain signal collection and product-level risk enforcement that characterises less mature platform architectures.

Keywords: on-chain reputation, sybil detection, wash trading, EVM normalisation, blockchain analytics, anomaly scoring, decentralised marketing, cross-chain provenance, wallet risk scoring, influencer fraud

1. Introduction

The growth of Web3-native marketing platforms introduces a category of fraud risk with no direct precedent in traditional digital advertising: the adversary does not merely fabricate engagement metrics in a centralised database but constructs entire on-chain identities whose transaction histories are designed to

mimic legitimate wallet behaviour. A KOL who controls a ring of wallets that trade amongst themselves can present the appearance of an active, economically engaged audience; a campaign creator who recycles previously flagged addresses through bridging protocols can re-enter a marketplace with a seemingly clean wallet history. These attacks exploit a gap that is pervasive in first-generation Web3 platforms: blockchain data is collected but not operationalised into automated controls.

The consequences of this gap are material. Fraudulent campaign approvals result in direct escrow losses; low-quality KOL audiences dilute return on investment for brands, increasing churn; and the absence of explainable on-chain risk signals prevents enterprise buyers from satisfying their internal due-diligence requirements. The problem is further complicated by multi-chain fragmentation: adversarial actors deliberately distribute activity across Ethereum mainnet, Layer-2 rollups (Arbitrum, Optimism, Base), and EVM-compatible sidechains (Polygon), and bridge between them to sever the provenance trail visible on any single chain.

This paper presents the Cross-EVM Reputation and Anomaly Scoring Engine (CRASE), an architecture that addresses these challenges by: (i) normalising on-chain telemetry across heterogeneous EVM chains into a unified entity model; (ii) computing multi-dimensional wallet reputation scores using graph topology, temporal pattern, cross-chain provenance, and activity normalisation components; and (iii) operationalising the resulting verdicts directly into marketplace discovery ranking and escrow gating, closing the loop between signal collection and product-level risk enforcement. The contributions of this work are:

- A formal normalisation schema for multi-chain EVM telemetry that resolves cross-chain entity relationships using deterministic address derivation heuristics and bridge event indexing.
- A four-component scoring model (TGTS: Transaction Graph, Temporal, Sybil, Settlement) with a composite scoring function and calibrated weight vector.
- An adversarial pattern taxonomy covering six sybil and wash-trading variants observed in DeFi exploit post-mortems and their CRASE detection mappings.
- An empirical evaluation on a 120,000-wallet labelled dataset demonstrating F1 scores of 0.88 (sybil) and 0.83 (wash-trading) with sub-850 ms P99 scoring latency.
- A product integration architecture demonstrating how CRASE verdicts feed directly into discovery ranking and pre-escrow settlement gates.

2. Background and Related Work

2.1 Sybil Attacks in Decentralised Networks

Douceur [1] introduced the Sybil attack as a general threat to distributed systems in which a single adversarial entity creates a large number of pseudonymous identities to subvert reputation or voting mechanisms. In blockchain ecosystems, the cost of creating new wallet addresses is negligible, making Sybil resistance a fundamental design challenge for any reputation system that relies on wallet-based identity. Leveraging on-chain economic activity as a Sybil-resistance signal was proposed by Buterin et al. [2] in the context of decentralised governance; however, the wash-trading attack (in which the adversary generates the economic activity themselves across controlled wallets) undermines naive activity-based scoring.

2.2 Wash-Trading Detection in Crypto Markets

Wash trading — the practice of simultaneously buying and selling the same asset across controlled accounts to generate the appearance of volume — was first documented at scale in NFT markets by Serneels [3] and subsequently characterised by von Wachter et al. [4], who identified cyclical transaction graph patterns (closed loops of value transfer between a small set of addresses) as the primary structural discriminant. Cong et al. [5] extended this analysis to centralised exchange token listing manipulation, demonstrating that temporal burst patterns (high-frequency, low-value transfers concentrated in narrow time windows) are reliable secondary signals. CRASE combines both structural and temporal signals in its scoring model.

2.3 Cross-Chain Attribution and Bridge Analytics

Cross-chain bridges transfer assets between blockchains through lock-and-mint or burn-and-release mechanisms, creating a deliberate break in the on-chain provenance trail. Xia et al. [6] demonstrated that bridge-mediated provenance laundering is now a standard technique in DeFi exploit execution, used to move stolen funds across chains before cashing out. From a marketing fraud perspective, the same technique is used to launder wallet history: a flagged address bridges its funds to a destination chain, receives them at a fresh address, and presents the fresh address for platform registration. CRASE addresses this by maintaining a Bridge Event Index that tracks lock/mint pairs across seven major bridge protocols and propagates risk scores across detected bridge relationships.

2.4 Blockchain Risk Scoring Systems

Commercial blockchain analytics platforms including Chainalysis [7], Elliptic [8], and TRM Labs [9] provide wallet risk scores primarily calibrated for Anti-Money Laundering (AML) compliance use cases: proximity to sanctioned addresses, darknet markets, and ransomware wallets. These signals are valuable but insufficient for marketplace trust decisions, where the adversarial patterns (sybil clusters, wash-trading rings, bridge laundering) are distinct from AML typologies and where the scoring output must be consumed in real time by automated product controls rather than by compliance analysts. CRASE is designed specifically for the marketplace trust context and integrates AML signals as one input component alongside marketplace-specific pattern detectors.

3. CRASE Architecture

3.1 System Overview

CRASE is a distributed scoring pipeline composed of five layers: a Multi-Chain Indexer, a Normalisation Layer, a Scoring Engine, a Verdict Store, and a Product Integration Layer. Data flows unidirectionally from raw chain events through normalisation to composite scores, which are then consumed by marketplace and escrow services via a low-latency API. Figure 1 (Table 1) summarises the layers and their responsibilities.

Layer	Components	Output
Multi-Chain Indexer	The Graph subgraphs for Ethereum, Polygon, Arbitrum, Base; custom Solana indexer (Geyser plugin); bridge event listeners for 7 protocols	Normalised event streams per chain

Normalisation Layer	Address resolver (deterministic derivation + bridge mapping), entity deduplication, cross-chain wallet graph builder	Unified Wallet Entity model with cross-chain edge set
Scoring Engine	Transaction Graph Analyser (TGA), Temporal Pattern Detector (TPD), Sybil Cluster Identifier (SCI), Settlement History Scorer (SHS)	Four component scores + composite CRASE Score (0.0 to 1.0)
Verdict Store	Redis sorted set cache (TTL: 4 hours, hot wallets), DynamoDB persistent store (full history), invalidation webhook publisher	Low-latency score lookups; score history audit trail
Product Integration Layer	Discovery Ranking API adapter, Escrow Gate API adapter, Brand Dashboard score explainer	Risk-adjusted discovery rankings; escrow approval/block verdicts; explainable score breakdowns

Table 1. CRASE Architecture Layers

3.2 Multi-Chain Normalisation

The normalisation layer resolves the fundamental heterogeneity of multi-chain data into a unified Wallet Entity model. Three mechanisms address cross-chain identity fragmentation. First, deterministic address derivation: EVM-compatible chains share the same address derivation algorithm (secp256k1 public key hashing), meaning the same private key generates the same address on Ethereum, Polygon, Arbitrum, and Base. CRASE groups identically addressed wallets across EVM chains into a single entity by default, while maintaining chain-specific activity records as subgraphs. Second, bridge event mapping: the Bridge Event Index tracks lock/mint and burn/release event pairs across seven major protocols (Polygon Bridge, Arbitrum Bridge, Optimism Gateway, Stargate, Hop Protocol, Across Protocol, and LayerZero). Wallets connected by confirmed bridge events receive a BRIDGE_LINKED edge with a confidence weight derived from the bridge's own security posture. Third, behavioural clustering: wallets that cannot be linked by address derivation or bridge events but share statistically improbable behavioural characteristics (identical transaction timing, matched value transfers, common counterparty sets) are flagged as probable sybil siblings with a probabilistic confidence score rather than a deterministic link.

3.3 Scoring Model

The CRASE composite score S(w) for wallet w is a weighted linear combination of four component scores, each normalised to the interval [0.0, 1.0] where 0.0 represents the lowest risk (best reputation) and 1.0 represents the highest risk:

$$S(w) = \alpha \cdot TGA(w) + \beta \cdot TPD(w) + \gamma \cdot SCI(w) + \delta \cdot SHS(w)$$

where alpha, beta, gamma, and delta are the component weights (alpha + beta + gamma + delta = 1.0). The weight vector was calibrated empirically against the labelled dataset described in Section 5, using isotonic regression to ensure score monotonicity with respect to ground-truth fraud labels. The calibrated weights are shown in Table 2 alongside component descriptions.

Component	Symbol	Weight	Signal Source	Primary Adversarial Pattern Detected

Transaction Graph Analyser	TGA(w)	0.35	On-chain transaction graph topology; Louvain clustering coefficient; cycle detection	Sybil clusters, wash-trading rings, coordinated airdrop farming
Temporal Pattern Detector	TPD(w)	0.25	Transaction timestamp distributions; inter-arrival time entropy; burst detection	Wash-trading bursts, bot-driven engagement mimicry, coordinated action timing
Sybil Cluster Identifier	SCI(w)	0.28	Bridge event index; address co-occurrence graph; probabilistic behavioural clustering	Bridge-mediated provenance laundering, address recycling, cross-chain sybil networks
Settlement History Scorer	SHS(w)	0.12	Prior platform escrow participation; dispute history; AML proximity score (Chainalysis/TRM integration)	Recidivist fraud, prior campaign disputes, sanctioned counterparty proximity

Table 2. CRASE Component Scores: Weights, Signal Sources, and Target Patterns

3.3.1 Transaction Graph Analyser (TGA)

TGA constructs a directed weighted transaction graph G_w for wallet w, comprising all counterparty addresses transacted with within a configurable lookback window (default: 180 days) and all second-degree counterparties (two hops). The TGA score is derived from three graph metrics: (i) the normalised clustering coefficient of w within G_w, measuring the density of transactions among w's counterparties independent of w (high density indicates a closed trading ring); (ii) the count of value-neutral cycles (sequences of transfers that return the same or similar value to the originating address within 72 hours); and (iii) the Gini coefficient of counterparty diversity (extremely low diversity — most value concentrated among very few counterparties — is characteristic of wash-trading). These three sub-signals are combined using a logistic function to produce TGA(w) in [0, 1].

3.3.2 Temporal Pattern Detector (TPD)

TPD analyses the timestamp series of all transactions originating from or received by wallet w. The core signal is the Shannon entropy of the inter-arrival time distribution: legitimate wallets exhibit high entropy (irregular, human-driven timing), while bot-driven sybil or wash-trading wallets exhibit low entropy (regular, scripted timing) or multi-modal distributions characteristic of scripted burst-and-pause cycles. TPD supplements entropy with a burst detection algorithm (sliding window, 1-hour epoch) that flags wallets whose 95th-percentile burst intensity exceeds three standard deviations from the distribution of similar-age wallets. Both signals are integrated via a two-feature logistic classifier trained on the labelled dataset.

3.3.3 Sybil Cluster Identifier (SCI)

SCI operates on the normalised cross-chain wallet graph produced by the normalisation layer. For each wallet w, SCI computes the maximum risk score of any wallet in w's k=2 hop neighbourhood, weighted by the edge confidence of the path (bridge link confidence, address derivation certainty, or behavioural clustering probability). This neighbourhood risk propagation ensures that a freshly bridged address inherits meaningful risk signal from its provenance wallet even if its own transaction history is thin. SCI additionally applies a community detection pass (Louvain algorithm, resolution parameter gamma = 1.2) to identify sybil communities and assigns each detected community a community risk score derived from the proportion of previously flagged members.

3.4 Product Integration: Discovery and Escrow Gates

The operationalisation of CRASE scores into product controls is implemented at two integration points. In discovery, CRASE scores are incorporated into the KOL and campaign ranking function as a multiplicative penalty: a KOL wallet with CRASE score above 0.6 (HIGH risk) has its discovery ranking suppressed by a configurable factor (default: 0.4x), reducing but not eliminating its visibility pending human review. This approach avoids hard exclusions that could produce false-positive brand friction while ensuring that high-risk profiles are systematically deprioritised.

In escrow gating, CRASE operates as a blocking control: campaign escrow initialisation is rejected for any participating wallet (campaign creator or KOL primary wallet) with CRASE score above 0.75 (CRITICAL risk), and requires manual compliance review for scores between 0.6 and 0.75. The CRASE API returns the composite score, the four component scores, the primary evidence signals (e.g., "3 wash-trading cycles detected within 30 days"), and a recommended action, providing the compliance reviewer with an auditable, explainable basis for their decision.

4. Adversarial Pattern Taxonomy

CRASE is designed against a taxonomy of six adversarial patterns drawn from documented DeFi exploit post-mortems and influencer marketing fraud case studies. Table 3 maps each pattern to its CRASE detection mechanism and the primary component score implicated.

Pattern	Mechanism	Primary CRASE Component	Detection Threshold
Wash-Trading Ring	Adversary controls 3 to 20 wallets that trade identical NFTs or tokens among themselves in cycles to simulate volume and engagement.	TGA (cycle detection)	2+ value-neutral cycles within 30 days; TGA > 0.65
Temporal Burst Bot	Scripted bot executes high-frequency micro-transactions at regular intervals to inflate transaction count while mimicking activity.	TPD (entropy, burst)	Inter-arrival entropy < 1.8 bits; burst intensity > 3 SD; TPD > 0.60
Bridge Provenance Launderer	Flagged wallet bridges assets to a fresh address on a destination chain, presenting the fresh address for platform registration.	SCI (bridge index)	Bridge link detected; source wallet SCI > 0.55; inherited SCI propagated
Airdrop Farmer Cluster	Adversary creates hundreds of wallets to qualify for token airdrops, generating artificial community appearance.	TGA (clustering coeff.) + SCI (community)	Louvain community with >60% airdrop-only activity; composite S(w) > 0.55
Sybil Audience Inflator	KOL controls sybil wallets that interact with their content	SCI (address co-occurrence)	co-occurrence graph density > 0.4 within 2-hop

	contracts to simulate an engaged on-chain audience.		neighbourhood; SCI > 0.70
Recidivist Address Recycler	Previously banned wallet re-registers using a slightly modified address or after a dormancy period intended to age-wash the account.	SHS (settlement history) + TPD (dormancy break)	Prior platform dispute record; dormancy break after flag; SHS > 0.75

Table 3. Adversarial Pattern Taxonomy and CRASE Detection Mappings

5. Empirical Evaluation

5.1 Dataset Construction

The evaluation dataset comprises 120,000 wallet records drawn from three sources: (i) 38,000 wallets publicly identified in DeFi exploit post-mortems and Chainalysis annual crime reports [7] as participating in sybil attacks, wash trading, or bridge-based laundering, labelled as fraud-positive; (ii) 12,000 wallets identified in influencer marketing fraud investigations published by HypeAuditor [10] and Modash [11], labelled as fraud-positive; and (iii) 70,000 wallets randomly sampled from Ethereum and Polygon mainnet transaction history, manually reviewed and labelled as fraud-negative after exclusion of any wallet within 3 hops of a known fraud cluster. The dataset contains 41.7% fraud-positive labels, deliberately oversampled relative to the estimated real-world base rate of 11-18% [7] to ensure sufficient minority-class signal for model calibration. Evaluation metrics are reported on a held-out test set of 24,000 records (20% stratified split) after training the logistic sub-classifiers within each CRASE component on the remaining 96,000 records.

5.2 Detection Performance

Detector	Precision	Recall	F1 Score	AUC-ROC
Sybil Cluster Detection (SCI-dominant)	0.91	0.85	0.88	0.94
Wash-Trading Detection (TGA + TPD)	0.86	0.80	0.83	0.91
Bridge Provenance Detection (SCI)	0.89	0.78	0.83	0.92
Composite CRASE Score (all patterns)	0.87	0.84	0.85	0.93
Baseline: AML Score Only (Chainalysis)	0.74	0.61	0.67	0.79
Baseline: Rule-Based Filters	0.63	0.72	0.67	0.71

Table 4. CRASE Detection Performance on Held-Out Test Set (n = 24,000)

The composite CRASE score achieves an F1 of 0.85 across all adversarial patterns, representing an 18 percentage point improvement over the best baseline (AML-only score, F1 = 0.67). The improvement is

most pronounced in recall: CRASE identifies 84% of fraud-positive wallets versus 61% for the AML baseline, reflecting CRASE's coverage of marketplace-specific attack patterns that AML typologies do not address. The AUC-ROC of 0.93 indicates strong separability across the full score range, supporting the use of variable threshold policies (discovery suppression at 0.6, escrow block at 0.75) without requiring binary classification at a single cut-point.

5.3 Scoring Latency

Scenario	Median Latency (ms)	P95 Latency (ms)	P99 Latency (ms)
Single-chain wallet (Ethereum, warm cache)	42	98	187
Multi-chain wallet (3 chains, warm cache)	118	310	520
Bridge-linked wallet (SCI propagation required)	220	490	830
Cold-start wallet (no cache, full graph traversal)	390	680	847
Batch scoring (100 wallets, parallel)	1,240 total / 12.4 per wallet	2,100 total	2,890 total

Table 5. CRASE Scoring Latency Under Production-Representative Load (200 req/s)

All scoring scenarios satisfy the sub-850 ms P99 latency target required for synchronous escrow gating decisions. Cold-start scoring (847 ms P99) approaches but does not exceed this threshold; in practice, the vast majority of wallets in active campaigns will have been pre-scored and cached, making cold-start latency a rare code path. Batch scoring at 12.4 ms per wallet supports nightly full-portfolio re-scoring of all registered KOL wallets without impacting API availability.

6. Discussion

6.1 Operationalisation as the Critical Differentiator

The technical literature on blockchain fraud detection is extensive, yet first-generation Web3 marketing platforms routinely fail to translate detection capability into automated product controls. The architectural gap is not detection accuracy but operationalisation: connecting the output of a risk scoring system to the platform decisions that determine which campaigns proceed, which KOLs appear in discovery, and which escrow requests are approved. CRASE is explicitly designed around this integration requirement. The scoring API is co-designed with the discovery ranking function and the escrow gateway, rather than being retrofitted as an advisory signal. This distinction is what converts blockchain analytics from a compliance checkbox into a product security capability.

6.2 Noise Handling and Score Stability

On-chain data is inherently noisy: failed transactions, contract interactions from DeFi protocol automation, and legitimate high-frequency trading can produce signal patterns superficially similar to adversarial behaviour. CRASE addresses this through two mechanisms. First, counterparty filtering excludes known protocol router addresses (Uniswap, Curve, Aave, and 200 additional whitelisted protocol contracts) from transaction graph construction, preventing protocol interaction patterns from contaminating the graph topology signal. Second, score smoothing applies an exponentially weighted moving average over a 7-day window for the SHS and TGA components, dampening transient signal spikes while preserving responsiveness to sustained pattern changes. The resulting score stability (standard deviation of daily score changes < 0.04 for legitimate wallets in a 30-day backtest) reduces false-positive discovery suppression events and builds platform trust in the system's outputs among compliance reviewers.

6.3 Adversarial Adaptation and Score Gaming

A sophisticated adversary aware of the CRASE scoring model might attempt to construct transaction histories that score below detection thresholds: spacing wash trades beyond the burst detection window, routing bridge transfers through a chain not yet indexed, or diluting sybil cluster density by adding benign-looking counterparties. CRASE is not designed to be static: the component weight vector and detection thresholds are subject to quarterly recalibration against newly labelled fraud cases, and the bridge event index expands as new bridging protocols achieve material adoption. Future work should evaluate the application of adversarial training techniques [12] to the TGA and TPD sub-classifiers to improve robustness under deliberate evasion.

7. Conclusion

This paper has presented CRASE, a cross-EVM reputation and anomaly scoring engine that normalises on-chain telemetry across multiple blockchain networks, computes multi-dimensional wallet risk scores using transaction graph topology, temporal pattern, sybil cluster, and settlement history components, and operationalises the resulting verdicts directly into marketplace discovery ranking and escrow gating controls. Evaluation on a 120,000-wallet labelled dataset demonstrates composite F1 of 0.85 and an 18 percentage point improvement over AML-baseline scoring, with P99 scoring latency of 847 ms meeting the requirements of synchronous escrow decision pipelines.

The central contribution of CRASE is not any individual detection algorithm but the end-to-end architecture that connects blockchain signal collection to product-level risk enforcement — a capability that remains absent from most Web3 marketplace platforms despite the maturity of underlying blockchain analytics research. The architecture described is generalisable to any decentralised marketplace in which wallet reputation must inform automated access and settlement decisions: NFT marketplaces, DeFi protocol governance, and peer-to-peer service platforms all share the structural requirement that CRASE is designed to satisfy.

References

[1] Douceur, J. R. (2002). The Sybil Attack. Proceedings of the 1st International Workshop on Peer-to-Peer Systems (IPTPS).

[2] Buterin, V., Hitzig, Z., & Weyl, E. G. (2019). A Flexible Design for Funding Public Goods. Management Science, 65(11).

[3] Serneels, S. (2022). Identifying Wash Trading in the NFT Market. Working Paper. University of Antwerp.

[4] von Wachter, V., Jensen, J. R., Regner, F., & Ross, O. (2022). NFT Wash Trading: Quantifying Suspicious Behaviour in NFT Markets. Proceedings of the 4th ACM International Conference on AI in Finance.

[5] Cong, L. W., Li, X., Tang, K., & Yang, Y. (2023). Crypto Wash Trading. Management Science, 69(11).

[6] Xia, P., Wang, H., Luo, X., Wu, L., Zhou, Y., Hua, G., ... & Xu, G. (2021). Trade or Trick? Detecting and Characterizing Scam Tokens on Uniswap Decentralised Exchange. Proceedings of the ACM on Measurement and Analysis of Computing Systems, 5(3).

[7] Chainalysis. (2024). The Chainalysis 2024 Crypto Crime Report. Chainalysis Inc.

[8] Elliptic. (2023). Financial Crime Typologies in DeFi: Elliptic Research Report.

[9] TRM Labs. (2024). TRM Blockchain Intelligence: Risk Scoring Methodology. TRM Labs Inc.

[10] HypeAuditor. (2024). State of Influencer Marketing Fraud: Annual Industry Report.

[11] Modash. (2024). Influencer Marketing Fraud Report: Fake Followers, Engagement Pods, and On-Chain Signals.

[12] Zugner, D., Akbarnejad, A., & Gunnemann, S. (2018). Adversarial Attacks on Neural Networks for Graph Data. Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.

[13] Blondel, V. D., Guillaume, J.-L., Lambiotte, R., & Lefebvre, E. (2008). Fast Unfolding of Communities in Large Networks. Journal of Statistical Mechanics, 2008(10).

[14] The Graph Protocol. (2023). The Graph: A Decentralized Query Protocol. Technical Whitepaper v2.

[15] Narayanan, A., Bonneau, J., Felten, E., Miller, A., & Goldfeder, S. (2016). Bitcoin and Cryptocurrency Technologies. Princeton University Press.

[16] Ron, D., & Shamir, A. (2013). Quantitative Analysis of the Full Bitcoin Transaction Graph. Proceedings of Financial Cryptography and Data Security.

[17] Meiklejohn, S., Pomarole, M., Jordan, G., Levchenko, K., McCoy, D., Voelker, G. M., & Savage, S. (2013). A Fistful of Bitcoins: Characterizing Payments Among Men with No Names. Proceedings of IMC.

[18] Association of National Advertisers. (2024). ANA Programmatic Media Supply Chain Transparency Study.

[19] Hu, X., Chen, T., & Zhang, Y. (2023). Cross-Chain Bridge Security: A Systematic Review of Attack Vectors and Mitigations. IEEE Transactions on Dependable and Secure Computing.

[20] Bartoletti, M., Carta, S., Cimoli, T., & Saia, R. (2020). Dissecting Ponzi Schemes on Ethereum: Identification, Analysis, and Impact. Future Generation Computer Systems, 102.