Blockchain and Cryptocurrency

Proceedings of the 5th Blockchain and Cryptocurrency Conference (B2C' 2026)

28-30 October 2026
Las Palmas de Gran Canaria, Spain



Blockchain and Cryptocurrency

Proceedings of the 5th Blockchain and Cryptocurrency Conference (B2C' 2026)

28-30 October 2026

Las Palmas de Gran Canaria, Spain

Edited by Sergey Y. Yurish

IFSA Publishing

Sergey Y. Yurish, *Editor*
Blockchain and Cryptocurrency. B2C' 2026 Conference Proceedings

Copyright © 2026
by International Frequency Sensor Association (IFSA) Publishing, S. L.

E-mail (for orders and customer service enquires): ifsa.books@sensorsportal.com

Visit our Home Page on http://www.sensorsportal.com

All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (IFSA Publishing, S. L., Barcelona, Spain).

Neither the authors nor International Frequency Sensor Association Publishing accept any responsibility or liability for loss or damage occasioned to any person or property through using the material, instructions, methods or ideas contained herein, or acting or refraining from acting as a result of such use.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identifying as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

ISBN: 978-84-09-93453-9
BN-20261020-XX
BIC: UNK; Thema: UNKD

Contents

Foreword **4**

An Exploratory Comparative Analysis of the Potential of Real Estate Tokenization to Enhance Liquidity, Mitigate Risk, and Support Smart Investment: International Experiences and a Proposed Framework for Kish Free Zone **5**
M. Noroozinejad Farsangi, A.Tariyan, J. Mahmoudi and S. M. Hashemi

Hidden Edge-State Inference in Bitcoin Lightning Network Topologies **15**
D. Patti, S. Monteleone, E. Russo and M. Palesi

Centralized Ledgers Versus Decentralized Permissionless Blockchains: A Game-Theoretic Comparison **19**
Kjell Hausken and Guizhou Wang

Return and Volatility Spillovers across Blockchain Layers **26**
C. Jeanneaux

A Framework based on ETL and Blockchain for Advanced Cyber Threat Intelligence Sharing **31**
A. Fotia, A. Salemi, A. Canino and G. Lax

Cybersecurity Awareness for Web3 Addressing Human-Centric Attack Surfaces in Blockchain Ecosystems **37**
N. P. R. Mendes and J. R. Almeida, Jr.

Privacy-Preserving Blockchain E-Voting System Based on Zero-Knowledge Proofs and Separated Identity Verification **42**
H. M. Sung, T. H. Wu, K. C. Tsao, T. C. Chiu, Q. Y. Chen, J. Y. Lin and Y. P. Hung

Embedded Cryptocurrency Wallets: A Comparative User Study on Security and Usability **48**
Dalibor Detko, Juraj Gergel, Martin Perešni and Ivan Homoliak,

Wallet Security Is Compositional: A Failure-Domain-Aware Evaluation of Cryptocurrency Wallet Configurations **50**
Tomáš Fišer, Martin Perešíni and Ivan Homoliak

Transactional and Deposit-Like Digital Money: Comparing U.S. Payment Stablecoins and China's e-CNY **53**
Saša Žiković,, Maria Čuljak, Bojan Tomić and Hoda Mansour

Algebraic Analysis of Groth16 and Multi-Framework Empirical Benchmarking of zk-SNARK Constraint Pipelines **57**
Michal Vaňo, Martin Perešni and Ivan Homoliak

Pre-Mempool Blockchain Transaction Compliance: Regulatory Convergence toward a Pre-Broadcast Screening Mandate **60**
Praveen Giri

Foreword

It is my great pleasure to welcome all authors, participants, reviewers, invited speakers and readers to the Proceedings of the 5th Blockchain and Cryptocurrency Conference (B2C' 2026), held on 28–30 October 2026 in Las Palmas de Gran Canaria, Spain.

Blockchain technology and digital currencies continue to stimulate important scientific, economic and societal discussions. What began largely as a new approach to transferring value has developed into a broad field involving distributed systems, cryptography, financial infrastructure, asset tokenization, cybersecurity and digital governance. Alongside these opportunities, essential questions remain concerning security, scalability, privacy, usability, regulatory compliance and the circumstances in which decentralization provides a genuine advantage over conventional solutions.

B2C' 2026 provides an international forum for researchers, engineers, economists and practitioners to exchange ideas, present recent achievements and critically examine emerging applications. The contributions included in this volume reflect both the diversity of the field and the need to connect theoretical advances with evidence from practical systems. They address subjects ranging from blockchain network analysis and digital finance to secure infrastructure, privacy-preserving protocols and human-centred technology design.

Several papers investigate the foundations and behaviour of blockchain-based networks and markets. Research on hidden edge-state inference in the Bitcoin Lightning Network and on return and volatility spillovers across blockchain layers deepens our understanding of network operation and interconnected financial risks. A game-theoretic comparison of centralized ledgers and permissionless blockchains raises a fundamental question: which architecture offers the most appropriate balance of incentives, access, trust and efficiency for a given application?

Other contributions explore how blockchain may support new services and forms of economic participation. The study of real estate tokenization considers potential improvements in liquidity and investment access, while comparative analysis of payment stablecoins and China's e-CNY places digital money within different institutional and regulatory models. Work on pre-mempool transaction compliance further illustrates the increasing importance of reconciling technological innovation with financial oversight and legal requirements.

Security and privacy are equally prominent themes in these Proceedings. The papers examine cyber threat intelligence sharing, human-centred attack surfaces in Web3, the usability and security of embedded cryptocurrency wallets, and the way wallet protection depends on the interaction of multiple failure domains. Research on zero-knowledge proofs addresses both privacy-preserving electronic voting and the performance of cryptographic constraint pipelines. Together, these studies remind us that trustworthy blockchain applications require more than sound cryptographic primitives: they also depend on careful system design, informed users and robust operational practices.

On behalf of the Organizing Committee, I sincerely thank all authors for sharing their work and contributing to the scientific programme. I am equally grateful to the reviewers, International Program Committee members, session chairs, invited speakers and everyone whose efforts made B2C' 2026 and these Proceedings possible.

I hope that this volume will provide a valuable reference for researchers and practitioners, encourage productive discussions and inspire new collaborations across disciplines. May the ideas presented at B2C' 2026 contribute to a more rigorous understanding and responsible development of blockchain technologies and digital finance.

Prof., Dr. Sergey Y. Yurish
B2C' 2026 Chairman

(002)

An Exploratory Comparative Analysis of the Potential of Real Estate Tokenization to Enhance Liquidity, Mitigate Risk, and Support Smart Investment: International Experiences and a Proposed Framework for Kish Free Zone

M. Noroozinejad Farsangi [1], A.Tariyan [2], J. Mahmoudi [3] and S. M. Hashemi [4]
[1] Director of Planning and Strategic Supervision, Kish Investment and Development Company, Kish, Iran
[2] Head of Sales and Commercial Affairs Department, Kish Investment and Development Company, Kish, Iran
[3] Deputy for Management Development, Kish Free Zone Organization, Kish, Iran
[4] Chief Executive Officer, Kish Investment and Development Company, Kish, Iran
Tel.: +989369801713
E-mail: mohsen.prince@gmail.com

Summary: This study examines real estate tokenization as a potential response to structural liquidity constraints and capital concentration in Kish Free Zone. Using a mixed approach, it compares international experiences and analyzes 203 real estate transactions (2021–2025), totaling 33588 billion Rials. Annual transaction value rose approximately 365 % to 14244 billion Rials in 2025, while average contract value increased from 76 to 273 billion Rials (about 259 %, or 3.6-fold), indicating higher entry barriers for retail investors. The findings motivate a conceptual framework for real estate tokenization within the KishConnect ecosystem, based on fractional ownership and aligned with Iranian regulations; KishConnect itself is not presented as the study's proposal. The framework could improve liquidity and broaden investment access, subject to legal, regulatory, and empirical assessment.

Keywords: Real-world asset tokenization, Fractional real estate ownership, Blockchain technology, KishConnect platform.

1. Introduction

Real estate is a unique and complex asset class and is recognized as one of the main drivers of economic growth and a vital indicator of a region's health and welfare [1]. The real estate market, in terms of asset base and transactional activity, is considered a significant segment of the global economy. Unlike other asset classes, this sector is frequently criticized for substantial transaction costs, land-use regulations, structural liquidity constraints, high investment prices, expensive intermediaries, and market conditions that can retain assets for extended periods [2-4].

Blockchain, as a versatile technology, has numerous potential applications across various industries, one of which is real estate. Due to its high level of security and transparency, it is particularly significant in real estate transactions [5]. The convergence of digital technologies with asset markets has given rise to a new paradigm known as real-world asset tokenization, in which physical ownership rights are transformed into blockchain-based digital tokens [6, 7]. Despite its large scale, the traditional real estate market faces challenges such as high transaction costs (5 to 7 percent), legal complexity, and long ownership transfer periods (30 to 90 days) [8, 9]. The indivisible nature of real estate results in low liquidity and capital concentration in large transactions, effectively excluding retail investors [3, 10]. Given the weaknesses present in the real estate ecosystem and their impact on overall market efficiency, this market becomes a primary candidate for this transformation. According to forecasts by Deloitte, the value of tokenized real estate is expected to reach 4 trillion dollars by 2035 [11, 12]. Tokenization, by dividing assets into smaller units, enables participation with minimal capital and, through smart contracts, reduces operational costs by up to 65 percent [13-15]. Studies indicate that this technology has the potential to enhance transparency and expand access to asset markets [16, 17].

In 2026, tokenization has moved beyond the experimental stage and has become a permanent infrastructure of capital markets. Leading experiences in Dubai through the VARA organization and the introduction of a new legal category of "asset-backed virtual assets" have provided the necessary framework to distinguish real estate tokens from securities. Similarly, Singapore, through the Project Guardian initiative, has enabled real-time settlement of transactions in a regulated environment [18, 19]. Free zones, as "living laboratories of innovation," provide an optimal platform for testing these models at scale [20, 21]. Kish Island, as one of Iran's economic hubs, despite the growth in transaction value, faces the challenge of "capital lock-in" [22]. As presented in Table 2, Statistical data analysis in this study shows that 3745 hectares (41.6 %) of the island's land remains either unallocated or vacant, while the 73 % concentration of transaction value in the past two years (2024–2025) indicates that the market has entered a phase of rapid price growth and asset inflation [23, 24]. Therefore, establishing a smart platform such as KishConnect to create a digital secondary market and enhance the efficiency of the island's real estate market is considered a strategic necessity [25, 26].

According to Table 1, the average monetary value of each contract increased from 76 billion Rials in 2021 to 273 billion Rials in 2025 (an approximately ~259 % increase (3.6-fold) in average contract value, increasing barriers to entry).

Table 1. Report of Real Estate Sales in 2021–2025 at Kish Investment and Development Company.

Row	Y.	No.	Area (m²)	Amount (Bln. Rls.)	Avg. Contract Value (Bln. Rls.)
1	2021	40	4383	3060	76
2	2022	28	2941	3170	113
3	2023	30	2520	2979	99
4	2024	53	5226	10135	191
5	2025	52	5281	14244	273
Total		203	20351	33588	165

Table 2. Land Use of Allocated and Utilized Lands in Kish Island up to 2024.

Percentage Share %	Area (ha)	Land Use	Category
4/7	420	Residential	Lands Under Operation
0/3	26	Tourism (Non-Hotel)	
5/3	481	Mixed-Use	
0.3	31	Commercial (Market)	
0.9	81	Administrative (Exhibition)	
0.8	72	Services (Non-Restaurant)	
2.6	236	Industrial, Industrial Services & Warehousing	
1.2	112	Infrastructure Facilities	
8.5	765	Road Network	
7	628	Airport	
1.7	154	Commercial & Fishing Ports	
1.5	135	Sports (Recreational)	
1.4	129	Cultural Heritage	
0.1	7	Cultural	
0.1	11	Religious	
0.7	66	Educational	
0.1	7	Health & Treatment	
0.1	12	Restaurant	
2.2	195	Hotel	
7.4	670	Military	
47	4238	Total	
7.6	683	Parks & Green Spaces	Open Green Spaces
0.2	16	Agriculture (Flower & Plant Cultivation)	
3.6	320	Protected Areas	
41.6	3745	Unallocated & Vacant Plots	
53	4764	Total	
100	9000	Grand Total	

Based on the fact that more than 41 percent of Kish's land has not yet been allocated, there is significant potential for financing unfinished construction projects through real estate tokens. However, a fundamental challenge in developing economies is a paradox in which growth in transaction value does not necessarily indicate increased market depth. Data analysis shows that in such markets, nominal liquidity growth is often driven by asset inflation and value concentration in large-scale contracts, leading to the exit of small investors. In this context, tokenization, through "fractional ownership," provides an efficient mechanism for absorbing dispersed capital and deepening the market.

Free zones, due to their relative autonomy in policymaking, serve as ideal "living laboratories of innovation" for testing FinTech models. Kish Island, with its long-standing infrastructure and tourism development background, has strong potential to adopt this emerging technology. Benchmarking against regions such as Malta and Dubai shows that free zones can become international investment hubs through digital assets. Moreover, the 2024 statistical yearbook of the Kish Free Zone Organization also emphasizes the necessity of leveraging this capacity to achieve sustainable development.

1.1. Problem Statement and Research Objectives

The fundamental problem of this study is to investigate the role of real estate tokenization in addressing structural liquidity challenges and capital concentration risk in the real estate market of Kish Island. Statistical data analysis indicates that the Kish market is facing a "growth paradox," whereby 73 percent of the total transaction value of 33588 billion Rials during the 2021–2025 period has been concentrated in the last two years alone. At the same time, the average contract value has surged by ~259 % (3.6-fold) to 273 billion Rials, indicating a shift toward high-value contracts and reduced liquidity for the broader investor base [23, 27]. Furthermore, according to land statistics from the Kish Free Zone Organization, 3745 hectares (41.6 %) of Kish Island's land area remain unallocated or vacant, providing substantial potential for crowd-based financing through tokenization. This trend has resulted in the exclusion of small investors and the freezing of capital in large-scale assets. To address this inefficiency, the establishment of the KishConnect digital ecosystem is a strategic necessity. A review of the literature indicates that tokenization has evolved from an experimental phase into a permanent infrastructure within financial markets [6, 28]. While traditional models face settlement periods of 30 to 90 days and transaction costs of 5 to 7 percent, modern blockchain-based solutions enable real-time settlement (T+0) and could potentially reduce transaction overhead (reportedly up to 60–65 % in automated DeFi protocols [14]); this is not validated for Kish [11, 28].

In recent years, research on real estate tokenization has increased significantly; however, despite consensus on its strategic importance as an efficient mechanism for mobilizing dispersed capital in developing economies such as free zones, there remains a fundamental research gap regarding how this

technology should be adapted to the institutional and regulatory structures of Special Economic Zones (SEZs) such as Kish Island [6, 20]. The success of this paradigm is not solely dependent on technology but is contingent upon the convergence of four layers: technology, law, economics, and governance.

Accordingly, given the importance of this issue, this study seeks to analyze leading experiences in Dubai and Singapore in order to provide a framework for linking "compliance technology" with the local requirements of the Kish real estate market [2]. By leveraging service-oriented architecture and standards such as ERC-3643, the KishConnect platform aims to provide the necessary infrastructure for managing tokenized assets, facilitating secondary market transactions, and implementing smart contracts, thereby significantly enhancing the efficiency of the Kish real estate market [26]. Therefore, the research objectives are structured around four axes: analysis of global RWA developments in 2026, diagnosis of Kish's transaction structure, assessment of retail participation capacity through fractional ownership, and utilization of the KishConnect platform architecture for intelligent asset management [2, 16].

2. Literature Review

The theoretical foundations and adoption framework of real estate tokenization define it as the process of converting property ownership rights into programmable digital tokens, thereby transforming illiquid assets into liquid financial instruments. Based on the TOE model, the success of this technology in free zones depends on the convergence of blockchain infrastructure, organizational readiness, and the regulatory environment [29]. In 2026, this technology has moved beyond the experimental stage and, according to Deloitte's report, is progressing toward a 4 trillion-dollar market [11]. The economic functions and international developments in this field in 2026 identify the main advantage of tokenization as the "democratization of investment" through fractional ownership, which reduces entry barriers by up to 99 percent [24, 30]. The use of smart contracts, in addition to automating profit distribution, reduces operational costs by up to 65 percent and increases capital circulation speed in unfinished projects (such as the Avid tokenization project in Kish) [25].

Global case analysis shows that tokenization has evolved from an experimental phase into a permanent infrastructure of capital markets [27]. Projects such as RealT and St. Regis Aspen, by fractionalizing large assets into $50 units, have enabled real estate liquidity through secondary digital markets [27, 31]. At the global level, regulatory frameworks have matured, including the following examples:

Dubai (VARA Model and ARVA Assets): Dubai, through the Virtual Assets Regulatory Authority (VARA), has established the most comprehensive multi-layered framework for RWA tokenization [15]. One of Dubai's most important achievements is the definition of a new legal category called ARVAs in May 2025 [15, 32]. This regulation separates real-asset-backed tokens (such as real estate) from traditional securities [32]. This distinction allows the Kish Free Zone to design regulations tailored to the nature of real estate without becoming entangled in the complexities of securities law.

In February 2026, the second phase of Dubai's tokenization project, involving integration with the Dubai Land Department (DLD), was launched, enabling secondary resale of real estate tokens [15]. This model demonstrated that liquidity is not created solely through tokenization; rather, it requires real-time integration between blockchain ledgers and government land registries [32]. The VARA-regulated framework has also created a pathway for startups to operate under the supervision of a licensed VASP (as a sponsor), thereby reducing initial costs and time-to-market [15, 32].

Singapore (Project Guardian and Real-Time Settlement): This project is led by the Monetary Authority of Singapore (MAS) and focuses on institutional infrastructure [36, 21]. In 2026, Singapore launched a pilot for tokenized government bonds using wholesale central bank digital currency (CBDC) for settlement [28]. This initiative is designed to mitigate settlement risk – an approach that could be replicated in Kish through similar "Kish Coin" models for real-time real estate settlement [26, 28].

In 2025, Singapore also released an operational roadmap for tokenized funds, under which major institutions such as JPMorgan and Standard Chartered manage assets within an environment known as GFF [18, 32]. This project emphasizes asset standardization and enables tokenized assets on one platform to be tradable or usable as collateral on other financial platforms [32].

Other Leading and Comparable Projects Switzerland (DLT Law): Switzerland adopts a technology-neutral approach. Under the Swiss DLT Act, the concept of "ledger-based securities" has been introduced, making on-chain token transfers legally equivalent to definitive ownership transfer in the physical world [32].

United Kingdom (Digital Securities Sandbox): The UK has established a Digital Securities Sandbox (DSS) under the supervision of the FCA, allowing banks to test the full lifecycle of an asset – from issuance to settlement – on blockchain infrastructure [32].

Hong Kong: The Securities and Futures Commission (SFC) has approved the first real estate tokenization products and implemented a strict licensing system to maintain the stability and value of RWA tokens [32].

Based on the data in the above table, it can be concluded that the successful development of the KishConnect ecosystem requires a hybrid benchmarking approach across these countries:

- Dubai Model (VARA): Kish can facilitate the issuance process by defining an independent legal category for real estate tokens outside the

constraints of traditional capital market regulations [32];

- Singapore Model (MAS): To address the challenge of high-value contracts (273 billion Rials), the use of real-time settlement mechanisms is essential to eliminate the 30–90 day registration waiting period [18, 28];
- Swiss Model: By adopting the concept of ledger-based securities, it can be ensured that token transfers within the KishConnect blockchain are legally recognized as equivalent to definitive ownership transfer in the Kish land registry system [32].

A summary comparison of various tokenization models across different countries is presented in Table 3 as follows:

Table 3. Comparative Summary of Real Estate Tokenization Models in Different Countries

Indicator	UAE (Dubai)	SGP	SWZ	UK
Main Regulatory Authority	VARA	Monetary Authority of Singapore (MAS)	FINMA	FCA
Reference Legal Framework	ARVA Regulations	Project Guardian	Swiss DLT Act	Digital Securities Sandbox (DSS)
Real Estate Token Classification	ARVAs (distinct from securities)	Institutional digital assets with standards	Ledger-based securities	Experimental digital securities
Settlement Mechanism	Secondary market linked to Land Department (DLD)	Real-time settlement using wholesale CBDC	Direct ownership transfer on ledger	Full lifecycle testing in sandbox environment
Key Advantage for Kish	Legal separation of tokens from securities markets	Elimination of counterparty risk via CBDC-based settlement	High legal enforceability and banking acceptance	Ability to test business models prior to final regulation
Governance Approach	Asset-specific and specialized	Infrastructure- and liquidity-oriented	Technology-neutral	Experimental and phased

The localization of these models within the KishConnect ecosystem and the Avid project is not merely a technical adaptation, but an attempt to establish "institutional alignment" between the legal authority of the Kish Free Zone and blockchain infrastructure to democratize investment [2, 26]. Therefore, in the literature review, the advantages of tokenization can be redefined along four fundamental dimensions [7]:

Financial dimension (democratization): Tokenization, by fractionalizing large assets into smaller units, reduces the minimum entry capital from large sums to minimal amounts (e.g., 100 USD). This structurally improves market liquidity by absorbing dispersed small-scale capital [7, 10].

Operational dimension (efficiency): Smart contracts may support a potential reduction in transaction overhead (reportedly up to 60–65 % in automated DeFi protocols [14]); this figure is not established for real-estate transactions in Kish. Automated protocols could also support a projected settlement-cycle reduction (from 30–90 days to under 10 seconds [14]), which has not been observed or validated for KishConnect.

Informational dimension (transparency): Blockchain infrastructure, by creating a "single source of truth," enables immutable tracking of ownership and asset returns, eliminating fraud risk and pricing opacity [7, 16].

Institutional dimension (financing): Tokenization has enabled new instruments such as Initial Real Estate Token Offerings (IRET) for phased financing of large-scale projects, preventing capital stagnation in non-productive land assets [31].

Recent systematic studies based on the TOE model identify the success of tokenization as dependent on the balance between technological maturity, organizational readiness, and the regulatory environment [29]. Key findings in 2026 include:

Creation of a new asset class: Real estate tokens exhibit behaviour different from real estate investment trusts (REITs), with strong potential for portfolio diversification and optimization of risk-return ratios [20].

Collateral mobility: In 2026, with regulatory reforms such as UCC Article 12, these tokens are recognized as valid collateral for bank loans, unlocking trillions of dollars in previously illiquid real estate capital.

Cost efficiency at Layer 2: Transferring profit distribution to Layer 2 networks (such as Polygon) has reduced operational costs from large amounts to approximately 0.01 USD, making retail participation economically viable.

Realizing these potentials in Kish requires a shift from "paper-based registration" to "programmable ownership," so that tokenization evolves beyond a transfer innovation into a core driver of market efficiency [25].

Despite the transformative potential of this approach, its implementation faces significant challenges. Technological risks, vulnerabilities in smart contracts, and oracle manipulation risks are among the threats that can compromise the security of tokenized assets [7]. In 2026, a major barrier identified for implementation is the legal and registry gap, which reflects the lack of real-time alignment between blockchain transactions and land registry records. Without harmonized legal frameworks, a token may be transferred technically but not recognized as a legal transfer of ownership [33].

On the other hand, adoption and user experience represent another challenge, as the complexity of digital wallet management and the lack of global valuation standards (such as RICS and IVS 2026) lead some traditional investors to remain cautious toward this model [7].

3. Research Methodology

The present study is applied in terms of its objective, aiming to evaluate the liquidity potential of

the Kish market through real-world assets (RWA), and employs a mixed-method approach (descriptive-analytical and comparative) based on secondary sources to investigate the liquidity gap in the real estate market [34, 35]. The analytical framework of the study is grounded in the technology acceptance framework based on the TOE model (Technology, Organization, Environment) in order to examine the interaction between blockchain infrastructure and the organizational capacities of the Kish Free Zone.

To explain the mechanism of attracting small-scale capital into large-scale assets in Kish, the economic concept of long-tail distributions was used to enable fractional ownership. Quantitative and documentary databases were utilized for data collection and analysis. The quantitative dataset is derived from the official transaction reports of the Kish Investment and Development Company (2021–2025), comprising details of 203 sales contracts with a total value of 33588 billion Rials, which were analyzed using statistical methods. In addition, the institutional basis of the research is founded on the 2024 statistical yearbook of the Kish Free Zone Organization, particularly data indicating a significant potential of 3745 hectares (41.6 %) of unallocated land for crowdfunding through RWA. The data analysis process in this study was implemented at three levels:

Trend and descriptive statistical analysis: Examination of the "liquidity paradox," where, despite fluctuations in the number of contracts, the average transaction value increased 3.6-fold from 76 billion Rials (2021) to 273 billion Rials (2025), creating a significant barrier to entry for retail investors.

Comparative and content analysis: Comparison of traditional market indicators with 2026 blockchain capabilities (such as the ERC-3643 standard). In this phase, potential reductions in administrative overhead and projected settlement-cycle reduction (from 30–90 days to under 10 seconds via automated protocols [14]) were considered as scenarios, not measured Kish outcomes.

Ecosystem modelling: The technical framework of the KishConnect ecosystem was used to connect long-tail retail capital to large-scale projects based on the concept of fractional ownership. The KishConnect ecosystem and the Avid project were analyzed based on ERC-3643 and ERC-1400 technical standards to evaluate the alignment between theoretical frameworks and real-world market implementation requirements [19, 26].

In this study, in order to operationally analyze tokenization in Kish and examine technology acceptance based on the TOE model, the analysis was conducted across three layers [29]:

Technological dimension: Focus on compliance-native protocols such as ERC-3643 and Layer 2 networks; any transaction-cost savings are prospective and require empirical evaluation in a Kish regulatory sandbox.

Organizational dimension: Evaluation of the institutional capacity of the Kish Free Zone and the readiness of the KishConnect ecosystem to manage programmable assets and automate profit distribution.

Environmental dimension: Examination of legal frameworks by benchmarking Dubai's ARVAs model to distinguish real estate tokens from traditional securities and enhance market acceptance.

4. Results and Data Analysis

The examination of real estate and land transaction data in Kish (2021–2025) was conducted to perform a structural analysis and investigate the liquidity paradox during this period based on three key variables: number of transactions, transaction area, and monetary value. The findings indicate a paradigm shift from "retail transactions" toward "capital concentration." The statistical findings extracted from the official reports of the Kish Investment and Development Company are summarized in Table 4 and Figs. 1 and 2.

Table 4. Annual Trend of Real Estate and Land Transactions at Kish Investment and Development Company (2021–2025).

Year	Area (m²)	Avg. Value per Contract (Bln. Rls.)	Transaction Value (Bln. Rls.)	Number of Contracts	Average Price per m² (Bln. Rls.)	Annual Growth Rate
2021	4383	76	3060	40	0.7	–
2022	2941	113	3170	28	1.0	4 %
2023	2520	99	2979	30	1.18	-6 %
2024	5226	191	10135	53	1.93	240 %
2025	5281	273	14244	52	2.69	41 %
Total	20351	165	33588	203	1.65	–

The analysis of real estate and land transaction data in Kish Island indicates a significant paradigm shift from stability toward "asset inflation" and "severe capital concentration." The statistical results presented in Table 5 and Figs. 1 and 2 reflect what can be described as the "liquidity paradox". The fundamental cause of this condition is the market's shift toward large-scale contracts. According to calculations, the average contract value increased 3.6 times, from 76 billion Rials to 273 billion Rials. From the perspective of 2026 international standards, this accumulation of value represents a "critical entry barrier," effectively excluding retail investors and concentrating the market in the hands of large-scale actors [3, 30].

Furthermore, the approximately 3.8-fold increase in price per square meter (from 0.7 to 2.69 billion Rials) indicates that market growth in Kish, unlike leading markets, is driven by "asset inflation" rather than "asset velocity," potentially increasing the risk of liquidity stagnation [10]. International experiences in Dubai and Singapore provide comparative examples, but do not establish Kish-specific outcomes; fractional

ownership is one potential approach to market access [7, 15, 18].

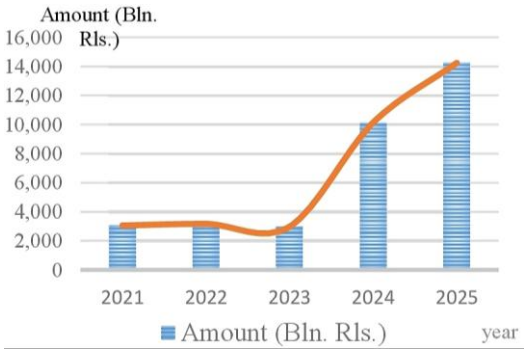


Fig. 1. Trend of Transaction Value and Asset Inflation in Kish Investment and Development Company (2021–2025).



Fig. 2. Analysis of Contract Value, Sold Area, and the Purchasing Power Gap (2021–2025).

The implementation of the KishConnect ecosystem based on the ERC-3643 standard, by converting these 273 billion Rials contracts into fractional tokens, can:

- Reduce entry barriers by up to 99 percent and enable participation from the "long-tail" of retail investors [19, 26];
- Projected settlement cycle reduction (from 30–90 days to under 10 seconds via automated protocols [14]); this is a prospective scenario;
- Utilize the potential of 3745 hectares (41.6 %) of unallocated land to operationalize the IRET model for financing development projects and prevent capital lock-in in undeveloped land assets [22, 25].

Analysis of temporal capital concentration and environmental drivers (Tables 5 and 6) indicates an abnormal surge in capital flow. This rapid accumulation reflects a transition into a phase of "inflationary growth," which in traditional models increases the risk of "liquidity freezing" due to long settlement cycles (30 to 90 days) [27, 28]. According to the 2024 statistical yearbook, Kish's structural dependence on smart tourism and international trade makes it an ideal environment for adopting FinTech tools. The existence of 16701 accommodation beds and high internet penetration further strengthen its potential to become a "living laboratory of innovation" [14, 15]. In this context, tokenization is not merely a transactional tool but a necessity for liquidating 3745 hectares (41.6 %) of unallocated land on the island, preventing capital lock-in in non-productive assets. By benchmarking the ARVAs model in Dubai, this large volume of capital can be directed toward unfinished development projects [17, 22].

Table 5. Comparative Share of Recession and Rapid Growth Periods in Total Transaction Value.

Transaction Value (Bln. Rls.)	Share of Total Market (%)	Market Condition	Time Period
9209	27 %	Recession and relative stability	2021–2023
24379	73 %	Rapid growth and asset inflation	2024–2025

Table 6. Analysis of Environmental Drivers Affecting the Kish Real Estate Market

Economic Driver	Strategic Impact on Real Estate Market	Necessity of Tokenization
Smart tourism	Increased demand for hospitality and hotel units	Crowdfunding of new projects
International trade	Growth in value of commercial and office land	Facilitation of foreign investment inflows
Infrastructure development	Increased attractiveness for residency and tourism	Tokenization of public urban assets
Free zones (SEZs)	Flexibility in currency and registration regulations	Creation of a regulatory sandbox

Comparative analysis of Tables 7 and 8 shows that the traditional Kish market, with a 273 billion Rial entry threshold, may constrain access for retail investors. The proposed KishConnect ecosystem based on ERC-3643 is a potential approach for addressing this bottleneck, subject to legal and empirical evaluation. Fractional ownership and a digital secondary market could support financial inclusion, but outcomes have not been measured in Kish. Potential transaction-overhead reductions reported in automated DeFi protocols [14] should not be interpreted as verified real-estate or Kish-specific cost savings. Token collateralization is likewise a prospective use case [17, 25].

The proposed business model for implementing real estate tokenization in Kish Island (based on the Business Model Canvas framework) defines tokenization as the process of converting property rights into programmable digital tokens on a blockchain, enabling fractional ownership, instant liquidity, and investment democratization [17, 18]. Given the availability of 3745 hectares (41.6 %) of unallocated land and the average contract value of 273 billion Rials in 2025, this model represents a

fundamental mechanism for attracting both retail and institutional capital. Its dimensions are outlined Table 9.

Table 7. Comparison of Market Efficiency Indicators (Traditional vs. Tokenized Market, 2026).

Market Efficiency Indicator	Traditional Kish Real Estate Market	KishConnect Ecosystem (Tokenized)
Entry threshold	High (avg. 273 billion Rials)	Very low (50–500 USD models)
Settlement time	30–90 days (physical registration)	Projected (not empirically validated): 30–90 days to under 10 seconds via automated protocols [14]
Transparency	Periodic and traditional reporting	Single Source of Truth (SSOT) on ledger
Liquidity	Limited to large-scale buyers	24/7 digital secondary market

Table 8. Strategic Benefits for Real Estate Market Participants in Kish.

Item	Economic Effect	Functional Interpretation in 2026
1	Smart collateralization using real estate tokens (DeFi collateral)	Smart collateralization
2	Potential reduction in transaction overhead (reportedly up to 60–65 % in automated DeFi protocols [14]); not established for Kish	Operational cost reduction
3	Protocol-level implementation of KYC/AML rules (ERC-3643)	Smart compliance
4	Fair distribution of property inflation gains among citizens	Democratization

The implementation roadmap proposed for establishing tokenization in the Kish Free Zone is designed as a five-stage process, inspired by Singapore's Project Guardian model, and includes the following phases [10, 37]:

Phase 1 (Selection and Evaluation): Identification of flagship projects (such as hotels and undeveloped land parcels in Kish) with transparent legal documentation and high profitability potential [15, 18];

Phase 2 (Sandbox Environment Creation): Deployment of a dedicated platform under the supervision of the Kish Free Zone Organization to test smart compliance mechanisms and on-chain identity verification [17, 33];

Phase 3 (Initial Real Estate Token Offering – IRET): Attraction of domestic retail capital through the democratization of ownership in land and infrastructure projects with minimal entry thresholds [7, 18];

Phase 4 (Secondary Market Activation): Launch of a 24/7 trading gateway with real-time settlement (T+0) to enhance liquidity in the island's real estate market [13, 18, 27];

Phase 5 (Internationalization): Leveraging the free zone's currency flexibility and Dubai's ARVAs model to attract foreign investors through digital payment infrastructure and smart collateralization mechanisms [13, 17, 34].

The establishment of the tokenization paradigm in Kish Island goes beyond a technical transformation; it represents a structural redesign of risk and return within the free zone economy [26, 32]. This model introduces both significant advantages and key challenges in the real estate market.

4.1. Key Advantages for Investors and the Financial Ecosystem

Democratization and financial inclusion: Through fractionalization of large-scale assets in Kish (where the average contract value reached 273 billion Rials in 2025), entry barriers are reduced by up to 99 percent [16, 18]. In 2026, leading platforms enable participation in luxury real estate with investments ranging from 50 to 1000 USD, thereby attracting the "long-tail" of retail capital into island development projects [18, 36].

Liquidity acceleration and real-time settlement: Unlike traditional 90-day transaction cycles, tokenization – through 24/7 secondary markets and Atomic Settlement mechanisms – could support a projected settlement cycle reduction (from 30–90 days to under 10 seconds via automated protocols [14]); this is not empirically validated for Kish. This significantly minimizes the risk of "capital lock-in" in medium-term development projects in Kish [7, 26].

Smart portfolio diversification: Investors can distribute their capital across multiple projects (commercial, hotel, and undeveloped land assets) instead of concentrating it in a single property unit, thereby optimizing risk-return profiles under market volatility conditions [7, 18].

Transparency and Single Source of Truth (SSOT): Immutable recording of all transactions, ownership records, and return histories on blockchain infrastructure eliminates fraud and document forgery risks while enhancing institutional trust [7, 16]. In addition, smart contracts enable automated rental income distribution without intermediaries [14, 18].

4.2. Challenges and Risks of Implementation

Legal and registry gap (Deed-Ledger Gap): The most significant barrier in 2026 is the lack of real-time alignment between blockchain ledgers and traditional land registry systems [33]. Without a regulatory sandbox in the free zone, token transfers may not necessarily be recognized as legal ownership transfers [7, 17, 38].

Technological risks and emerging threats: Smart contract vulnerabilities and oracle manipulation risks in real estate pricing represent critical challenges [28, 33]. In 2026, the use of MPC security models and continuous AI-based auditing has become essential to mitigate cyber risks [12, 19].

Initial liquidity risk and market depth limitations: The success of the secondary market in Kish depends on sufficient participation and the presence of market makers; otherwise, early-stage trading may suffer from high spreads and low liquidity [13, 17].

Social adoption paradox: The transition from "paper-based deeds" to "programmable assets" requires a fundamental shift in investor risk perception. Lack of education and the complexity of digital wallet management may reduce adoption rates in early phases [32, 33].

Table 9. Value Creation Dimensions of Real Estate Tokenization Implementation in Kish Island (Business Model Canvas Framework).

Row	Dimension	Value Creation in Real Estate Tokenization in Kish Island
1	Value Proposition	Structural liquidity: Converting illiquid assets into tradable digital units in a 24/7 secondary market, with a projected settlement cycle reduction (from 30–90 days to under 10 seconds via automated protocols [14]); not validated empirically for Kish.Fractionalization: Reducing minimum entry capital from billions to small units (e.g., 100 USD), enabling mass participation in large projects [18, 27].Compliance-as-Code: Use of ERC-3643 standard to automate KYC/AML and tax regulations at protocol level, ensuring transparency and legal security [13, 17].Activation of idle assets: Establishing crowdfunding infrastructure for undeveloped land and unfinished projects (such as the Avid project) to unlock liquidity [28, 34].
2	Customer Segments	Retail investors (long tail): Individuals with limited capital seeking inflation hedging in Kish's high-yield market [7].Institutional investors and funds: Banks and investment funds requiring digital asset portfolio management and smart collateralization [13, 27].Developers and construction companies: Project operators seeking faster and cheaper financing than traditional bank loans [7, 15].Asset owners and stakeholders: Individuals willing to liquidate part of their asset value without full sale [18].
3	Channels	KishConnect platform: Core Token Asset Lifecycle Management (TALM) system based on service-oriented architecture for issuance and exchange [26, 28]. Digital brokerage network: Cooperation with licensed financial entities in the free zone to enhance secondary market depth [13]. Smart tourism ecosystem: Use of airport and hotel information gateways to attract non-resident investors (incoming travelers) [17, 26].
4	Revenue Streams	Issuance fees: Charges for converting physical assets into tokens based on legal complexity [7]. Secondary market commissions: Percentage fees from each transaction [7, 13]. Management and compliance fees: Periodic fees for automated dividend distribution, ledger maintenance, and on-chain KYC services [7, 26].
5	Key Activities	Compliance protocol development: Implementation of ERC-3643 to embed free zone regulations and transfer restrictions in smart contracts [13, 17].Secondary market management: Ensuring instant liquidity via Atomic Settlement models to eliminate counterparty risk [13, 18].AI valuation services: Use of data oracles (e.g., Chainlink) for real-time NAV updates and market pricing alignment [12, 28].Legal support services: Establishment of SPVs for each project to ensure asset separation and investor protection [13, 17].
6	Key Resources & Partners	Resources: Digital trading platform, legal blockchain-real estate experts, and access to 2024 statistical yearbook data for risk analysis [15, 17]. Government partners: Kish Free Zone Organization as regulator for sandbox licensing and land registry integration; Kish Investment and Development Company as project implementer [13, 17]. Technical and financial partners: Blockchain developers, licensed crypto exchanges, RWA-friendly banks, and smart auditing institutions [13, 26].
7	Cost Structure	Upfront costs: Legal structuring of SPVs, smart contract security audits, and Layer 2 infrastructure deployment [7, 13]. Operational costs: Network maintenance, regulatory updates, and digital marketing for attracting international capital [7, 28].

5. Conclusion and Recommendations

The findings of this study indicate that the real estate market of Kish Island between 2021 and 2025 has experienced a structural paradox; that is, alongside an approximately 365 % increase in annual transaction value reaching 14244 billion Rials in 2025, market accessibility for retail investors has significantly declined. Statistical data show that 73 percent of total liquidity over the period is concentrated in the last two years, reflecting the dominance of large-scale actors and an increased risk of capital lock-in within high-value contracts.

Accordingly, the fundamental challenge in Kish is not merely price inflation, but a structural shift toward a market with limited liquidity and reduced public participation. This condition is consistent with the findings of Baum (2020) and OECD reports, which consider tokenization not merely a technical tool but a paradigm shift in the ownership and financing of

real-world assets [6, 27]. Within the context of Kish Island, converting physical land assets into digital tokens can bridge the gap between retail capital and large-scale projects through the concept of "tail connection," thereby steering the market toward a smart digital economy [7, 18].

From a strategic perspective, given that 3745 hectares (41.6 %) of Kish Island's land remains unallocated, tokenization offers substantial potential for financing construction and infrastructure projects without full reliance on bank loans. In 2026, benchmarking Singapore's Project Guardian and Dubai's ARVAs model has demonstrated that the success of this paradigm in free zones depends more on "smart regulation" than on technical infrastructure, particularly in addressing the Deed-Ledger Gap between physical deeds and blockchain ledgers [10, 13]. The implementation of the ERC-3643 standard within the KishConnect ecosystem can embed compliance rules at the protocol level, ensuring legal security in ownership transfers and strengthening institutional trust [19, 38].

The results motivate real estate tokenization as a potential mechanism for addressing structural liquidity challenges and democratizing investment in the Kish Island real estate market [6, 7]. By enabling fractional ownership and establishing a 24/7 digital secondary market, this model has the capacity to channel dispersed retail capital toward infrastructure development and flagship projects such as the Avid project [26, 28]. However, the operational success of this model requires a transition from the pilot phase to the establishment of a "regulatory sandbox" within the Kish Free Zone Organization, enabling the testing of real-time settlement mechanisms and on-chain identity verification in a controlled environment [34, 38]. If these requirements are met and advanced 2026 protocols are adopted, Kish Island can emerge as a "living laboratory of innovation" and a regional pioneer in tokenized real-world assets, while simultaneously attracting domestic and international investment and ensuring sustainable economic growth and financial transparency [14, 17, 28].

The validation of this study's findings must consider three fundamental limitations:

First, reliance on a transaction database from a single reference company, which limits the generalizability of the results to the broader Kish real estate market [27];

Second, limited access to detailed microdata from the 2024 statistical yearbook and the 2025 performance reports, which constrained the analysis to macro-level indicators [15, 24];

Third, the secondary nature of the data and the absence of field surveys on risk perception and behavioural preferences of retail investors, which may be addressed in future research [29]. International references and sources such as Anderson [36], Creswell [34], NeosLegal [32], and Zheng [33] are used only as theoretical, methodological, or jurisdictional analogs; they do not provide direct empirical validation of a Kish tokenization platform. Claims about settlement speed, overhead, investor thresholds, and cross-border flows are therefore projections or benchmarks, not observed KishConnect performance.

Based on the findings of this study, it is recommended that the Kish Free Zone Organization, leveraging its legal authority in investment facilitation, establish a regulatory sandbox for digital assets. This environment would enable the testing of "compliance-by-code" mechanisms based on the ERC-3643 standard, allowing on-chain identity verification and real-time settlement processes to be evaluated in a controlled and low-risk setting [10, 19, 25]. The key advantage of this approach lies in addressing the Deed-Ledger Gap between traditional land registry systems and blockchain-based ledgers through institutional learning by policymakers [18, 33].

Ultimately, due to its strategic position and regulatory flexibility, Kish Island is well-positioned to become the first "living laboratory of innovation" for tokenized real-world assets in Iran and the region [10, 14, 18].

References

[1]. S. Shen, Resilience and sustainability: Engineering solutions in disaster-affected real estate markets, *Sustainable Development*, Vol. 32, Issue 6, 2024, pp. 6748-6762.

[2]. Asset tokenization in financial markets: The next generation of value exchange, *World Economic Forum*, Geneva, 2025.

[3]. A. Baum, Tokenisation – The future of real estate investment? Oxford Future of Real Estate Initiative Research Report, Saïd Business School, *University of Oxford*, Oxford, 2020.

[4]. A. El Jaouhari, A. Samadhiya, A. Kumar, H. Chokshi, et al., Tokenization and the future of property investment: A new paradigm for real estate, *International Journal of Strategic Property Management*, Vol. 29, Issue 4, 2025, pp. 297-315.

[5]. V. J. Oza, A. Nikte, V. Bhanushali, U. Rote, Smart contracts and tokenization: Revolutionizing real estate transactions with blockchain technology, in *Proceedings of the 7th International Conference on Inventive Computation Technologies (ICICT)*, 2024, pp. 1163-1168.

[6]. The tokenisation of assets and potential solutions to addressing low real estate liquidity, *OECD Publishing*, Paris, 2020.

[7]. S. Kampakis, A CAIA mini course: A real estate focus on the crypto tokenization of real assets, *Portfolio for the Future*, 2026.

[8]. Emerging trends in real estate: Global outlook, *PwC*, London, 2023.

[9]. Real estate market dynamics and investment barriers, *World Bank*, Washington, DC, 2022.

[10]. Tokenization of real estate assets: Opportunities for fractional ownership, *Ernst & Young*, London, 2023.

[11]. Commercial real estate tokenization: The future of property investment, *Deloitte*, London, 2025.

[12]. From real estate to real world assets: The expanding tokenization opportunity, *McKinsey & Company*, New York, 2024.

[13]. R. Garcia-Teruel, H. Simón-Moreno, The digital tokenization of property rights: Challenges and opportunities, *Computer Law & Security Review*, Vol. 41, 2021, 105543.
[14]. F. Schär, Decentralized finance: On blockchain and smart contract-based financial markets, *Federal Reserve Bank of St. Louis Review*, Vol. 103, Issue 2, 2021, pp. 153-174.
[15]. Asset-referenced virtual assets (ARVA) rulebook, *Dubai Virtual Assets Regulatory Authority (VARA)*, Dubai, 2025.
[16]. Tokenization of real world assets: Opportunities and challenges, *McKinsey & Company*, New York, 2023.
[17]. D. Y. Aharon, S. Ali, M. Brahim, Connectedness at extremes between real estate tokens and real estate stocks, *Research in International Business and Finance*, Vol. 69, 2024, 102220.
[18]. Project Guardian: Enabling open and interoperable networks for asset tokenisation, *Monetary Authority of Singapore (MAS)*, Singapore, 2023.
[19]. The architecture of liquidity: T-REX protocol standards, *ERC-3643 Association*, 2026.
[20]. T. Farole, Special Economic Zones: Comparing Performance and Learning from Global Experience, *World Bank*, Washington, DC, 2011.
[21]. D. Schuurman, Bridging the gap between open and user innovation: Living labs, PhD Thesis, *Ghent University*, Ghent, 2015.
[22]. Land use and infrastructure report, in Statistical Yearbook of Kish Island, *Kish Free Zone Organization (KFZO)*, 2024.
[23]. Real estate and land transaction report (2021–2025), *Kish Investment & Development Company*, 2025.
[24]. Six-month performance report, *Kish Free Zone Organization*, 2025.
[25]. Whitepaper: Token business plan and infrastructure, *AVID Project*, 2024.
[26]. S. M. Hashemi, M. Noroozinejad, KishConnect: A blockchain-based service-oriented ecosystem for FinTech and tourism synergy" In *Proceedings of the 9th International Conference on the Development of Geographical Sciences, Tourism, and Sustainable Development in Iran*, Tehran, 2025. https://civilica.com/doc/2637108
[27]. Tokenization of real estate assets using blockchain technology, *SolidBlock*, 2026.
[28]. Project Guardian: Operational roadmap for tokenized funds, *Monetary Authority of Singapore (MAS)*, Singapore, 2025.
[29]. L. G. Tornatzky, M. Fleischer, The Processes of Technological Innovation, *Lexington Books*, Lexington, 1990.
[30]. Real estate tokenization in 2026: Facts and trends, *ScienceSoft*, 2026.
[31]. A. Jreisat, M. Mili, Blockchain in Real Estate: Theoretical Advances, *Springer*, Cham, 2024.
[32]. RWA tokenisation in the UAE: The complete legal & regulatory guide, NeosLegal, Dubai, 2026.
[33]. Z. Zheng, S. Xie, H. Dai, X. Chen, et al., Blockchain challenges and opportunities: A survey, *International Journal of Web and Grid Services*, Vol. 14, Issue 4, 2018, pp. 352-375.
[34]. J. W. Creswell, Research Design: Qualitative, Quantitative, and Mixed Methods Approaches, *Sage Publications*, Thousand Oaks, 2014.
[35]. R. K. Yin, Case Study Research and Applications: Design and Methods, *Sage Publications*, Thousand Oaks, 2018.
[36]. C. Anderson, The Long Tail: Why the Future of Business is Selling Less of More, *Hyperion*, New York, 2006.
[37]. D. Tapscott, A. Tapscott, Blockchain Revolution, *Penguin Random House*, New York, 2016.
[38]. Digital Securities Sandbox: Guidance and framework documentation, *UK Financial Conduct Authority*, London, 2023.

(003)

Hidden Edge-State Inference in Bitcoin Lightning Network Topologies

D. Patti [1], S. Monteleone [2], E. Russo [1] and M. Palesi [1]
[1] Department of Electrical, Electronic and Computer Engineering, University of Catania,
V.le A. Doria 6, 95125 Catania, Italy
[2] Department of Engineering, Niccolo Cusano University, Via Don Carlo Gnocchi 3, 00166 Roma, Italy
E-mail: davide.patti@unict.it, salvatore.monteleone@unicusano.it, enrico.russo@phd.unict.it,
maurizio.palesi@unict.it

Summary: The Bitcoin Lightning Network (LN) exposes topology and routing policies, but not directional liquidity. This work introduces a methodology using graph neural networks (GNNs) to infer amount-dependent directional channel-side viability from public graph features. Using the open-source Ultraviolet simulation platform, we generated 100-, 500-, and 1000-nodesynthetic LN topologies and transformed them into PyTorch Geometric-compatible graph tensors. We compare GraphSAGE with edge-only and endpoint-aware MLP baselines across 10 splits that keep both directions of each physical channel together. GraphSAGE achieves higher mean ROC-AUC and average precision in all nine size/amount settings, with a smaller advantage over the endpoint-aware baseline. These introductory results support further investigation of graph-based liquidity analysis for Bitcoin's second layer; their present scope is limited to synthetic channel-opening states, before payment traffic redistributes balances.

Keywords: Lightning network, Graph neural networks, PyTorch Geometric, Edge-state inference.

1. Introduction and Background

To preserve decentralization and trust lessness properties of the Bitcoin Network while allowing global-level scalability, the Lightning Network (LN) layer was proposed in 2016 [2]. The basic idea is to move transactions into off-chain channels [2], using the base layer primitives for funding and settlement. Once channels have been established, nodes can exchange a potentially unlimited number of signed transactions, while pathfinding and routing allow payments across multiple channels. Although nodes, channel capacities, and policies are publicly announced, the actual liquidity distributions between the two sides of a channel remain private [3]. Hidden directional balances make payment feasibility uncertain from public topology alone. They do not prevent the use of classical graph-routing algorithms: implementations combine Dijkstra's algorithm or its variants with estimated success probabilities and payment-attempt history [16, 8].

The idea proposed in this work is to consider the LN liquidity distribution as a graph-learning problem: topology, fees, capacity, and neighborhood position may help infer latent channel state. Prior work documents capacity concentration and a core-periphery structure in LN, motivating investigation of structural information [15]. Active probing can reveal channel balances, illustrating the tension between liquidity inference and privacy [14]. Here, predictions use graph features without active probing; simulator-visible balances provide training and evaluation labels.

Machine-learning models have been used to interpolate channel balances [5], temporal link classification has been applied to channel closures [6], and recent topology studies stress that LN remains sparse and dynamic for routing [7, 8]. Feichtinger et al. benchmark GNNs on LN node and edge prediction tasks and show that topology and neighboring-node information can improve prediction performance [13]. The present study examines amount-dependent directional spendability using controlled liquidity ground truth, complementing those public-property benchmarks. Other efforts, such as the open-source Ultra-Violet (UV) platform used here, provide environments for simulating LN while exposing channel state to re- searchers, supporting controlled pathfinding and routing studies [4].

The contributions of this introductory work are: (i) a controlled dataflow from profile-based synthetic LN bootstrap to GNN-ready directed graphs (Fig. 1, Table 1); (ii) amount-aware edge-state labelling for predicting channel-side viability; and (iii) a demonstrative comparison of GraphSAGE with two MLP baselines across different payment-amount thresholds and net- work sizes, using physical-channel-grouped splits (Table 2, Fig. 2). This provides a basis for later routing- support and liquidity-management studies.

2. Experimental Setup

Fig. 1 summarizes the proposed methodology: topology bootstrap produces a synthetic LN graph, the graph is transformed into directed tensors with amount-threshold labels, and the learning stage outputs usability scores and metrics. The same graph snapshots, features, labels, and train/test partitions are used for the three models.

Topology and initial balances: The bootstrap phase samples UV node profiles derived from real LN capacity, channel-size, and fee distributions [4].

Table 1 summarizes the profiles used for synthetic topology generation. The generated networks contain 100, 500, and 1000 nodes, with 885, 5187, and 10797 physical channels, respectively. Each channel is initially funded entirely by the opening node, with no initial transfer to its peer and a reserve of 1 % of channel capacity. Snapshots are collected after bootstrap, be- fore payment traffic redistributes balances; no balance randomization or rebalancing is applied. The resulting one-sided funding state is a deliberate introductory scenario, not an empirical distribution of operating LN balances. The task is evaluated for 1k, 10k, and 100k sats.

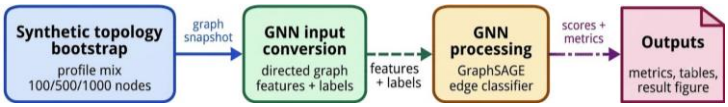


Fig. 1. Methodology dataflow.

Table 1. UV node profiles used in topology bootstrap.

Profile	Share	Node capacity	Channels / size
Small	60 %	100k–10M sats	1–10 / 1k–10M sats
Medium	30 %	10M–200M sats	10–100 / 1M–10M sats
Large	10 %	> 200M sats	100–500 / 10M–100M sats

Table 2. Amount-aware scores: mean ± standard deviation over ten physical-channel-grouped splits.

N	Amount	Usable	ROC-AUC			Average precision (AP)		
	(sats)		GraphSAGE	Endpoint MLP	Edge MLP	GraphSAGE	Endpoint MLP	Edge MLP
100	1k	50.0 %	0.901 ± 0.019	0.882 ± 0.022	0.732 ± 0.023	0.901 ± 0.019	0.885 ± 0.021	0.690 ± 0.021
100	10k	43.0 %	0.883 ± 0.016	0.865 ± 0.021	0.787 ± 0.016	0.844 ± 0.022	0.812 ± 0.024	0.690 ± 0.021
100	100k	40.8 %	0.894 ± 0.013	0.880 ± 0.018	0.811 ± 0.010	0.844 ± 0.021	0.812 ± 0.030	0.693 ± 0.020
500	1k	50.0 %	0.947 ± 0.005	0.919 ± 0.007	0.691 ± 0.019	0.947 ± 0.005	0.919 ± 0.007	0.654 ± 0.018
500	10k	43.7 %	0.938 ± 0.006	0.908 ± 0.007	0.752 ± 0.008	0.917 ± 0.009	0.873 ± 0.010	0.649 ± 0.008
500	100k	41.7 %	0.942 ± 0.005	0.912 ± 0.004	0.773 ± 0.009	0.917 ± 0.006	0.866 ± 0.006	0.652 ± 0.014
1000	1k	50.0 %	0.978 ± 0.002	0.955 ± 0.004	0.832 ± 0.010	0.978 ± 0.002	0.955 ± 0.003	0.804 ± 0.011
1000	10k	43.4 %	0.969 ± 0.003	0.944 ± 0.005	0.864 ± 0.006	0.960 ± 0.004	0.923 ± 0.008	0.793 ± 0.009
1000	100k	41.4 %	0.974 ± 0.003	0.953 ± 0.003	0.875 ± 0.003	0.963 ± 0.004	0.929 ± 0.004	0.793 ± 0.006

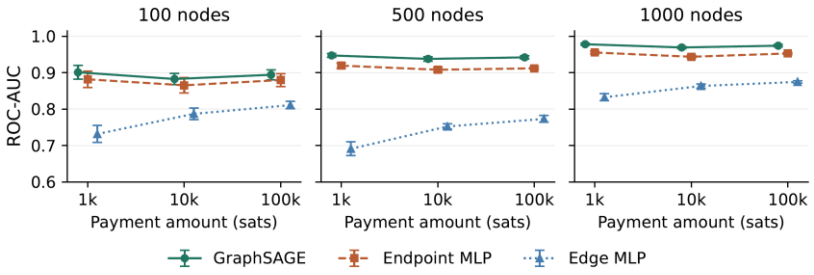


Fig. 2. ROC-AUC by network size and payment threshold. Error bars denote one standard deviation across grouped splits and model initializations; the vertical axis starts at 0.60.

Inputs: The GNN input is the known directed topology, node features (degree, incident capacity, local channel count, node capacity), and edge features (capacity, base fee, fee rate, time lock delta, policy flags, direction). Node capacity is the sum of incident channel capacities, not a node's private wallet balance. Local/remote balances and source/destination balances and spendable balances are excluded because they are not visible to external nodes. Features are standardized over the visible snapshot without using labels. Channel identifiers group the evaluation partitions but are not model features.

Output: Each channel is evaluated in both directions. For amount a, $y^a_{u\to v} = 1$ means that the u side can spend at least a sats toward v after reserve. With source balance $B_{u\to v}$ and channel reserve R, the label is

$$y^a_{u\to v} = \mathbf{1}[\max(0, B_{u\to v} - R) \geq a]$$

A zero label denotes insufficient source-side liquidity for that amount. A positive label captures a local

liquidity condition, not guaranteed end-to-end payment delivery, which also depends on other hops and routing policies. UV supplies this private ground truth for supervised training and evaluation; the models produce usability scores from the permitted features.

Models and training: The topology-aware model uses GraphSAGE [9], related to GCNs [10], in PyTorch Geometric [11]. GraphSAGE uses two neighborhood aggregation layers with 32 hidden channels and an edge decoder combining source and destination embeddings with edge features. The endpoint-aware MLP replaces these aggregation layers with two node-wise linear layers of the same width, keeping the activation pattern and edge decoder. It therefore receives both endpoints' node features and the edge features, but no neighboring-node information. The edge-only MLP retains the same hidden width and uses only edge features.

All three models use 0.10 dropout in their decoder, Adam with learning rate 0.01 and weight decay 10^{-4}, and class-weighted binary cross-entropy, with class weights estimated from training labels only. Each model is trained separately for each amount threshold for 100 epochs. Settings are fixed across network sizes and amounts, with no test-set tuning. The endpoint-aware baseline controls access to node information and atches encoder depth and decoder structure; parameter counts are 4769 for GraphSAGE, 3585 for the endpoint MLP, and 289 for the edge MLP, so model capacity is not exactly matched.

Evaluation protocol: For each network size, one synthetic snapshot generated with bootstrap seed zero is reused across 10 splits and model initializations. We split physical channels approximately 70/30 between training and testing, stratifying by the number of usable directions per channel. Both directions of every physical channel remain in the same partition, and all models share each partition. This is a transductive task: the full observed topology and feature matrix remain visible, but both directional labels of every test channel are withheld from training. It evaluates unknown channel states on a known graph, rather than generalization to unseen networks. Standard deviations describe variation across splits and initializations of that snapshot, not independent replications of the network-generating process.

Metrics: ROC-AUC, the area under the receiver operating characteristic curve, measures how well scores rank usable directions above unusable ones across decision thresholds. Average precision (AP) summarizes precision over increasing recall, making it informative when the usable and unusable classes are unequal [12]. Accuracy is also recorded using a score threshold of 0.5. Table 2 reports ROC-AUC and AP; Fig. 2 visualizes the ROC-AUC comparison.

Artifacts Reproducibility: The Ultraviolet repository includes a dedicated artifact folder with the scenario configurations and graph export/conversion workflow needed to regenerate the experiment inputs [4].

3. Results

Table 2 shows that payment amount changes the target itself: usable directions decrease from 50.0 % at 1k sats to 40.8 %–41.7 % at 100k sats. The exactly balanced 1k case follows from one-sided initialization: every channel has one funded side able to cover that amount. At larger amounts, some funded sides also become unusable because their reserve-adjusted capacity is insufficient. These proportions describe the selected boot- strap state, not measured liquidity availability on the operating LN.

The table also separates the effect of endpoint information from the larger comparison against edge only features. At 500 nodes and 1k sats, ROC-AUC is 0.947 for GraphSAGE, 0.919 for the endpoint MLP, and 0.691 for the edge MLP. Endpoint features account for a substantial improvement over the edge only baseline; the remaining GraphSAGE advantage is smaller. Across all nine settings, its mean ROC-AUC advantage over the endpoint MLP ranges from 0.014 to 0.030, consistent with additional predictive information in neighborhoods under this setup. This comparison does not establish that aggregation alone explains the entire residual gap, given the different parameter counts.

AP provides a complementary view of the ranked usable directions. At 500 nodes and 100k sats, Graph SAGE reaches 0.917 AP, compared with 0.866 for the endpoint MLP and 0.652 for the edge MLP, when only 41.7 % of directions are usable. Thus, the advantage is also visible in precision–recall performance for the stricter liquidity condition. At 1000 nodes and 100k sats, the corresponding AP values are 0.963, 0.929, and 0.793. At the fixed 0.5 decision threshold, accuracy in that setting is 0.927, 0.886, and 0.852, respectively.

Fig. 2 highlights the same model ordering in mean ROC-AUC for every network size and amount. Edge MLP ROC-AUC rises with the amount threshold, for example from 0.691 to 0.773 at 500 nodes. A plausible interpretation is that public capacity becomes more informative as the required amount increases. GraphSAGE and the endpoint MLP instead show a dip at 10k sats followed by a partial recovery at 100k sats, illustrating that discrimination need not vary monotonically with the payment amount. Each threshold defines a different classification target, so the curves do not measure a single fixed task becoming uniformly easier or harder.

The figure's error bars are smaller for the 1000-node snapshot: GraphSAGE obtains 0.978 ± 0.002 at 1k sats, versus 0.901 ± 0.019 at 100 nodes. This indicates less variation across the evaluated partitions and initializations in that snapshot. Because only one topology is sampled at each size, these results do not establish that increasing network size causes better pre- diction or that performance is stable across independently generated topologies. The overlapping repeated splits are not independent network replications, and no statistical-significance claim is made.

Scope and interpretation: This is an introductory work intended to demonstrate the methodology and motivate more realistic experiments. Initial one-sided balances can create correlations between funding direction, node profiles, and neighborhood structure; the observed performance may partly reflect these bootstrap-specific relationships. The results do not yet demonstrate inference of payment-driven balance dynamics, transfer to live LN data, or improved routing success. Likewise, withholding balances from model inputs does not establish a privacy guarantee: inference itself can reveal information about hidden channel states.

For Bitcoin's second layer, the practical motivation is to assess whether graph-derived scores can eventually supplement routing decisions, identify liquidity bottlenecks, and guide rebalancing. Such uses require validation on evolving balances and actual payment outcomes. A single-channel liquidity label is an intermediate analytical tool, rather than a direct measure of network payment reliability.

4. Conclusions and Future Work

This introductory study demonstrates an amount-aware graph-learning workflow for Lightning channel usability. The grouped evaluation retains an advantage for GraphSAGE, while the stronger MLP baseline shows that endpoint information explains a substantial part of the improvement over edge-only features. The findings motivate further study of graph context for liquidity analysis in Bitcoin second-layer systems.

Future work will include several scenarios closer to operating LN conditions: independently generated topologies and balance allocations, payment traffic that evolves channel states, temporal train/test splits, and real-topology-based simulations with partial observer knowledge. Payment-delivery and rebalancing benchmarks will then assess whether the inferred scores can improve liquidity use and payment reliability.

References

[1]. S. Nakamoto, Bitcoin: A peer-to-peer electronic cash system, 2008, https://bitcoin.org/bitcoin.pdf

[2]. J. Poon, T. Dryja, The Bitcoin Lightning Network: Scalable off-chain instant payments, 2016, https://lightning.network/lightning-network-paper.pdf

[3]. A. M. Antonopoulos, O. Osuntokun, R. Pickhardt, Mastering the Lightning Network, *O'Reilly Media*, Sebastopol, 2021.

[4]. D. Patti, S. Monteleone, E. Russo, M. Palesi, Timechain-level modeling and analysis of the Bitcoin Lightning Network, *Computer Networks*, Vol. 275, 2025, 111891.

[5]. V. Davis, V. Singh, E. Rossi, Channel balance interpolation in the Lightning Network via machine learning, in *Proceedings of the IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, 2025, pp. 1-7.

[6]. S. Antonelli, V. Davis, H. Rush, A. Potdevin, et al., Predicting channel closures in the Lightning Network with machine learning, *arXiv*, 2026, arXiv:2605.12759.

[7]. D. Valko, J. M. Gomez, Topology and network dynamics of the Lightning Network: A comprehensive analysis, *arXiv*, 2025, arXiv:2512.20641.

[8]. D. Valko, R. Paranjpe, J. M. Gomez, Outperforming Dijkstra on sparse graphs: The Lightning Network use case, *arXiv*, 2025, arXiv:2509.13448.

[9]. W. L. Hamilton, R. Ying, J. Leskovec, Inductive representation learning on large graphs, in Advances in Neural Information Processing Systems 30 (I. Guyon, U. von Luxburg, S. Bengio, H. Wallach, et al., Eds.), *Curran Associates, Inc*., Red Hook, 2017, pp. 1024-1034.

[10]. T. N. Kipf, M. Welling, Semi-supervised classification with graph convolutional networks, in *Proceedings of the 5th International Conference on Learning Representations (ICLR)*, 2017.

[11]. M. Fey, J. E. Lenssen, Fast graph representation learning with PyTorch Geometric, in *Proceedings of the ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.

[12]. J. Davis, M. Goadrich, The relationship between precision-recall and ROC curves, in *Proceedings of the 23rd International Conference on Machine Learning (ICML)*, 2006, pp. 233-240.

[13]. R. Feichtinger, F. Grötschla, L. Heimbach, R. Wattenhofer, Benchmarking GNNs using Lightning Network data, *arXiv*, 2024, arXiv:2407.07916.

[14]. S. Tikhomirov, R. Pickhardt, A. Biryukov, M. Nowostawski, Probing channel balances in the Lightning Network, *arXiv*, 2020, arXiv:2004.00333.

[15]. J.-H. Lin, K. Primicerio, T. Squartini, C. Decker, et al., Lightning Network: A second path towards centralisation of the Bitcoin economy, *New Journal of Physics*, Vol. 22, Issue 8, 2020, 083022.

[16]. B. Vu, The node operator's guide to the Lightning Galaxy, part 2: Node scoring and pathfinding, Lightning Labs, 2019, https://blog.lightning.engineering/posts/2019/11/07/routing-guide-2.html

(004)

Centralized Ledgers versus Decentralized Permissionless Blockchains: A Game-Theoretic Comparison

Kjell Hausken and Guizhou Wang
Faculty of Science and Technology, University of Stavanger, 4036 Stavanger, Norway
Tel.: +47 51 831632, fax: +47 51 831550
E-mail: kjell.hausken@uis.no; fintwang@outlook.com

Summary: This paper develops a three-period game with four unitary players: centralized intermediary *I*, aggregate proof-of-stake validation sector *V*, unitary attacker *A* (one pre-existing malicious coalition), and user *U*. The model infrastructure contains centralized regime *C*, permissionless blockchain regime *B*, and outside option *O*. In period 1, *I* chooses access and monitoring while *V* chooses validation stake; in period 2, *A* chooses nonnegative attack efforts; in period 3, *U* chooses a regime. Backward induction yields attack-deterrence conditions and a user choice set. Under constant absolute risk aversion, residual-loss variance lowers the user's certainty equivalent by $\rho/2$ times the variance. The user chooses *B* only when its risk-adjusted certainty equivalent, incorporating access, assurance, service quality, residual net loss, and the exogenous fee, is strictly highest.

Keywords: Blockchain security, Centralized intermediation, Proof of stake, Access, Credibility, Regime choice, Game theory.

1. Introduction

Digital ledger infrastructure allocates control over records, access, and execution. Centralized regime C can be fast and operationally efficient because one organization coordinates records and disputes, but an insider or external coalition may freeze accounts, censor transactions, reverse entries, or reinterpret contested records. Permissionless blockchain regime B replaces a single record keeper with public validation, cryptographic evidence, and token incentives, while introducing fees and governance or concentration risks across validators, interfaces, and protocols [1-4].

The analysis is deliberately user-centered. The feasible outcomes are centralized ledger regime *C*, permissionless proof-of-stake blockchain regime *B*, and the outside option *O*. Regimes *C* and *B* are not players. The unitary players are the legitimate centralized intermediary *I*, the representative validation sector *V*, the unitary attacker *A* and the representative user *U*. In *C*, the access choice a_C is the share of otherwise eligible users and transactions to which *I* supplies timely, nondiscriminatory ledger service. The research question is whether each regime makes profitable attack unattractive while delivering credible access at acceptable risk-adjusted user cost. Figure 1 depicts the full model infrastructure. The framework applies to payments, registries, digital identity, decentralized finance, and related ledger settings [5, 6].

2. Contribution

This paper develops a unified three-period game in which centralized access and monitoring, permissionless stake and attack deterrence, and a risk-averse user's regime choice are jointly endogenous. The contribution is the joint determination of regime-specific attack deterrence and the user's equilibrium regime choice, rather than a claim that one governance technology is universally superior.

References [1-6] provide the protocol, Bitcoin economics, proof-of-stake, governance, institutional, and decentralized-finance foundations. None combines endogenous centralized access and monitoring, validation stake, target-specific attacks against both regimes, and a risk-averse user's final choice. Hausken and Wang [7] is thematically related but structurally distinct: it studies a static multi-country macro-financial equilibrium across currencies, borrowing, taxation, inflation, and eleven asset classes.

3. Model and Sequence of Play

3.1. Unitary Players, Primitives and Three-Period Timing

All four players are unitary reduced-form decision makers with fixed identities throughout the game. Player *I* is one legitimate centralized intermediary. Player *V* is one aggregate proof-of-stake validation provider whose only strategic action is the validation-stake choice s_B; *V* has no attack or misconduct action in the baseline. Player *A* is a separate unitary attacker, equivalently one pre-existing perfectly coordinated malicious coalition; coalition formation and membership are fixed outside the model. Player *A* alone chooses the attack efforts e_C and e_B. Player *U* is one representative marginal user. No player changes identity or becomes another player. The target-specific attacks are separable, and *A* has no common effort budget in the baseline.

The compact strategy sets are $x_I = (a_C, m_C) \in X_I = [0, 1]\times[0, m^{max}]$, $s_B \in X_V = [0, s^{max}]$, $x_A = (e_C, e_B) \in X_A = [0, e_C^{max}]\times[0, e_B^{max}]$, and $j \in X_U = \{C, B, O\}$; all upper bounds are finite and positive. The choice a_C is the share of eligible users and transactions served in C, and m_C is I's costly monitoring effort. The choice s_B is the aggregate validation stake posted by V to validate B, at cost $C_s(s_B)$; it scales the attack-controlled position ηsB. The efforts e_C and e_B are the nonnegative resources A directs at compromising C and B, respectively. Because they are independently chosen, A may attack both, either one, or neither. Player U chooses exactly one $j \in \{C, B, O\}$.

The access choice $a_C \in [0, 1]$ is, for the representative U, the probability of receiving timely and nondiscriminatory account opening, continued use, and transaction execution in C. A larger a_C benefits U and can raise I's adoption revenue, but I incurs the nonnegative access-capacity cost $C_a(a_C)$, and the centralized attack rent $R_C(a_C)$ may rise. The monitoring choice $m_C \in [0, m^{max}]$ includes access-control review, audit trails, record-integrity controls, and evidence retention. It costs I through $C_m(m_C)$, raises attack detection $d_C(m_C)$, and may raise user-perceived assurance $h_C(m_C)$.

Blockchain access $a_B \in (0, 1]$ is broad and exogenous: U and similarly situated eligible users can submit and verify transactions subject to protocol rules and the fee f_B, and no player strategically rations a_B in the baseline. The user pays f_B, whereas V bears validation-resource cost $C_s(s_B)$. For $j \in \{C, B\}$, $q_j \geq 0$ is exogenous nonsecurity service quality – speed, throughput, interface usability, dispute handling, and interoperability – which benefits U through X_j; any provision cost is embedded in fees, revenues, or cost primitives. The continuous nondecreasing assurance functions $h_C:[0, m^{max}]\to[0, 1]$ and $h_B:[0, s^{max}]\to[0, 1]$ convert observable monitoring or stake into a normalized index of U's confidence that an accepted transaction is correctly validated and recorded, remains final, and remains available. Thus, what is assured is ledger integrity, finality, and availability – not zero loss. Their resource costs are already represented by C_m and C_s. Residual net loss, defined below, is the post-deterrence stochastic monetary loss borne by the user and must not duplicate fees, service-quality effects, or the strategic attack losses modeled elsewhere. Interface usability is one component of q_j and is not an additional payoff term.

Before play, common-knowledge primitives are fixed exogenously. A legal or contractual sanction $F_C>0$ is imposed on A after a detected and attributable attack on C. The detection function $d_C:[0, m^{max}]\to[0, 1]$ is continuous and nondecreasing. For B, $d_B \in [0, 1]$ is the probability that public evidence identifies a slashable attack; $\phi \in (0, 1]$ is the protocol-enforced fraction of the attacker-controlled stake confiscated or burned; and $\eta>0$ is the required attack-stake multiple relative to V's validation stake s_B. Thus, to attack B, A must control or procure the slashable position ηs_B and is assumed to internalize its expected loss. A slashing rule is an automatic protocol penalty applied to that attack-controlled stake when the attack is detected. The baseline does not endogenize a misconduct choice by V; adding such a choice would require enlarging V's strategy set and payoff function.

For each target $j \in \{C, B\}$, e_j^{max} denotes e_C^{max} when $j = C$ and e_B^{max} when $j = B$. Attack success $p_j:[0, e_j^{max}]\to[0, 1]$ and effort cost $c_j:[0, e_j^{max}]\to[0, \infty)$ are continuous and nondecreasing, c_j is convex, and $p_j(0) = c_j(0) = 0$. The centralized attack-rent function R_C maps [0, 1] into nonnegative real values, and the blockchain attack-rent function R_B maps nonnegative secured values into nonnegative real values; both are continuous and nondecreasing. The parameter $K_B \geq 0$ is the exogenous value secured by B. Exogenous per-use fees satisfy $f_C \geq 0$ and $f_B \geq 0$. Residual nonstrategic net loss in regime j is $\tilde{\ell}_j \sim N(\mu_j, \sigma_j^2)$, where μ_j is real and $\sigma_j^2 \geq 0$. It is residual because it remains after A's strategic attack has been deterred or the attacked regime has been removed from U's choice set. It includes post-deterrence outages, processing or finality errors, software or governance faults, unauthorized losses outside A's modeled attack, and recovery costs, net of recoveries. Fees, quality effects, and modeled strategic attack losses are excluded to avoid double counting. Positive realizations are losses; negative realizations are net recoveries. All monetary quantities use a common unit. Fig. 1 summarizes the model infrastructure, regimes, players, choices, and primitives.

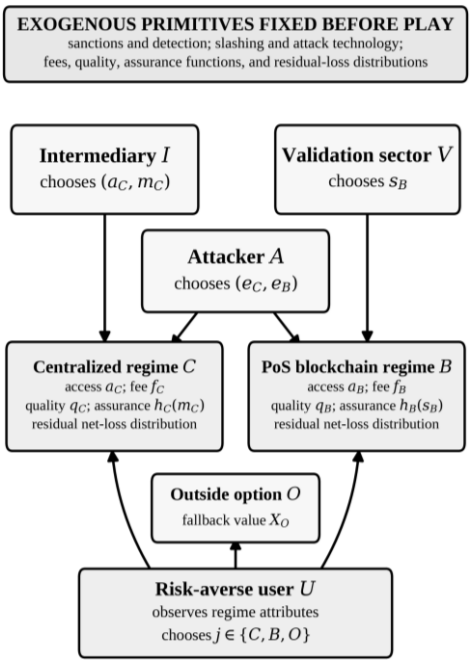


Fig. 1. Model infrastructure, regimes, unitary players, strategic choices, and exogenous primitives. C and B are regimes, not players.

The game has exactly three periods. In period 1, I and V choose simultaneously. In period 2, A observes both choices and selects its attack-effort vector. In period 3, U observes the induced attributes and

chooses C, B, or O; payoffs then follow. Fig. 2 shows the strategy sets and timing. A pure-strategy subgame-perfect equilibrium, when it exists, is found by solving the period-3 choice of U and then the period-2 choice of A by backward induction, while I and V satisfy a period-1 Nash fixed point in continuation payoffs.

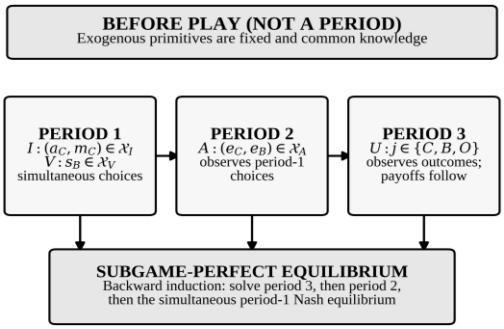


Fig. 2. Three-period game, strategy sets, and exogenous pre-play primitives.

3.2. Period-1 Objectives of I and V

Let $D_j \in \{0,1\}$ equal one when U selects regime $j \in \{C,B\}$ and zero otherwise. The period-1 players anticipate the continuation choices of A and U. Their adoption-contingent payoffs are revenue net of access, monitoring, and staking costs:

$$\Pi_I = D_C r_I(a_C, f_C) - C_a(a_C) - C_m(m_C), \quad (1)$$

$$\Pi_V = D_B r_V(s_B, f_B) - C_s(s_B) \quad (2)$$

The reduced-form revenues $r_I(a_C, f_C)$ and $r_V(s_B, f_B)$ are continuous and nonnegative on the stated domains. The nonnegative cost functions C_a, C_m, and C_s are continuous, increasing, and convex on the corresponding strategy intervals, with zero cost at zero. Capacity and security costs are incurred even if U selects another regime. The exogenous fees f_C and f_B may enter the revenues of I and V, but neither player sets a fee strategically.

3.3. Period-2 Objective of A and Attack Deterrence

For $j \in \{C,B\}$, define the pre-penalty attack payoffs by $H_C(e_C) = p_C(e_C)R_C(a_C) - c_C(e_C)$ and $H_B(e_B) = p_B(e_B)R_B(K_B) - c_B(e_B)$. These pre-penalty attack-payoff opportunities are present at the system level before the representative marginal user chooses, so D_j does not enter A's baseline payoff. The expected penalties conditional on initiating an attack are

$$L_C(m_C) = d_C(m_C)F_C, L_B(s_B) = d_B\phi\eta s_B \quad (3)$$

In C, $L_C(m_C)$ is the detection-weighted sanction borne by A. In B, $L_B(s_B)$ is the expected slashing loss borne and internalized by A on the attack-controlled position ηs_B. This loss is distinct from V's validation-stake cost $C_s(s_B)$ in (2). The target-specific and total attack payoffs are

$$\pi_{Aj}(e_j) = H_j(e_j) - \mathbf{1}_{\{e_j>0\}}L_j, j \in \{C,B\}, \quad \Pi_A = \pi_{AC}(e_C) + \pi_{AB}(e_B), \quad (4)$$

where $L_j = L_C(m_C)$ for $j = C$ and $L_j = L_B(s_B)$ for $j = B$. In (4), $\mathbf{1}_{\{e_j>0\}}L_j$ means that the indicator equals one when A initiates a positive-effort attack against regime j and zero when $e_j = 0$; the loss L_j is therefore imposed only when an attack is initiated.

Because the two targets are separable, A maximizes the sum of the target-specific payoffs. The compact strategy set and the payoff's upper semicontinuity – its value at zero effort is no smaller than the limiting payoff from arbitrarily small positive effort when $L_j > 0$ – ensure that a best response exists. Define the maximum pre-penalty attack payoffs as follows:

$$x_A^\star = (e_C^\star, e_B^\star) \in \arg\max_{x_A \in \mathcal{X}_A} \Pi_A(x_A), \quad G_C(a_C) = \max_{e_C \in [0, e_C^{\max}]} H_C(e_C), \quad G_B(K_B) = \max_{e_B \in [0, e_B^{\max}]} H_B(e_B) \quad (5)$$

With the tie-breaking convention that A chooses zero effort when indifferent, the attacker's target-specific best response is zero exactly when

$$e_C^\star = 0 \Leftrightarrow G_C(a_C) \le d_C(m_C)F_C, \quad e_B^\star = 0 \Leftrightarrow G_B(K_B) \le d_B\phi\eta s_B \quad (6)$$

Equation (6) separates strategic attack risk from residual operational and governance risk. Legal sanctions deter A in C, whereas stake exposure and slashing deter A in B.

3.4. Period-3 Objective of U and Equilibrium Closure

The baseline uses an attack-contingent period-3 choice rule. If A chooses positive equilibrium attack effort against C, U removes C from its choice set; if A chooses positive equilibrium attack effort against B, U removes B; if both equilibrium attack efforts are positive, only O remains. This is a behavioral choice rule, not technological unavailability. Let $a_j = a_C$, $h_j = h_C(m_C)$, and $f_j = f_C$ for $j = C$, while $a_j = a_B$, $h_j = h_B(s_B)$, and $f_j = f_B$ for $j = B$. Let $X_j(a_j, q_j, h_j)$ denote a real-valued deterministic gross user benefit and be continuous and nondecreasing in access, service quality, and assurance. Conditional on zero equilibrium attack effort, the user's monetary payoff is $X_j - f_j - \tilde{\ell}_j$. Under constant absolute risk aversion, $u(y) = -\exp(-\rho y)$, where y is monetary payoff and $\rho > 0$ is the coefficient of absolute risk aversion. With

normally distributed residual loss, the single-symbol certainty equivalent is exact:

$$\Psi_j = X_j - f_j - \mu_j - \frac{\rho}{2}\sigma_j^2,\ j \in \{C, B\},$$
$$\Psi_O = X_O \tag{7}$$

The outside-option value X_O is deterministic. Because $u'(y) > 0$ and $u''(y) < 0$, player U values more monetary payoff but, at a fixed mean, dislikes additional payoff variability. Its Arrow–Pratt coefficient of absolute risk aversion is $-u''(y)/u'(y) = \rho$. Holding $X_j - f_j - \mu_j$ fixed, higher σ_j^2 lowers Ψ_j by $(\rho/2)\sigma_j^2$. Thus larger ρ means stronger risk aversion. As $\rho \downarrow 0$, Ψ_j approaches expected net payoff, the risk-neutral limit. Player U solves

$$\mathcal{J}_U = \{O\} \cup \{C : a_C > 0, e_C^\star = 0\} \cup$$
$$\cup \{B : s_B > 0, e_B^\star = 0\},$$
$$j^\star = \tau\left(\arg\max_{j \in \mathcal{J}_U} \Psi_j\right), D_j = \mathbf{1}_{\{j^\star = j\}} \tag{8}$$

Here J_U always contains O; it contains C if $a_C > 0$ and $e_C^\star = 0$, and contains B iff $s_B > 0$ and $e_B^\star = 0$. Thus $s_B = 0$ makes B unavailable to U. The tie rule τ ranks O first, C second, and B third. Equation (8) defines $j^\star$ and D_j; exactly one indicator equals one.

The indicators D_C and D_B in (1)-(2) are therefore generated by U's choice rather than imposed externally. A period-1 Nash equilibrium in continuation payoffs satisfies

$$x_I^\star \in \arg\max_{x_I \in \mathcal{X}_I} \Pi_I(x_I; s_B^\star),$$
$$s_B^\star \in \arg\max_{s_B \in \mathcal{X}_V} \Pi_V(s_B; x_I^\star) \tag{9}$$

A profile satisfying (4)-(9) is a pure-strategy subgame-perfect equilibrium. The results below condition on existence of the period-1 Nash fixed point in (9); compact action sets alone do not guarantee that fixed point because the user's discrete choice can make continuation payoffs discontinuous at indifference boundaries. The surplus measure used below is net user surplus rather than aggregate social welfare:

$$S_j = \Psi_j - \Psi_O, j \in \{C, B\},$$
$$\Delta_{BC} = \Psi_B - \Psi_C = S_B - S_C \tag{10}$$

Fees remain costs in S_j even when they are transfers to I or V. A social-welfare calculation would require aggregating all four players' payoffs and is not claimed here.

4. Equilibrium Results

Proposition 1: Centralized deterrence feasibility. Suppose d_C is continuous and strictly increasing and G_C is nondecreasing in a_C. Define the set of monitoring levels that make centralized access a_C credible and its smallest element by

$$M_C(a_C) = \{m \in [0, m^{\max}] \mid$$
$$\mid d_C(m)F_C \geq G_C(a_C)\},$$
$$m_C^{\min}(a_C) = \min M_C(a_C) \text{ if } M_C(a_C) \neq \emptyset \tag{11}$$

The set $M_C(a_C)$ contains exactly the feasible monitoring choices that induce $e_C^* = 0$ at access a_C. If the set is nonempty, $m_C^{min}(a_C)$ is its smallest and therefore least-cost element because C_m is increasing. The set is empty exactly when $d_C(m^{max})F_C < G_C(a_C)$: even maximum feasible monitoring then fails to deter A. Under the attack-contingent choice rule, this does not mean that physical access is impossible; it means that C is absent from U's choice set at that access level because the attacker would choose positive equilibrium effort. Player I may still choose a lower a_C for which $M_C(a_C)$ is nonempty.

Proof sketch. Continuity on $[0, m^{\max}]$ makes $M_C(a_C)$ closed, so its minimum exists whenever it is nonempty. If $G_C(a_C) \leq d_C(0)F_C$, then $m_C^{\min} = 0$. If $d_C(0)F_C < G_C(a_C) < d_C(m^{\max})F_C$, the unique interior threshold solves $d_C(m)F_C = G_C(a_C)$. If $G_C(a_C) = d_C(m^{\max})F_C$, then $m_C^{\min} = m^{\max}$; above that pre-penalty attack payoff the set is empty. Wider access weakly raises the threshold when G_C is nondecreasing; a larger F_C or more productive detection lowers it. Equation (11) is A's zero-profit threshold, not I's indifference condition. Given V's strategy, I compares its maximum continuation payoff over credible pairs with the normalized zero payoff from withdrawal. It provides a credible C only if that maximum is nonnegative and U selects C.

Proposition 2: Permissionless attack deterrence. If $d_B > 0$, $\phi > 0$, and $\eta > 0$, the minimum validation stake that makes A's best blockchain attack unprofitable is

$$s_B^{\min}(K_B) = \frac{G_B(K_B)}{d_B \phi \eta} \tag{12}$$

Proof sketch. The inequality in (6) binds at the minimum validation stake that deters the attacker. The threshold rises with the maximum pre-penalty attack payoff and secured value and falls with d_B, ϕ, η, or any higher attack cost that lowers G_B. If $s_B^{min}(K_B) > s^{max}$, B cannot be secured. Equation (12) is A's zero-profit threshold, not V's indifference point: at s_B^{min}, A is indifferent and the attacker's tie-breaking convention selects zero effort. Given I's strategy, V compares its best payoff among attack-deterring stake choices under (2), (8), and (9) with zero at $s_B = 0$. If revenue is fixed once B is credible and C_s is increasing, V chooses s_B^{min} and is indifferent when $D_B r_V(s_B^{min}, f_B) - C_s(s_B^{min}) = 0$; otherwise its optimal validation stake may exceed s_B^{min}. If V chooses sB = 0, B is inactive and absent from U's choice set; hence U can choose B only when V participates, supplies positive stake, and A chooses zero attack effort.

Proposition 3: Conditional regime choice. Assume $a_C > 0$, $s_B > 0$, and $e_C^\star = e_B^\star = 0$, equivalently that both regimes are active and both deterrence inequalities on the right-hand side of (6) hold. Evaluating X_C and X_B at the induced equilibrium

attributes, the blockchain-centralized certainty-equivalent difference is

$$\Delta_{BC} = (X_B - X_C) - (f_B - f_C) - \\ -(\mu_B - \mu_C) - \frac{\rho}{2}(\sigma_B^2 - \sigma_C^2) \quad (13)$$

Holding the induced nonfee attributes X_C and X_B and the residual-loss moments fixed, the pairwise blockchain fee at which U is indifferent between B and C is

$$f^{\star}_{B,BC} = f_C + (X_B - X_C) - (\mu_B - \mu_C) - \\ -\frac{\rho}{2}(\sigma_B^2 - \sigma_C^2) \quad (14)$$

Proof sketch. By $\Delta_{BC} = \Psi_B - \Psi_C = S_B - S_C$ in (10), substitution from (7) yields (13). Holding nonfee attributes and residual-loss moments fixed, $\Delta_{BC} = 0$ gives the pairwise B–C indifference fee in (14). Under the same tie-breaking rule, U chooses B if and only if $S_B > \max\{0, S_C\}$; chooses C if and only if $S_C > 0$ and $S_C \geq S_B$; and chooses O if and only if $0 \geq \max\{S_B, S_C\}$. At pairwise B–C equality, C is selected if their common payoff exceeds the payoff of O; otherwise O. All security parameters can matter in equilibrium because they determine whether a regime enters U's choice set and may change access or assurance. Equation (13) isolates the direct B–C comparison only after the equilibrium gross user benefits X_j have been fixed; conditional on that step, the remaining direct terms are fees and residual-loss moments.

5. Symbolic Benchmark and Sensitivity

The benchmark is symbolic rather than numerically calibrated. It assumes both regimes are secure and $\Psi_B > \Psi_C > \Psi_O$, while allowing $f_C < f_B$. The blockchain advantage is measured by Δ_{BC}, not by an undefined qualitative label such as "modestly higher."

Tables 1 and 2 separate two analytical steps. First, sanctions, detection, slashing, and provider participation determine whether regimes C and B enter U's choice set and may alter gross user benefits through access or assurance. Second, conditional on both regimes being available and on fixed X_C and X_B, equation (13) compares U's risk-adjusted payoffs through fees and residual-loss moments. This conditional decomposition does not imply that the other parameters are irrelevant to equilibrium.

Table 1. Model closure and symbolic benchmark for the user's choice among C, B, and O.

Element	Formal definition or benchmark	Interpretation
Players and timing	Unitary I, V, A, U; primitives fixed before play; I, V move in period 1, A in period 2, and U in period 3.	Regimes C and B belong to the model infrastructure; D_C and D_B are generated by U's regime choice.
Strategy ranges	$x_I \in \mathcal{X}_I, s_B \in \mathcal{X}_V, x_A \in \mathcal{X}_A$, and $j \in \mathcal{X}_U$	Compact strategy sets; the period-1 outcome is a mutual best response in continuation payoffs.
Fees and quality	$f_C, f_B, q_C, q_B \geq 0$ are exogenous and observed before U chooses.	No baseline pricing or quality-setting game is assumed; fees are user costs.
Central security	$G_C(a_C) \leq d_C(m_C)F_C$	A chooses $e^{\star}_C = 0$; monitoring is costly to I and can increase assurance for U.
Blockchain security	$G_B(K_B) \leq d_B \phi \eta s_B$	A chooses $e_B^* = 0$; the expected slashing loss is borne by A on the attack-controlled position ηs_B. Player V remains a distinct supplier of validation stake.
User objective	User's choice set J_U and certainty equivalent Ψ_j in (7)–(8).	Regime O is always available. Regime C requires $a_C > 0$ and $e_C^{\star} = 0$; regime B requires $s_B > 0$ and $e_B^{\star} = 0$. Holding expected net payoff fixed, the variance penalty is $(\rho/2)\sigma_j^2$.
Net user surplus	$S_j = \Psi_j - \Psi_O$ and $\Delta_{BC} = S_B - S_C$	This is user surplus, not social welfare.
Symbolic benchmark	Both regimes active and attack-deterred, with $\Psi_B > \Psi_C > \Psi_O$.	Regime B is chosen because Ψ_B is strictly highest. Access, assurance, quality, residual net loss, and fees are combined in Ψ_j. Holding nonfee attributes fixed, the pairwise ranking reverses when $f_B > f_{B,BC}^{\star}$ or another regime has the higher certainty equivalent.

6. Extension and Limitations

6.1. Permissioned Hybrid as an Extension

The baseline choice set is limited to C, B, and O. A permissioned hybrid ledger H with restricted validators, audit evidence, legal liability, or collateral is not solved and is therefore not reported as a benchmark result. A formal extension must define validator admission, audit and liability parameters, the attacker's rent and payoff, an attack-deterrence condition for H, and a certainty-equivalent payoff Ψ_H. Only then can the choice set be expanded to $\{C, B, H, O\}$.

6.2. Limitations

The four players remain distinct and unitary. Player V aggregates validation supply and chooses only s_B; it does not choose misconduct. Player A is a separate pre-existing malicious coalition and alone chooses e_C and e_B. Player V never becomes A. A model of individual validators choosing validation, abstention, or misconduct would require additional players and strategies and would assign slashing to those choosing misconduct. Attacks are separable and A has no common effort budget.

The attack-contingent choice rule removes a targeted regime from U's choice set when equilibrium attack effort is positive, but does not make it physically unavailable. A continuous-loss extension could keep it selectable and subtract expected strategic loss. Fees and quality are exogenous; residual net loss is post-deterrence operational and governance loss under CARA-normal preferences. The model has one representative user and omits token-price dynamics, validator entry, repeated governance, learning, and calibration. A pure period-1 Nash fixed point is assumed; discontinuous user choice may require mixed strategies or smoothed demand. S_j is user surplus, not social welfare.

Table 2. Sensitivity of the blockchain-centralized margin and the security constraints.

Parameter change	Direct effect on Δ_{BC}	Security or regime-availability mechanism
Blockchain fee f_B increases	$\partial\Delta_{BC}/\partial f_B = -1$	Reduces U's B–C payoff difference one-for-one and favors C or O.
Centralized fee f_C increases	$\partial\Delta_{BC}/\partial f_C = +1$	Raises U's B–C payoff difference one-for-one and favors B.
Blockchain loss moments increase	$\partial\Delta_{BC}/\partial\mu_B = -1$, $\partial\Delta_{BC}/\partial\sigma_B^2 = -\rho/2$	Lowers the blockchain payoff; the variance effect is stronger for larger risk aversion.
Central loss moments increase	$\partial\Delta_{BC}/\partial\mu_C = +1$, $\partial\Delta_{BC}/\partial\sigma_C^2 = +\rho/2$	Raises U's B–C payoff difference and favors B.
Blockchain gross value X_B increases	$\partial\Delta_{BC}/\partial X_B = +1$	Better blockchain access, quality, or assurance favors B, provided B remains available.
Central gross value X_C increases	$\partial\Delta_{BC}/\partial X_C = -1$	Better centralized access, quality, or assurance favors C.
F_C or d_C improves	No direct effect conditional on fixed X_C, X_B, and both regimes available.	Lowers minimum monitoring or permits broader access in C; induced changes in a_C or m_C can alter X_C.
d_B, ϕ, or η improves	No direct effect conditional on fixed X_C, X_B, and both regimes available.	Lowers minimum stake or expands the parameter region in which B is available; induced changes in s_B can alter X_B.

7. Conclusion

The model compares centralized regime C and permissionless proof-of-stake regime B in a three-period game. Players I and V move simultaneously in period 1, A moves in period 2, and U chooses among C, B, and O in period 3.

Regime C is attack-deterred when A's best pre-penalty payoff does not exceed its expected legal or contractual penalty. Regime B is attack-deterred when that payoff does not exceed expected slashing on the attack-controlled stake. Player U then compares available regimes through Ψ_j, net of fees and residual net loss; holding expected payoff fixed, greater loss variance lowers this certainty equivalent.

The result is conditional: B is chosen only when its risk-adjusted certainty equivalent is strictly highest. Stronger centralized sanctions may support broader access in C; higher blockchain fees or risk may move U to C or O. No single player designs both regimes. The model instead identifies when each separately supplied regime is operational and attack-deterred and predicts which regime U selects.

Acknowledgements

No funding was received for this research. The authors declare no competing interests.

References

[1]. S. Nakamoto, Bitcoin: A peer-to-peer electronic cash system, 2008, https://bitcoin.org/bitcoin.pdf

[2]. R. Böhme, N. Christin, B. Edelman, T. Moore, Bitcoin: Economics, technology, and governance, *Journal of Economic Perspectives*, Vol. 29, Issue 2, 2015, pp. 213-238.

[3]. F. Saleh, Blockchain without waste: Proof-of-stake, *The Review of Financial Studies*, Vol. 34, Issue 3, 2021, pp. 1156-1190.

[4]. P. De Filippi, M. Mannan, W. Reijers, Blockchain as a confidence machine: The problem of trust and challenges of governance, *Technology in Society*, Vol. 62, 2020, 101284.

[5]. S. Davidson, P. De Filippi, J. Potts, Blockchains and the economic institutions of capitalism, *Journal of*

Institutional Economics, Vol. 14, Issue 4, 2018, pp. 639-658.

[6]. F. Schär, Decentralized finance: On blockchain- and smart contract-based financial markets, *Federal Reserve Bank of St. Louis Review*, Vol. 103, Issue 2, 2021, pp. 153-174.

[7]. K. Hausken, G. Wang, A game-theoretic model of cryptocurrencies, fiat currencies and assets, *International Review of Financial Analysis*, Vol. 117, 2026, 105226.

(006)

Return and Volatility Spillovers across Blockchain Layers

C. Jeanneaux
UCLy (Lyon Catholic University), ESDES, Lyon, France
UCLy (Lyon Catholic University), UR CONFLUENCE: Sciences et Humanités (EA1598), Lyon, France
E-mail: cjeanneaux@univ-catholyon.fr

Summary: This paper introduces a layer-based analytical framework for measuring return and volatility spillovers across blockchain infrastructure. Using the Diebold and Yilmaz's (2012) [1] connectedness methodology, we construct equal-weighted daily indices for three layers, Layer 0 (interoperability protocols), Layer 1 (settlement blockchains), and Layer 2 (scaling solutions), from 18 crypto-assets over 2 217 daily observations spanning April 2019 to April 2025. The total spillover index reaches 45.0 % for returns and 46.2 % for volatility, indicating substantial heterogeneity in cross-layer exposure across L0, L1, and L2. This paper's contribution is to apply directional return and volatility connectedness measures to aggregate L0, L1, and L2 indices, documenting a potentially time-varying L0/L1/L2 transmission hierarchy. Over the full sample, Layer 0 emerges as the dominant net transmitter of volatility shocks and Layer 1 the primary net receiver, a pattern that challenges the conventional view of Layer 1 as the ecosystem's risk epicenter. Rolling-window analysis reveals sharp connectedness spikes during the Terra Luna collapse (May 2022) and the FTX bankruptcy (November 2022). The exclusion of Bitcoin and Ethereum from the Layer 1 index confirms the full-sample hierarchy is not a large-cap artefact. However, sub-period analysis reveals that this hierarchy is itself time-varying: over 2023-2025, L1 shifts from net receiver to net transmitter while L0's transmitter position weakens substantially, consistent with the maturation of the L1 ecosystem following Bitcoin ETF adoption.

Keywords: Blockchain layers, Cryptocurrencies, Spillovers, Contagion, Bitcoin, Layer 1, Layer 2.

1. Introduction

The cryptocurrency ecosystem is organized into technologically distinct architectural strata: Layer 0 (L0) interoperability protocols such as Polkadot and Cosmos provide cross-chain infrastructure; Layer 1 (L1) blockchains such as Bitcoin and Ethereum serve as consensus and settlement layers; and Layer 2 (L2) scaling solutions such as Arbitrum and Optimism are built atop L1 to enhance throughput. Despite this structural heterogeneity, the existing spillover literature treats crypto-assets as a homogeneous class, entirely ignoring whether financial shocks propagate along technological dependency chains or reflect common macro-financial exposures. This paper fills that gap.

We ask: do return and volatility spillovers across blockchain layers follow the technological hierarchy of the stack, consistent with Allen and Gale's (2000) [2] network structure-based transmission channel, or do they instead reflect fundamental-based synchronicity driven by shared investor bases and common macro-financial exposures? The theoretical prior, grounded in Acemoglu et al. (2015) [3], is that the structural position of a node determines its role as transmitter or absorber. This prior generates a testable hypothesis for blockchain layers: if L0 protocols provide the foundational infrastructure for cross-chain connectivity, they should occupy a distinct position in the transmission hierarchy relative to L1 settlement layers and L2 scaling solutions. Whether this position is that of a dominant transmitter, and whether it remains stable as the ecosystem matures, are empirical questions we address.

The literature on cryptocurrency spillovers has documented several stylized facts relevant to our study. Ji et al. [4] confirm Bitcoin's centrality using a network-based extension of the Diebold and Yilmaz (2012) [1] framework, while Yi et al. (2018) [5] document the secular rise in volatility connectedness and question whether Bitcoin's dominance is eroding as the ecosystem matures. Antonakakis et al. (2019) [6] show that connectedness is highly state-dependent, with periods of high uncertainty corresponding to sharply elevated spillovers, a pattern further documented by Conlon et al. (2023) [7] around the FTX collapse, where correlations between crypto-assets surged in the weeks following the exchange's bankruptcy. Despite this rich body of evidence, all existing cryptocurrency spillover studies share a common limitation: they treat crypto-assets as a homogeneous class or organize the analysis by market capitalization sub-categories, ignoring the technological stratification that structures the blockchain ecosystem. A structural channel for such stratification is not merely theoretical: Adams (2024) [8] provides direct empirical evidence that transaction costs on L1 are approximately 35 times higher than on L2, and that this cost differential shapes user behaviour across layers, with trading volume migrating back to L1 precisely during periods of market volatility. A recent paper from Harvey et al. (2026) [9] models the economic interactions between L1 and L2. Our paper builds on this categorization and proposes an analysis of volatility and return spillover across layers.

2. Data and Methodology

We construct equal-weighted daily indices for three blockchain layers from 18 crypto-assets (the number of available assets varies and details can be found in table 1) over 2217 daily observations spanning April 2019 to April 2025, the earliest window over which a three-layer analysis is economically meaningful, given the embryonic nature of the L2 ecosystem prior to 2021. Layer indices are computed as equal-weighted averages over the subset of assets with valid observations on each trading day, so the number of constituents varies as tokens are progressively listed (Table 1). Data are sourced from CoinGecko historical archives and layer classification follows CoinGecko's taxonomy. L1 includes the ten largest settlement blockchains by market capitalization as of March 2026. L2 includes all five scaling protocols with a native token and sufficient price history over our sample period, a binding constraint given that most L2 tokens were launched after 2021. L0 is necessarily limited to three assets as only Polkadot, Cosmos and Internet Computer unambiguously qualify as pure interoperability protocols under major taxonomies over our sample period. We acknowledge this as a limitation and conduct a robustness test excluding Bitcoin and Ethereum to verify that the main results are not driven by the dominance of these two assets within the L1 index.

Table 1. Number of assets included in the categories.

	2019	2020	2021	2022	2023	2024	2025
L0	1	2	3	3	3	3	3
L1	9	10	10	10	10	10	10
L2	2	2	3	5	5	5	5

Note: Crypto assets and their year of data availability are shown in parentheses below the table. L0 → ATOM (2019), DOT (2020), ICP (2021); L1 → BTC (2019), ETH (2019), XRP (2019), BNB (2019), SOL (2019), ADA (2019), TRX (2019), BCH (2019), LTC (2019), AVAX (2020); L2 → POL (2019), LRC (2019), ARB (2021), OP (2022), IMX (2022). While all 18 assets remained continuously listed and liquid throughout the sample period, the selection of layer constituents reflects the set of projects that achieved and maintained prominence within their category by 2025, which may not fully represent the broader, more volatile population of projects that existed at earlier stages of each layer's development.

We apply the Diebold and Yilmaz (2012) [1] connectedness framework, building on Diebold and Yilmaz (2009) [10], Koop et al. (1996) [11], and Pesaran and Shin (1998) [12], based on a three-variable VAR and the generalized forecast error variance decomposition (GFEVD), which is invariant to variable ordering, a critical property given the absence of an established causal hierarchy among blockchain layers. We compute total spillover indices, directional FROM/TO measures, and net positions identifying net transmitters and receivers. Robustness is assessed through: (i) a 200-day rolling window analysis; (ii) exclusion of Bitcoin and Ethereum from the L1 index; and (iii) a sub-period analysis over 2023-2025 when all 18 tokens are simultaneously available. We adopt a 200-day rolling window, consistent with the original specification of Diebold and Yilmaz (2012) [1]. The literature seems to have no clear consensus around the length of the window selected to run a rolling-window analysis in cryptocurrency markets. This window length, corresponding approximately to one trading year in traditional finance, balances estimation precision against sensitivity to regime shifts, allowing spillover measures to adapt to structural breaks without being dominated by distant historical observations.

3. Main Results

Cross-layer financial integration is economically large: the total spillover index (Table 2) reaches 45.0 % for returns and 46.2 % for volatility (Table 3), substantially above the 26.6 % and 20.8 % reported by Koutmos (2018) [13] for individual cryptocurrencies in 2015-2018, consistent with the secular deepening of crypto-market interconnectedness. These results indicate substantial heterogeneity in cross-layer exposure across L0, L1 and L2.

Table 2. Spillover table for layer returns.

	L0	**L1**	**L2**	**FROM others (%)**
L0 (Layer 0)	55.54	32.11	12.35	44.46
L1 (Layer 1)	32.61	45.19	22.20	54.81
L2 (Layer 2)	14.98	20.76	64.26	35.74
TO others (%)	**47.59**	**52.87**	**34.55**	-
NET (%)	***+3.13***	***-1.94***	***-1.19***	-
Spillover index				***45.01 %***

Table 3. Spillover table for layer volatility.

	L0	**L1**	**L2**	**FROM others (%)**
L0 (Layer 0)	61.17	27.94	10.90	38.83
L1 (Layer 1)	29.49	41.55	28.96	58.45
L2 (Layer 2)	19.20	22.16	58.64	41.36
TO others (%)	**48.69**	**50.09**	**39.86**	-
NET (%)	***+9.86***	***-8.35***	***-1.50***	-
Spillover index				***46.21 %***

The central finding concerns the directional asymmetry of volatility transmission. Layer 0 is the dominant net transmitter of volatility shocks (+9.9 %), while Layer 1 is the primary net receiver (-8.4 %), with Layer 2 near-neutral (-1.5 %). This finding survives the exclusion of Bitcoin and Ethereum from the L1 index, with L0's net transmitter position strengthening to

+11.1 %, confirming it is not a large-cap artefact. In the return system, by contrast, net positions are near-symmetric across layers (maximum spread: 5.1 percentage points), inconsistent with directional technological transmission and instead supporting fundamental-based synchronicity.

Rolling-window analysis reveals sharp connectedness spikes during the Terra Luna collapse (May 2022) and the FTX bankruptcy (November 2022), with the return spillover index approaching 65 % and the volatility index 70 %, indicating exceptionally elevated connectedness across layers during systemic stress, consistent with the rational cascading mechanism and eliminating layer diversification benefits precisely when they are most needed. Notably, the sub-period analysis over 2023-2025 reveals an evolution in the volatility hierarchy: L1 shifts from net receiver to net transmitter (+7.2 %), while L0's transmitter position weakens substantially (+1.5 %), consistent with the maturation of the L1 ecosystem following Bitcoin ETF adoption and deepening institutional participation.

4. Robustness Tests

We now examine the temporal dynamics and robustness of cross-layer connectedness. First, we apply a rolling-window analysis to assess how the transmission hierarchy evolves over time and around major market events (Fig. 1 and 2). We then verify that our results are not driven by Bitcoin and Ethereum's dominance within the Layer 1 index by re-estimating the model on their exclusion.



Fig. 1. Total Return Spillover Index This figure plots the total return spillover index computed over a 200-day rolling window using the generalized forecast error variance decomposition (GFEVD) at a 10-day horizon (H = 10). Sample: 6 April 2019 - 30 April 2025.

Fig. 1 and 2 plot the rolling total spillover index for both returns and volatility across the full sample. Two findings emerge. First, both indices are manifestly time-varying, oscillating between approximately 25 % and 70 %. These results confirm time-variation of spillover indexes which is consistent with the regime-dependent connectedness documented by Antonakakis et al. (2019) for individual cryptocurrencies. Second, sharp spikes in both indices coincide with the Terra Luna collapse (May 2022) and the FTX bankruptcy (November 2022).

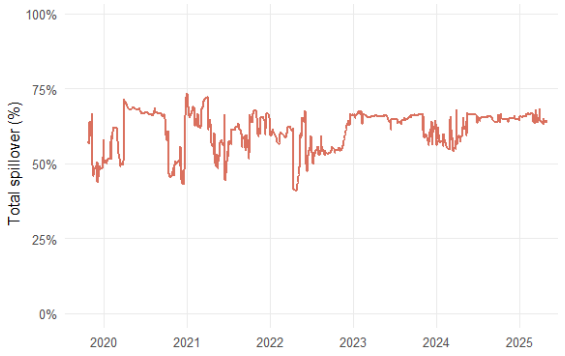


Fig. 2. Total Volatility Spillover Index. This figure plots the total volatility spillover index over the same rolling window and horizon. Volatility is measured as the 30-day rolling standard deviation of log-returns for each layer index, annualized by √365.

4.1. Rolling-Window Analysis

To examine whether the full-sample spillover hierarchy holds under stress, we extract the rolling net positions at three major crisis dates (Table 4). During the Terra Luna collapse (10 May 2022), triggered by the algorithmic de-pegging of the UST stablecoin and the subsequent death spiral of the LUNA token, which wiped approximately $40 billion in market capitalization within days; L1 becomes a massive net receiver (-49.1 %) while L2 switches to net transmitter (+40.7 %). During the FTX bankruptcy (6 November 2022), the collapse of the second-largest centralized cryptocurrency exchange, which triggered a cascade of counterparty failures across the industry; net positions are near-neutral across all three layers despite a total spillover index of 54.4 %, suggesting that this exchange-level shock propagated simultaneously across all layers rather than directionally. During the Silicon Valley Bank failure (10 March 2023), the largest US bank collapse since 2008, which generated contagion fears across the broader financial system and triggered sharp selloffs in crypto markets, the total spillover index reaches its highest level of the three episodes (65.5 %), yet directional positions remain weak, indicating that this macro-financial shock triggered parallel liquidations across all crypto layers rather than structural cross-layer transmission. Taken together, these crisis-period results suggest that the dominant mechanism during stress episodes is not technological shock propagation but rather common-factor exposure, all layers reacting simultaneously to ecosystem-wide information shocks, regardless of their architectural position in the blockchain stack.

4.2. Exclusion of BTC and ETH

Our principal robustness concern is the dominance of Bitcoin and Ethereum within the L1 index. BTC and

ETH jointly represent over 70 % of L1 market capitalization throughout the sample, raising the possibility that the L1 index's net receiver position is driven by idiosyncratic BTC/ETH dynamics rather than by broad L1 layer characteristics. To address this, we re-estimate the VAR excluding BTC and ETH from the L1 index (using the remaining eight L1 assets: XRP, BNB, SOL, ADA, TRX, BCH, LTC, AVAX). Tables 5 and 6 report the results and seem to confirm previous conclusions.

Table 4. Net Volatility Spillover Positions by Layer During Crisis Episodes (Rolling Window, 200 days).

Episode	Date	NET L0 (%)	NET L1 (%)	NET L2 (%)	Total Spillover (%)
Full sample	Apr 2019 Apr 2025	9.86	-8.35	-1.50	46.21
Terra Luna	May 2022	8.38	-49.06	40.68	49.09
FTX bankruptcy	Nov 2022	-4.66	0.10	4.56	54.36
SVB failure	March 2023	-0.30	-3.95	4.25	65.47

NET values extracted from the 200-day rolling generalized FEVD at the date of each event. Positive NET = net transmitter; negative NET = net receiver. Ukraine war (24 Feb 2022) excluded due to insufficient L2 data within the rolling window at that date. SVB Failure stands for the Silicon Valley Bank Failure.

Table 5. Return spillover index exclusion BTC and ETH from L1.

	L0	**L1 (ex-BTC/ETH)**	**L2**	**FROM others (%)**
L0	55.93	32.38	11.68	44.07
L1 (ex-BTC/ETH)	32.11	46.64	21.25	53.36
L2	15.29	19.63	65.09	34.91
TO others (%)	**47.40**	**52.01**	**32.93**	—
NET (%)	***3.33***	***-1.35***	***-1.98***	—
Spillover index				***44.11 %***

Table 6. Volatility spillover index exclusion BTC and ETH from L1.

	L0	**L1 (ex-BTC/ETH)**	**L2**	**FROM others (%)**
L0	61.60	27.66	10.74	38.40
L1 (ex-BTC/ETH)	29.85	42.04	28.11	57.96
L2	19.62	22.18	58.20	41.80
TO others (%)	**49.48**	**49.84**	**38.85**	—
NET (%)	***11.08***	***−8.13***	***−2.95***	—
Spillover index				***46.06 %***

5. Conclusions

This paper makes three contributions. First, it introduces a layer-based analytical framework for cryptocurrency spillover analysis, providing a theoretically grounded taxonomy that maps financial interdependencies onto blockchain architectural structure. Second, it shows that returns and volatility do not behave symmetrically across layers: while return spillovers are near-symmetric and consistent with fundamentals-based synchronicity, volatility spillovers exhibit a directional pattern originating from Layer 0 over the full sample, a pattern consistent with a technological transmission channel but which our data cannot causally establish. Third, and most importantly, we document that this hierarchy is not a stable structural property but an evolving feature of the ecosystem: the L0 net-transmitter position weakens substantially over 2023-2025 as L1 shifts from net receiver to net transmitter, consistent with the maturation of the L1 ecosystem following Bitcoin ETF adoption and deepening institutional participation.

The practical implications are threefold. For investors, the high cross-layer connectedness during crises implies that layer diversification provides limited systemic risk protection. For risk managers, the finding that volatility spillovers substantially exceed return spillovers during stress implies that Value-at-Risk models calibrated on unconditional correlations will systematically underestimate tail risk. For policymakers, the finding that L0 bridging infrastructure has constituted a disproportionate source of ecosystem-wide volatility uncertainty over most of our sample, despite its lower market capitalization, suggests that systemic risk monitoring should extend beyond settlement-layer protocols; the recent shift of L1 toward a transmitter position further indicates that monitoring frameworks should not assume static layer roles.

6. References

[1]. F. X. Diebold, K. Yilmaz, Better to give than to receive: Predictive directional measurement of volatility spillovers, *International Journal of Forecasting*, Vol. 28, Issue 1, 2012, pp. 57-66.

[2]. F. Allen, D. Gale, Financial contagion, *The Journal of Political Economy*, Vol. 108, Issue 1, 2000, pp. 1-33.

[3]. D. Acemoglu, A. Ozdaglar, A. Tahbaz-Salehi, Systemic risk and stability in financial networks, *American Economic Review*, Vol. 105, Issue 2, 2015, pp. 564-608.

[4]. Q. Ji, E. Bouri, C. K. M. Lau, D. Roubaud, Dynamic connectedness and integration in cryptocurrency markets, *International Review of Financial Analysis*, Vol. 63, 2019, pp. 257-272.

[5]. S. Yi, Z. Xu, G.-J. Wang, Volatility connectedness in the cryptocurrency market: Is Bitcoin a dominant cryptocurrency?, *International Review of Financial Analysis*, Vol. 60, 2018, pp. 98-114.

[6]. N. Antonakakis, I. Chatziantoniou, D. Gabauer, Cryptocurrency market contagion: Market uncertainty,

market complexity, and dynamic portfolios, *Journal of International Financial Markets, Institutions and Money*, Vol. 61, 2019, pp. 37-51.

[7]. T. Conlon, S. Corbet, Y. Hu, Understanding the contagion effects of the FTX exchange collapse, *SSRN*, 2023, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4640797.

[8]. A. Adams, Layer 2 be or Layer not 2 be: Scaling on Uniswap v3, *arXiv*, 2024, arXiv:2403.09494.

[9]. C. R. Harvey, F. Saleh, R. Sverchkov, An economic model of the L1-L2 interaction, *SSRN*, 2026, https://ssrn.com/abstract = 5163823

[10]. F. X. Diebold, K. Yilmaz, Measuring financial asset return and volatility spillovers, with application to global equity markets, *The Economic Journal*, Vol. 119, Issue 534, 2009, pp. 158-171.

[11]. G. Koop, M. H. Pesaran, S. M. Potter, Impulse response analysis in nonlinear multivariate models, *Journal of Econometrics*, Vol. 74, Issue 1, 1996, pp. 119-147.

[12]. H. H. Pesaran, Y. Shin, Generalized impulse response analysis in linear multivariate models, *Economics Letters*, Vol. 58, Issue 1, 1998, pp. 17-29.

[13]. D. Koutmos, Return and volatility spillovers among cryptocurrencies, *Economics Letters*, Vol. 173, 2018, pp. 122-127.

(007)

A Framework based on ETL and Blockchain for Advanced Cyber Threat Intelligence Sharing

A. Fotia [1], A. Salemi [1], A. Canino [2] and G. Lax [1]
[1] DIIES Department, University of Reggio Calabria, Italy
[2] University College London (UCL), United Kingdom
Tel.: ++39 0965 1693304
E-mail: lax@unirc.it

Summary: The rise of the Ransomware-as-a-Service model has escalated the frequency and sophistication of cyber-extortion campaigns, with threat actors frequently leveraging dark web leak sites on the Tor network for double-extortion tactics. Effectively tracking these threats requires robust Cyber Threat Intelligence (CTI); however, current CTI sharing mechanisms rely on fragmented, centralized platforms that suffer from single points of failure, lack of cross-organizational trust, and severe blind spots due to the volatility of anonymous networks. To overcome these limitations, this paper proposes a comprehensive and decentralized framework for the automated aggregation, normalization, and sharing of ransomware threat intelligence. The first pillar of the architecture features a resilient Extract, Transform, and Load (ETL) pipeline that employs scraping engines to reliably navigate the Tor network, along with local storage for high-speed data extraction, deduplication, and real-time alerting. The second pillar democratizes access to this refined intelligence by anchoring the standardized data to a permissioned consortium blockchain utilizing Hyperledger Fabric. By eliminating centralized intermediaries and transaction costs of public ledgers, the proposed system provides a scalable, trusted environment for organizations to maintain real-time situational awareness and collaboratively mitigate modern ransomware operations.

Keywords: Ransomware-as-a-service, RaaS, Hyperledger fabric, Tor, Dark web.

1. Introduction

The rapid acceleration of digital transformation has significantly expanded the attack surface for organizations around the world, leading to a proliferation of sophisticated cyber threats. In the broader cybersecurity landscape, one of the main threats is *ransomware*. This term refers not only to a specific type of malware, but also to a set of cyber-extortion techniques that involve blocking access to a victim's data or systems accompanied by a demand for payment to restore them or to avoid further harmful consequences.

Driven by the Ransomware-as-a-Service (RaaS) model, modern threat actors have shifted from opportunistic infections to highly targeted, multi-stage campaigns. These campaigns frequently employ double extortion tactics, combining data exfiltration with data encryption, and the threat of publication on dark web leak sites. Monitoring and mitigating such a diverse array of threats requires actionable Cyber Threat Intelligence (CTI).

Recent advances in cybersecurity have led to the development of various techniques to detect ransomware before or during the encryption phase [1]. These approaches can be broadly classified into Machine-Learning-based and Non-Machine-Learning-based methods, each leveraging different behavioral, structural, or network-level features of the malware.

Machine Learning (ML) models learn from historical data to identify malicious patterns, such as API calls, system calls, and I/O request behaviors. These models are highly effective at detecting both known and zero-day variants [2].

Among non-Machine-Learning-based methods, decoy-based approaches involve deploying *honey files* or trap systems that mimic valuable user data or network resources. Detection is triggered immediately when unauthorized processes attempt to read, write, or modify these decoys. A significant contribution in this area is made by Mehnaz et al. [3], who developed *RWGuard*, a honeyfile-based strategy that detects and blocks suspected ransomware in near-real-time. Software-Defined Networking methods passively monitor network traffic to identify malicious activity, such as communication with Command and Control servers or anomalous file-sharing patterns. Akbanov et al. [4] developed a solution based on OpenFlow that detects suspicious activities through network traffic monitoring and blocks infected hosts by adding flow table entries to OpenFlow switches in real-time.

Beyond the intrinsic difficulty in ransomware detection, another critical aspect is that the current CTI sharing mechanisms are highly fragmented. Organizations typically rely on isolated centralized Threat Intelligence Platforms, which suffer from single point of failure, lack of trust among competing entities, and severe blind spots due to the volatility of the infrastructures used by cybercriminals (often based on anonymity services such as the Tor network [5]).

To overcome the limitations of isolated monitoring and trust asymmetry, we propose a comprehensive framework that bridges data collection techniques with a decentralized sharing infrastructure. The first pillar of our solution is a robust Extract, Transform, and Load (ETL) pipeline designed to autonomously retrieve intelligence across heterogeneous and volatile environments. Our collection engine routes traffic

through Tor SOCKS5 proxies to securely access dark web leak sites and employs headless browser automation to render dynamic, JavaScript-heavy content. This allows the system to continuously monitor the public activity of RaaS affiliates and *Initial Access Brokers*, entities that acquire and resell compromised access obtained through stolen credentials, corporate VPNs, or exposed vulnerabilities. This procedure enables the interception of new victims, stolen data proofs, and impending leak countdowns. The system inherently manages the unreliability of the Tor network through adaptive retry mechanisms and intelligent request scheduling, ensuring continuous observation of the threat landscape.

Raw data collected from these disparate sources is highly unstructured and requires meticulous modeling to be useful for downstream consumers, such as Security Operations Centers and Computer Emergency Response Teams. Our framework implements a dedicated Normalization Engine based on a registry pattern. This engine applies group-specific parsing logic to extract and standardize key entities from the raw data. The intelligence is modeled into a highly structured format that explicitly defines the ransomware group, victim name, industrial sector, geographic location (inferred heuristically via Top-Level Domains), publication timestamps, and cryptographic hashes of the leaked payloads. By structuring the CTI in this standardized manner, we ensure that any participating organization can easily search, filter, and correlate the exact intelligence they need (for example, querying all recent double-extortion attacks). Deduplication is strictly enforced at this stage using cryptographic hashing to prevent redundant alerts and database bloat.

To democratize access to this refined intelligence while ensuring data integrity, provenance, and accountability, the second pillar of our architecture relies on a permissioned consortium blockchain. We propose the use of Hyperledger Fabric, which provides a robust alternative to public blockchains like Ethereum. In a public ledger, sensitive intelligence might be exposed to the threat actors themselves, and the financial burden of gas fees for high-volume data anchoring is unsustainable. A consortium blockchain, instead, restricts participation to authenticated, trusted entities (e.g., allied corporations, governmental agencies, and cybersecurity vendors). This permissioned nature eliminates transaction costs, allows for high-throughput data sharing, and ensures compliance with data protection regulations by controlling who can view specific intelligence channels. The orchestration of the shared intelligence on the consortium blockchain is managed through dedicated smart contracts. When the ETL pipeline identifies and normalizes a new threat event, the system invokes the smart contract to anchor the structured data to the immutable ledger.

By integrating automated dark web scraping, rigorous data modeling, and consortium blockchain technology, this framework provides a resilient and scalable environment for threat intelligence sharing. Organizations can collaboratively maintain real-time situational awareness against both traditional ransomware operations and emerging Web3 threats, without relying on a centralized intermediary.

The remainder of this paper is organized as follows. In the next section, we review related work. Section 3 defines the threat scenario and the problem addressed. Section 4 details the proposed framework, which includes the ETL pipeline and the blockchain integration. Finally, Section 5 concludes the paper and outlines future developments.

2. Literature Review

Blockchain technology has been proposed as a foundation for trustworthy CTI exchange. Homan et al. [6] design a permissioned network on Hyperledger Fabric in which channels confine what each subset of participants can see. Their prototype exchanges STIX 2.0 objects and is validated functionally, showing that non-members cannot query restricted blocks. Riesco et al. [7] and Menges et al. [8] address the reluctance to share by turning the CTI exchange into a decentralized marketplace where contributors are rewarded with tokens, on Ethereum and EOS respectively. Both rely on permissionless ledgers, so participation is open and transaction metadata are publicly visible, and the Ethereum-based design additionally pays gas for every write. Gong and Lee [9] focus on the reliability of contributors and shared content with a framework resistant to Sybil attacks. Badsha et al. [10] add fine-grained access control to the information exchanged. All these proposals start from intelligence that the members already hold, and concentrate on how it is exchanged, protected, or rewarded. None of them addresses where that intelligence comes from. Collecting it from volatile adversary-controlled sources, such as leak sites, is left to each participant, and a report submitted to the platform says nothing about how it was obtained. Our framework closes this gap by joining the two sides. The pipeline gathers leak-site publications automatically over Tor, normalizes and deduplicates them, and anchors every structured event on the consortium ledger. What members share is therefore not a report to be taken on trust, but an observation whose integrity and provenance any member can verify.

3. Scenario and Problem Formulation

The contemporary ransomware landscape has evolved into a highly specialized, decentralized criminal supply chain. Unlike early monolithic threat actors, the modern Ransomware-as-a-Service ecosystem operates on a modular, platform-like business model [11]. In this distributed architecture, core developers maintain malicious code, encryption tools, Command and Control servers, and leak sites, while independent affiliates and Initial Access Brokers

execute the actual network compromises. This operational fragmentation significantly reduces the barrier to entry for cybercriminals and increases the overall resilience of the ecosystem, making attribution and mitigation highly complex.

Within this framework, dark web leak sites serve a critical dual purpose: they act as direct instruments of psychological and reputational coercion against victims, and they operate as propaganda hubs to recruit new affiliates [12]. Tracking the activities on these platforms is essential for anticipating campaigns and prioritizing defenses. However, continuous observation presents severe structural challenges. Threat actors deliberately rotate their. onion infrastructures, generating a highly volatile environment characterized by sudden outages, mirror migrations, and intentional data manipulation, such as the publication of fake or incomplete leaks.

Consequently, the fundamental problem addressed in this paper is the pronounced operational asymmetry between offensive and defensive cybersecurity strategies. While cybercriminals collaborate fluidly through decentralized, specialized service models, organizational defense remains largely compartmentalized. Security Operations Centers and Computer Emergency Response Teams struggle to maintain a unified situational awareness because the raw intelligence is scattered across unstable sources, highly unstructured, and ultimately siloed within proprietary, centralized threat intelligence platforms.

This fragmentation leads to critical blind spots and high latency in incident response. Overcoming this asymmetry requires a paradigm shift: defenders must transition from manual, reactive monitoring in isolated silos to a proactive, automated, and trustless intelligence-sharing model capable of matching the agility of the RaaS ecosystem.

The solution we propose to address this challenge is detailed in the next section.

4. Proposed Solution

The core objective of the proposed architecture is to overcome the fragmentation of current Cyber Threat Intelligence platforms by providing a resilient environment for the automated tracking, normalization, and decentralized sharing of ransomware campaigns. The designed system is composed of two foundational pillars:

1. An automated, resilient ETL (Extract, Transform, Load) pipeline engineered to continuously monitor dark web leak sites via the Tor network, normalizing highly unstructured data into standardized threat events;
2. A decentralized, permissioned blockchain infrastructure (Hyperledger Fabric) that serves as an immutable, shared ledger for the CTI, eliminating centralized intermediaries and enabling data sharing among consortium members.

The proposed framework relies on a layered architecture that integrates data acquisition, processing, and decentralized ledger storage. The pipeline follows a structured ETL methodology followed by the persistence:

1. *Data Collection (Extract):* Access to *Tor* and *Clearnet* leak sites, handling network latency, CAPTCHAs, and dynamic JavaScript rendering;
2. *Data Processing (Transform)*: HTML parsing, data normalization, deduplication, and semantic structuring of the threat intelligence;
3. *Data Storage (Load):* Local high-speed persistence using suitable structures to support immediate dashboard queries and alert generation;
4. *Blockchain Anchoring*: Orchestration via smart contracts to record the normalized payload of the intelligence on a consortium blockchain.

The overall system architecture is depicted in Fig. 1 and detailed in the following sections.

4.1. Data Collection

Ransomware threat actors typically publish their claims, victims' data proofs, and extortion countdowns on dedicated leak sites. These platforms are primarily hosted on the Tor network to guarantee the operators' anonymity, although they are sometimes replicated on the Clearnet (i.e., the publicly accessible internet) to maximize their extortionate reach and public exposure. Because Tor is an anonymizing network that routes traffic through a series of intermediate nodes to obscure the physical location and identity of the hosting servers, continuous monitoring of these environments is highly challenging due to the intrinsic unreliability and structural volatility of these networks. Specifically, *.onion* infrastructures are characterized by high latency, frequent timeouts, and rapid rotation, as sites are often temporarily unreachable, deliberately taken offline, or migrated to different mirror addresses to evade law enforcement.

To overcome the unreliability of the Tor network, the extraction module is divided into two distinct scraper engines: the *Static Scraper* and the *Dynamic API Scraper*. From an implementation perspective, for leak sites that serve static HTML, the system utilizes the Python `requests` library configured to route traffic through a local Tor SOCKS5 proxy. The `socks5h` protocol ensures that DNS resolution is delegated to the Tor node, preventing DNS leaks that could expose the origin of the intelligence gathering operation. However, an increasing number of ransomware leak sites generate their content dynamically via complex JavaScript execution. To handle these advanced anti-scraping measures, we implemented a Dynamic Scraper utilizing `Playwright`. This scraper launches a headless Firefox browser, routes it through the Tor proxy, and waits for network stability (`networkidle`) before

extracting the fully rendered Document Object Model (DOM).

4.2. Data Processing

Due to the heterogeneous nature of leak site layouts, a rigid parsing script is insufficient. The Transformation layer implements a *Registry Pattern* within the Parser Engine to ensure architectural scalability and maintainability.

From an architectural perspective, this pattern decouples the core routing logic from the group-specific extraction rules. For each monitored ransomware group (e.g., LockBit, ALPHV, Clop), a dedicated parsing method is registered in a centralized structure as a pair of keywords and corresponding functions. When the raw HTML is received, a dispatcher evaluates the source URL's host, matches the identifying keyword, and dynamically delegates the extraction task to the appropriate parser.

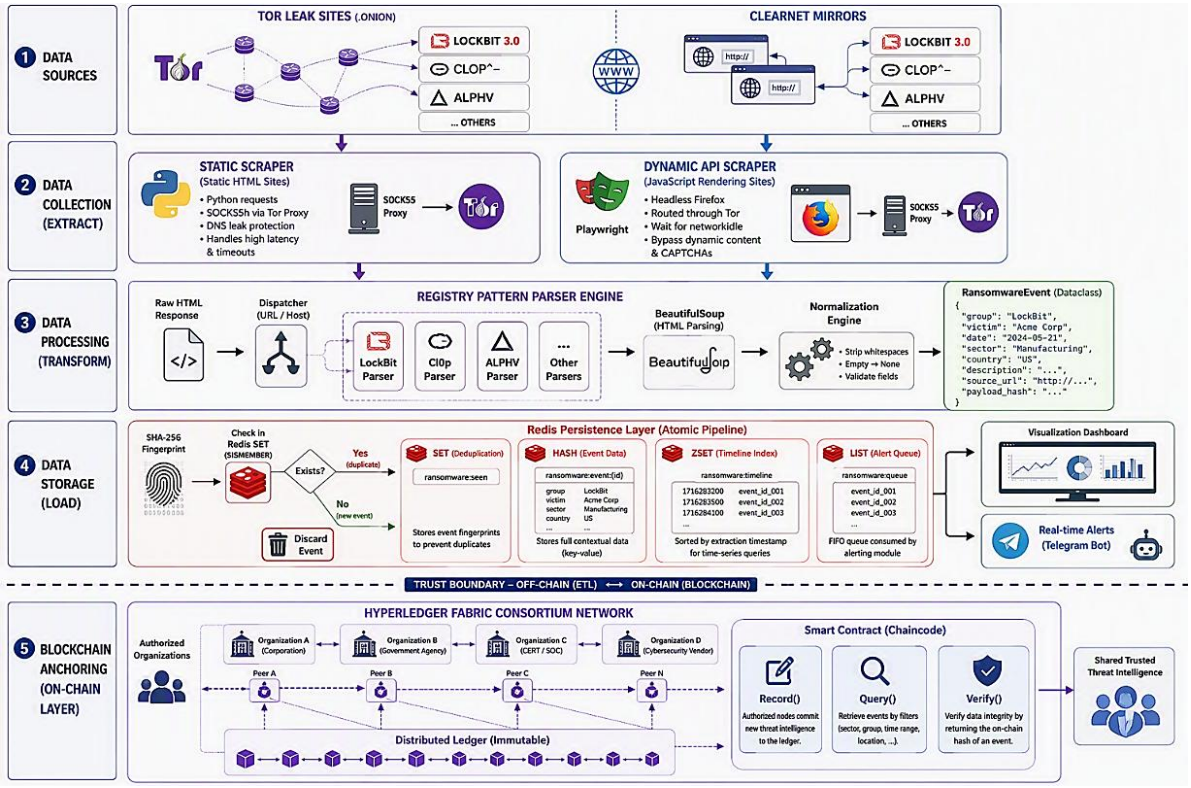


Fig. 1. Architecture of the proposed system.

This modular design allows the system to seamlessly add support for newly emerging ransomware groups without modifying the existing routing infrastructure. The ultimate goal of this phase is to map the highly unstructured data into a standardized `RansomwareEvent` dataclass. This normalized data model strictly defines mandatory fields – namely, the ransomware group, victim name, and source URL – alongside optional contextual fields such as publication date, attack description, industrial sector, and geographic location.

At the implementation level, the extraction is performed using the `BeautifulSoup` library, which navigates the Document Object Model (DOM) tree of the HTML response to locate the elements of interest. For instance, the specific parser for the LockBit group iterates over HTML blocks identified by specific CSS selectors, such as `div.post-card`, and targets child elements like `.post-title` or `.post-date` to extract the required intelligence. To handle the frequent inconsistencies and missing data in the raw HTML, the parser integrates a Normalization Engine implemented via a dedicated static method that strips superfluous whitespaces and safely converts empty strings to `None` values. By sanitizing the inputs before populating the data class, the Normalization Engine guarantees that only clean, standardized, and validated data is propagated to the subsequent phases of the pipeline, implicitly validating the mandatory fields and preventing the propagation of invalid data.

4.3. Data Storage

The primary architectural objective of this layer is to enforce strict data integrity, organize information for rapid retrieval, and prevent data bloat. To achieve this, the storage architecture explicitly decouples write-oriented tasks, such as deduplication, from read-oriented tasks, like temporal aggregation and querying. A critical requirement at this stage is

rigorous deduplication: because volatile leak sites frequently mirror content or republish past claims, the system must uniquely identify threat events to prevent redundant alerts and database inconsistencies. By staging and structuring the intelligence locally, the architecture guarantees that only verified, unique, and well-formatted events are propagated to the alerting modules and eventually anchored to the immutable blockchain.

At the implementation level, the Persistence Layer leverages Redis, an in-memory NoSQL key-value store selected for its exceptional read and write speeds. To enforce the deduplication requirement, the system computes a SHA-256 cryptographic hash for each normalized event by combining the ransomware group, victim name, and publication date. This unique fingerprint is subsequently verified against a Redis `SET` using the `SISMEMBER` command; if the fingerprint already exists, the insertion is aborted. When an event is identified as new, an atomic transaction utilizing a Redis pipeline simultaneously populates multiple native structures without race conditions.

This storage mechanism operates across three primary data structures:

Event Hash: Contextual fields of threat intelligence are stored in a Redis `HASH` (`ransomware:event:{id}`) to allow structured key-based retrieval.

Temporal Index: The event identifier is added to a `ZSET` (`ransomware:timeline`), which is automatically sorted by its extraction timestamp. This chronological index enables highly efficient paginated time-series queries for the visualization dashboard using the `ZREVRANGE` command.

Alert Queue: The event identifier is also pushed to a `LIST` (`ransomware:queue`) that functions as a First-In-First-Out (FIFO) queue. This queue is immediately consumed by a dedicated alerting module to dispatch real-time notifications via other channels (e.g., we included a Telegram bot).

4.4. Blockchain Anchoring

While the ETL layer operates locally to aggregate and normalize intelligence, sharing these data across organizational boundaries introduces complexities regarding trust, provenance, and data integrity.

To address these challenges, our architecture integrates a second pillar represented by a blockchain, which acts as a decentralized sharing infrastructure to democratize access to refined intelligence while ensuring data integrity, provenance, and accountability. This design choice is necessary due to the severe limitations of both centralized threat intelligence platforms and public ledgers when handling leak-site intelligence. Indeed, traditional centralized platforms rely on a single operator for trust, governance, and data protection, inherently creating a single point of failure. In these systems, data access and write attribution are entirely operator-dependent and typically gated behind restrictive subscription models. Conversely, while public blockchains (e.g., Ethereum) eliminate the single point of failure through anonymous validators, they introduce two critical flaws for threat intelligence. First, because the ledger is entirely public, sensitive intelligence could be inadvertently exposed to the very threat actors being monitored. Second, high-volume data anchoring incurs volatile and unsustainable transaction costs (gas fees) that scale directly with the adversary's activity.

A consortium blockchain elegantly solves the shortcomings of both paradigms. By distributing trust and governance across a majority of approved members, it eliminates the single point of failure while guaranteeing verified write attribution tied to known member organizations rather than pseudonymous addresses. Crucially, it restricts participation and read access strictly to authenticated, trusted entities (e.g., allied corporations, governmental agencies, and cybersecurity vendors). This permissioned nature ensures robust data protection, keeps sensitive intelligence hidden from adversaries, and completely eliminates per-record financial barriers, enabling high-throughput data sharing regardless of the threat actors' activity volume.

At the implementation level, the framework utilizes Hyperledger Fabric to orchestrate the shared intelligence. The interaction with the immutable ledger is governed by dedicated smart contracts (referred to as chaincode within the Fabric ecosystem). Once the ETL pipeline processes, deduplicates, and stages a new threat event successfully, the system invokes the smart contract to anchor the structured data on-chain. We designed three core functions to interact with the ledger:

- *Record*: A core recording function allows authorized nodes to securely commit new intelligence to the state database. It receives the normalized threat details (i.e., the ransomware group, victim identity, affected sector, geographic location, timestamp, and payload hash) and writes them to the ledger;
- *Query*: To facilitate seamless data discovery, the smart contract exposes granular retrieval functions. These interfaces take specific filtering criteria (e.g., an industrial sector, the name of a threat actor, or a specific time range) as input and return a structured list of all matching threat events recorded in the chain;
- *Verify*: A dedicated verification function enables consumers to independently validate the integrity of data. By inputting a unique event identifier, the function outputs the corresponding cryptographic fingerprint securely recorded on the blockchain. This allows users to verify that the detailed threat data fetched from off-chain storage repositories (such as local databases) perfectly match the immutable on-chain record.

This hybrid architecture (i.e., local ETL and Redis for high-speed processing, paired with Hyperledger Fabric for immutable attestation) ensures high throughput without sacrificing the cryptographic guarantees of decentralized ledgers.

5. Conclusions

The structural volatility of dark web leak sites and the heavily fragmented nature of current CTI sharing mechanisms severely hamper the collective defense against modern ransomware campaigns. In this paper, we proposed a comprehensive framework that addresses these challenges through the synergy of an advanced ETL pipeline and the technology of consortium blockchain.

By implementing resilient Tor scraping strategies capable of bypassing dynamic rendering obfuscations and standardizing heterogeneous data via a Registry-pattern parser, the system ensures high-quality intelligence aggregation. The integration of Redis guarantees efficient deduplication and high-speed data access, while the Hyperledger Fabric blockchain democratizes intelligence, providing an immutable, trusted, and cost-effective ledger for consortium members.

Future developments of this framework will focus on the integration of Large Language Models (LLMs) to construct adaptive parsers that automatically infer HTML structural changes.

References

[1]. S. Razaulla, C. Fachkha, C. Markarian, A. Gawanmeh, et al., The age of ransomware: A survey on the evolution, taxonomy, and research directions, *IEEE Access*, Vol. 11, 2023, pp. 40698-40723.

[2]. S. Jawad, H. M. Ahmed, Machine learning approaches to ransomware detection: A comprehensive review, *International Journal of Safety and Security Engineering*, Vol. 14, Issue 6, 2024, pp. 1827-1834.

[3]. S. Mehnaz, A. Mudgerikar, E. Bertino, RWGuard: A real-time detection system against cryptographic ransomware, in *Proceedings of the 21st International Symposium on Research in Attacks, Intrusions, and Defenses (RAID)*, 2018, pp. 114-136.

[4]. M. Akbanov, V. G. Vassilakis, M. D. Logothetis, Ransomware detection and mitigation using software-defined networking: The case of WannaCry, *Computers & Electrical Engineering*, Vol. 76, 2019, pp. 111-121.

[5]. A. Montieri, D. Ciuonzo, G. Aceto, A. Pescape, Anonymity services Tor, I2P, JonDonym: Classifying in the dark (web), *IEEE Transactions on Dependable and Secure Computing*, Vol. 17, Issue 3, 2020, pp. 662-675.

[6]. D. Homan, I. Shiel, C. Thorpe, A new network model for cyber threat intelligence sharing using blockchain technology, in *Proceedings of the 10th IFIP International Conference on New Technologies, Mobility and Security (NTMS)*, 2019, pp. 1-6.

[7]. R. Riesco, X. Larriva-Novo, V. A. Villagra, Cybersecurity threat intelligence knowledge exchange based on blockchain: Proposal of a new incentive model based on blockchain and smart contracts to foster the cyber threat and risk intelligence exchange of information, *Telecommunication Systems*, Vol. 73, Issue 2, 2020, pp. 259-288.

[8]. F. Menges, B. Putz, G. Pernul, DEALER: Decentralized incentives for threat intelligence reporting and exchange, *International Journal of Information Security*, Vol. 20, Issue 5, 2021, pp. 741-761.

[9]. S. Gong, C. Lee, Bl oCIS: Blockchain-based cyber threat intelligence sharing framework for Sybil-resistance, *Electronics*, Vol. 9, Issue 3, 2020, 521.

[10]. S. Badsha, I. Vakilinia, S. Sengupta, Blocynfo-share: Blockchain based cybersecurity information sharing with fine grained access control, in *Proceedings of the 10th Annual Computing and Communication Workshop and Conference (CCWC)*, 2020, pp. 317-323.

[11]. C. Patsakis, D. Arroyo, F. Casino, The malware as a service ecosystem, in Malware: Handbook of Prevention and Detection, *Springer*, 2024, pp. 371-394.

[12]. L. de Marcos, A. Domínguez-Díaz, C. Cilleruelo, D. Rodríguez, Ransomware threats in manufacturing industry: Patterns based on dark web Tor groups, in *Proceedings of the 55th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops (DSNW)*, 2025, pp. 79-82.

(008)

Cybersecurity Awareness for Web3 Addressing Human-Centric Attack Surfaces in Blockchain Ecosystems

N. P. R. Mendes and J. R. Almeida, Jr.
Department of Electrical Engineering, Polytechnic School, University of São Paulo,
Electrical Engineering Building, Ground Floor, Block B, Av. Prof. Luciano Gualberto,
Travessa 3, No. 158, University City, São Paulo, SP 05508-010, Brazil
Tel.: +55 11 3091-5733
E-mails: niltonpr@usp.br, jorgerady@usp.br

Summary: Web3 uses blockchain for more than cryptocurrency transfers. It also includes wallets, decentralized applications, smart contracts, and digital assets. In these systems, users have more direct control over approvals, funds, and account recovery. This also means that users must make security decisions that are different from normal Web2 systems. This paper presents a simple Web3 cybersecurity awareness framework based on studies about human security behavior and real operational cases. The comparison shows that basic security practices, such as checking sources and keeping devices safe, are still important. However, they must be adapted to transaction signing, token approvals, seed phrases, irreversible transfers, and wallet recovery. The framework has six control areas and is tested against four common situations: normal credential phishing, malicious token approval and wallet draining, private-key compromise, and stablecoin intervention by an issuer. The results show which controls come from Web2 and which ones are specific to Web3 authorization and recovery.

Keywords: Blockchain security, Web3, Cybersecurity awareness, Human factors, Transaction signing, Wallet security.

1. Introduction

Blockchain is now used for more than digital money. It supports decentralized applications, digital assets, smart contracts, and other online services. This wider environment is usually called Web3. In this paper, Web3 means systems where users work directly with wallets, decentralized applications, smart contracts, and blockchain assets. This does not mean that every service is fully decentralized. It means that blockchain transactions and wallet approvals are a central part of how the user interacts with the system.

Many studies about blockchain security focus on consensus, cryptography, smart contracts, and the correctness of the ledger. These topics are important. However, many real problems do not start because the blockchain protocol is broken. They start when users are convinced to approve, sign, share, or transfer something without fully understanding the result. In these cases, the system may work correctly, but the user authorizes an action that they did not really want.

This changes the role of the user. In Web2, the user usually logs in to an account managed by a platform. In Web3, the user may also control the private key, approve spending permissions, choose the destination address, and decide if a smart-contract action should continue. Because of this, the user becomes part of the authorization process. A simple decision on the screen can cause a final blockchain transaction.

Traditional cybersecurity awareness is still useful because phishing, fake identities, unsafe devices, and social engineering also affect blockchain users. The problem is not that Web2 security rules are useless. The problem is that they do not fully explain what is being approved, what kind of credential is being used, how blockchain execution works, and what recovery options exist in Web3. The main research question is: which awareness controls can come from ordinary cybersecurity, and which controls must be changed or added for Web3?

This paper has three main contributions. First, it compares Web2 and Web3 attacks that depend on human decisions. Second, it proposes a framework with six awareness areas related to wallets, keys, decentralized applications, and incident response. Third, it gives an initial validation by matching the framework with previous studies and four practical scenarios. The goal is not to replace technical blockchain security. The goal is to identify what users need to understand when they approve blockchain actions.

2. Human-Centered Security in Web3

Research about cryptocurrency users shows that many people do not fully understand security decisions when they need to make them. Mai et al. interviewed 29 people and found misunderstandings about private keys, anonymity, and fees. Their study shows that current tools do not always protect users from wrong ideas, even when these wrong ideas can cause money loss or privacy problems [2].

In their empirical analysis of mobile cryptocurrency wallets, Voskobojnikov et al. examined a large dataset of 45821 user reviews. From this corpus, they identified 6859 reviews specifically related to user experience, which they used to investigate usability issues in cryptocurrency wallet applications [4]. This matters because users must

understand the consequences before they sign a transaction.

Si et al. studied protective concerns in blockchain-based environments by looking at how individual participants recognize and handle risk, instead of treating the problem only as a matter of protocol or platform weaknesses. Their analysis included threats faced by users when dealing with ledger infrastructure, distributed software services, social interaction channels, and external crypto-asset providers [1]. Their work shows important differences between Web2 and blockchain-based ecosystems, especially because more responsibility for safe use is transferred to the individual [1]. Houy et al. also reached a similar conclusion when studying wallet security. Their review separates threats to keys, wallet software, devices, and users, showing that wallet security cannot be solved by only one technical measure [3].

These studies show that Web3 needs its own awareness layer. At the same time, they also show that Web3 is not completely different from Web2. Users still need to check identities, sources, devices, and signs of attacks. The difference is the object being approved and the possible consequence. For example, the user may not be typing a password into a fake website. Instead, the user may be giving a smart contract permission to move tokens, signing data that is hard to read, or exposing a seed phrase that controls many accounts.

2.1. Web2 and Web3 Attack Conditions

Table 1 shows the main differences between Web2 and Web3 security conditions for users. The comparison is not perfect in every case. Centralized exchanges, custodial wallets, stablecoin issuers, and hosted websites are also part of Web3. Also, some Web2 systems use strong cryptography and final actions. The table shows the most common pattern that users face, not a strict technical rule.

Table 1. Comparison of Web2 and Web3 human-centered security conditions.

Point	Web2	Web3
Credential	Password; account	Key; seed; wallet
Approval	Login/action	Signature/allowance
Execution	Server validated	Contract executed
Recovery	Reset/support	Move/revoke
Finality	Often reversible	Usually final
Visibility	Clear platform UI	Unclear call data

The comparison shows that the main difference is not only decentralization. The important point is where the approval happens and what can be done after a mistake. In Web3, the user often approves the action directly, but the real effect may be hidden behind wallet messages, long addresses, contract functions, and token permissions. Awareness must therefore help users understand an action before it becomes valid, not only teach them to identify suspicious messages.

3. Methodological Development

The framework was created using a structured design process based on literature. It is not presented as a full systematic review or as a completed user experiment. This method was chosen because the first goal is to define the awareness controls and explain how they are different before testing them with a larger group of users.

The first step was choosing the sources. The analysis used peer-reviewed studies about cryptocurrency mental models, wallet usability, user risks in Web3, and threats to wallet security [1-5]. These sources were selected because they give different types of evidence, such as interviews, large review analysis, literature review, and comparison between Web2 and Web3 risks.

The second step was comparison. The findings were organized into six points: type of credential, object being approved, execution model, finality of the transaction, recovery authority, and institutional intervention. This produced the Web2-Web3 comparison in Table 1. A control was considered Web3-specific when the user behavior depended on blockchain features, such as signing authority, seed phrases, token approvals, smart-contract execution, or limited recovery.

The third step was creating the controls. Each repeated threat or misunderstanding was connected to a user decision and then to a preventive or response action. For example, confusion about private keys was connected to seed phrase and key protection. Wrong ideas about irreversible transactions were connected to checking transaction intent. Fake decentralized-application interfaces were connected to awareness about smart-contract interaction.

The fourth step was an initial validation. Four scenarios were used to see if the framework separates Web2 controls from Web3 controls and if it covers both prevention and response. The scenarios are: normal credential phishing, malicious token approval and wallet draining, private-key compromise, and intervention by a stablecoin issuer. This validation is analytical. It checks coverage and logic, but it does not prove that the framework reduces losses in real life.

4. A Web3-Specific Awareness Framework

The proposed framework has six control areas. These areas are connected, but each one helps users at a different moment in the decision process. The first two are close to normal Web2 awareness. The other areas are more specific to wallets, blockchain approvals, and recovery problems.

4.1. Identity and Source Verification

Identity and source verification means checking whether the website, application, support channel, or request is legitimate before connecting a wallet or

signing. This extends normal phishing awareness to Web3, where an attacker may ask for an approval or signature instead of a password.

4.2. Wallet Hygiene and Access Security

Wallet hygiene includes safe devices, updated software, hardware wallets when needed, separated wallets, and small exposed balances. Wallet separation is especially important in Web3: one wallet can be used for testing decentralized applications, while another protects long-term assets. This reduces losses from unsafe approvals or websites [5].

4.3. Seed Phrase and Private-Key Protection

Seed phrases and private keys are not the same as normal passwords. A password is usually checked by a service and can often be reset. A seed phrase can recreate the wallet authority itself. If someone sees it, that person may control the wallet without using the original app. Users must know that real support teams do not need the seed phrase, that screenshots and cloud notes can create unsafe copies, and that typing the recovery phrase into an unknown website can compromise all accounts linked to it. This responds to the user misunderstandings found in earlier studies [2, 3].

4.4. Transaction Intent and Signature Validation

Transaction-intent validation means comparing what the user wants to do with what the wallet is asking the user to approve. The user should check the asset, destination, amount, network, permission level, and expected contract. Users must also know the difference between connecting a wallet and approving tokens. Connecting a wallet may only show an address, while a token approval can allow a contract to move tokens. Unlimited approvals can create large losses if the contract or interface is malicious. Because of this, Web3 awareness must teach what approvals and signatures mean, not only how to notice suspicious links.

4.5. Decentralized-Application Interaction Awareness

Decentralized-application awareness means understanding how wallets interact with smart contracts. Users should know that a nice-looking website is not the same thing as the actual code that will run on the blockchain. The contract address, method, network, and approval scope decide the real on-chain effect. Crypto drainers show this problem. They copy real projects and convince users to connect wallets and approve harmful transactions. Different from normal credential phishing, the attacker may not need to steal a username and password. The victim may approve the theft with a valid signature [7].

4.6. Incident Recognition and Response Readiness

Incident response in Web3 depends on what was compromised. If only a token approval was given, removing that approval may stop more loss. If the private key or seed phrase was compromised, removing approvals is not enough because the attacker has the same signing power as the user. Any remaining assets may need to be moved to a new wallet, and the old wallet should not be trusted. Public guidance for self-custody wallets also says that transactions usually cannot be reversed and that a compromised recovery phrase requires moving away from the affected wallet [11]. This means that users must distinguish between a bad approval, a leaked secret, and an asset that may involve an issuer or intermediary.

Emergency response may include moving remaining assets quickly, using higher-priority or private transactions, and revoking unsafe approvals when possible [10]. However, these actions are not guaranteed because attackers may use automated sweepers or competing transactions.

Some assets also include issuer controls. For example, stablecoin issuers may block or freeze addresses in defined cases [8, 9]. This does not make all transfers reversible, but it shows that Web3 recovery depends on the asset, issuer, legal basis, and timing.

5. Preliminary Validation

The initial validation connects the six control areas to four scenarios. The goal is to check if the framework can separate Web2 controls from Web3 controls, cover the main risks, and give a clear response logic. Table 2 identifies the main controls with a direct preventive or response role in each scenario. It does not mean that the controls give complete protection.

Table 2. Preliminary scenario-based validation of the proposed framework.

Scenario	Controls	Web3 result
Credential phishing	Source; device	Mostly Web2 unless signing appears.
Approval/drainer	Source; intent; dApp; revoke	Loss may follow a valid signature.
Private-key leak	Seed/key; hygiene; response	Reset fails; move assets.
Issuer stablecoin	Record; escalate	Freezing may depend on issuer.

The phishing scenario confirms that part of the framework comes from Web2: fake login pages still require source checking and safe-device habits. The Web3 difference appears when the same social

engineering asks for wallet connection, seed phrase disclosure, or transaction approval. The malicious-approval scenario shows that identity checking is not enough, because a user can give a valid signature that allows a smart contract to move assets. Therefore, transaction-intent validation and decentralized-application awareness are necessary additions to ordinary phishing training.

The private-key scenario shows the limits of normal account recovery: changing a password does not remove an attacker's signing power, so assets may need to be moved and the old wallet abandoned. The stablecoin scenario shows that some Web3 assets have mixed recovery paths, where issuer or legal intervention may be possible but not guaranteed. These cases show why the framework includes both prevention and response controls.

According to the FBI's 2025 IC3 Annual Report, cryptocurrency-related investment fraud produced reported losses of approximately USD 7.2 billion, with many cases involving deception, impersonation, fraudulent platforms, and user-initiated transfers rather than failures in blockchain protocols [6].

6. Discussion

The comparison and validation show that Web3 awareness should extend, not replace, normal cybersecurity awareness. Source checking, device safety, and resistance to social engineering remain basic controls, but they are not enough when users must understand blockchain approvals or respond without a normal recovery service. Wallets and applications should therefore explain transaction intent, permission scope, contract information, and risks in ways users can understand [2].

Another point is the meaning of decentralization. Web3 includes self-custody wallets, centralized exchanges, hosted websites, bridges, stablecoin issuers, and applications with administrative keys. The recovery process depends on the system and on the asset. Awareness programs should explain these differences before an incident happens. Users need to know which controls are personal, which depend on contracts, and which depend on an issuer or intermediary.

A practical implication is that Web3 awareness should be taught as a decision process, not only as a list of warnings. Before signing, users should identify the asset, network, contract, permission scope, and expected result. After an incident, they should identify whether the problem is a stolen password, an unsafe approval, a leaked seed phrase, or an asset with issuer-level controls. This distinction helps avoid weak responses, such as changing an application password when the attacker already has signing authority. It also helps institutions design training that separates ordinary phishing behavior from blockchain-specific authorization mistakes.

The framework has limitations. The sources were selected purposefully, not through a full systematic review, and the validation is scenario-based rather than statistical. It also focuses on end users, not developers, validators, governance, or institutional custody. For this reason, the framework should be read as an initial awareness model that still needs empirical testing.

Future work should turn the six areas into learning objectives and test them with users. Experiments could compare normal cybersecurity training with Web3-specific training and measure task accuracy, confidence, false alarms, and retention.

7. Conclusions

Blockchain security cannot be studied only by looking at protocols and cryptography. In Web3, users often control the credentials and approvals that create blockchain effects. A transaction can be technically valid and still be a security failure if the user signed it because of a wrong idea or deception.

This paper explained the difference between normal cybersecurity awareness and Web3-specific awareness. The result is not a complete separation. Identity checking, source checking, safe devices, and incident recognition are still necessary. Web3 adds other needs, such as checking transaction intent, understanding approvals, protecting seed phrases, understanding smart-contract interaction, and responding to incidents where transactions may be irreversible and signing authority may be shared with an attacker.

The proposed six-area framework was created from peer-reviewed studies about human-centered security and was first checked against four practical scenarios. The analysis shows where Web2 controls are enough, where they need changes, and where new Web3 controls are needed. The broader goal is to make cybersecurity awareness match the real authorization and recovery conditions of blockchain systems.

Acknowledgements

This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES), Finance Code 001.

References

[1]. J. J. Si, T. Sharma, K. Y. Wang, Understanding user-perceived security risks and mitigation strategies in the Web3 ecosystem, in *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI)*, 2024, 974.

[2]. A. Mai, K. Pfeffer, M. Gusenbauer, E. Weippl, et al., User mental models of cryptocurrency systems: A grounded theory approach, in *Proceedings of the 16th Symposium on Usable Privacy and Security (SOUPS)*, 2020, pp. 341-358.

[3]. S. Houy, P. Schmid, A. Bartel, Security aspects of cryptocurrency wallets: A systematic literature review, *ACM Computing Surveys*, Vol. 56, Issue 1, 2024, 4.

[4]. A. Voskobojnikov, O. Wiese, M. M. Koushki, V. Roth, et al., The U in crypto stands for usable: An empirical study of user experience with mobile cryptocurrency wallets, in *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI)*, 2021, 65.

[5]. Y. Yu, T. Sharma, S. Das, Y. Wang, Don't put all your eggs in one basket: How cryptocurrency users choose and secure their wallets, in *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI)*, 2024, 353.

[6]. 2025 IC3 annual report, Federal Bureau of Investigation, Internet Crime Complaint Center, Washington, 2026, https://www.ic3.gov/AnnualReport/Reports/2025_IC3Report.pdf

[7]. Understanding crypto drainers, Chainalysis Blog, 2023, https://www.chainalysis.com/blog/crypto-drainers/

[8]. USDC terms, Circle Internet Financial, LLC, 2026, https://www.circle.com/legal/usdc-terms

[9]. Legal terms and risk disclosure statement, Tether Operations, S.A. de C.V., 2026, https://tether.to/legal/

[10]. Flashbots protect documentation: Private transactions and bundle mode, Flashbots, 2026, https://docs.flashbots.net/flashbots-protect/

[11]. MetaMask Support: Unauthorized transactions and compromised wallets, Consensys Software Inc., 2026, https://support.metamask.io/stay-safe/protect-yourself/ive-been-hacked-scammed-unauthorized-transactions-on-my-account/

(009)

Privacy-Preserving Blockchain E-Voting System Based on Zero-Knowledge Proofs and Separated Identity Verification

H. M. Sung [1], T. H. Wu [1], K. C. Tsao [2], T. C. Chiu [2], Q. Y. Chen [1], J. Y. Lin [3] and Y. P. Hung [1]
[1] National Taiwan University, Graduate Institute of Networking and Multimedia, Taiwan
[2] National Taiwan University, Department of Computer Science and Information Engineering 3 National Taiwan University, Trans-disciplinary Bachelor Degree Program, Taiwan

Summary: Voting is applied in diverse contexts, ranging from small-group decision-making to nationwide referendums, many of which require voters choices to remain confidential. Traditional voting relies largely on institutional trust to protect ballot secrecy but often demands substantial labor and cost. Electronic voting can reduce these burdens; however, ensuring anonymity and gaining voter trust remain major challenges. This paper presents a flexible, privacy-preserving blockchain-based e-voting framework that employs Zero-Knowledge Proofs (ZKPs) through the Semaphore protocol and separates identity verification from ballot casting. Organizers can create voting events, select suitable verification methods, and sponsor blockchain transaction fees to reduce user participation barriers. The proposed framework supports diverse anonymous voting scenarios while reducing labor costs and providing transparency, verifiability, and anonymity. Furthermore, this paper discusses how different configuration choices influence the applicability of the proposed framework across various voting scenarios.

Keywords: Blockchain e-voting, Zero-knowledge proofs, Semaphore protocol, Identity verification, Privacy-preserving cryptography.

1. Introduction

Electronic voting (e-voting) has become an increasingly important research topic because of its potential to improve voting efficiency, reduce administrative costs, and simplify voting procedures. Compared with traditional paper-based elections, e-voting systems provide faster ballot collection and automated vote counting while reducing human intervention. However, public trust remains a major challenge because voters expect both ballot secrecy and transparent, verifiable voting results. Ensuring voter anonymity without sacrificing verifiability has therefore become one of the primary research objectives of modern e-voting systems [1].

Blockchain technology has attracted considerable attention due to its decentralized architecture, immutable ledger, and public verifiability. Existing blockchain-based e-voting systems utilize cryptographic techniques such as homomorphic encryption and zero-knowledge proofs to improve voting integrity and voter privacy [2-5]. Although prior studies have proposed different mechanisms to address privacy and usability challenges, identity verification and blockchain interaction remain closely coupled with the underlying voting architecture in some systems. Such coupling may limit deployment flexibility and increase the participation barrier for users unfamiliar with blockchain technologies.

This paper presents a flexible privacy-preserving blockchain e-voting framework based on the Semaphore protocol [6]. By separating identity verification from ballot casting, the proposed framework allows organizers to adopt different identity-verification methods without modifying the underlying voting protocol. Furthermore, the framework incorporates configurable result disclosure and organizer-sponsored transaction relaying to improve practical deployment without requiring voters to directly pay blockchain transaction fees. The proposed framework provides a flexible solution for diverse voting scenarios while preserving transparency, voter anonymity, and public verifiability.

2. Related Work

2.1. Blockchain Voting

Blockchain has been widely explored for e-voting because its decentralized and immutable ledger can improve transparency, auditability, and resistance to vote tampering [7, 8]. Existing studies have demonstrated the feasibility of blockchain-based voting, while subsequent surveys have compared such systems from different perspectives. Huang et al. classified existing approaches according to blockchain type, consensus mechanism, and participant scale [9], whereas Benabdallah et al. examined voter authentication, cryptographic protection, attack resistance, and security properties [10]. Despite these advantages, several challenges remain. Privacy protection and transaction speed are repeatedly identified as major concerns, while voter authentication, scalability, security, and verifiability are also important requirements for practical deployment [2, 3, 10, 11]. Therefore, although blockchain strengthens voting transparency and integrity, additional cryptographic mechanisms are still needed to protect voter privacy and enable anonymous eligibility verification.

2.2. Zero-Knowledge Proofs and Anonymous Eligibility Verification

A zero-knowledge proof (ZKP) lets a prover establish a statement without revealing its private witness [12]; in voting this supports two distinct obligations, proving that a voter belongs to an eligible set and proving that an encrypted ballot obeys the election rules. A common construction for the first is the commitment-tree membership proof: registered voters publish commitments to private secrets, the commitments form a Merkle tree, and a zk-SNARK shows knowledge of a secret behind some leaf while a nullifier prevents reuse. Some systems apply this pattern to voter identification on a blockchain [13], while others achieve anonymity through Idemix credentials and blind signatures on a permissioned chain [14]. Semaphore [6] packages the membership-proof approach as a reusable protocol with Ethereum contracts, where an election-specific scope and a contract-enforced nullifier yield at most one submission per registered identity. Semaphore proves only membership, but the rules for admitting a commitment are external to it, so one submission per identity equals one person, one vote only if admission prevents duplicate identities. Our system adopts this division: an organizer-selected policy governs admission to the election's Semaphore group, Semaphore governs anonymous casting, and a separate zk-SNARK proves validity of the encrypted ballot.

2.3. ZKP-Enabled Blockchain Voting

Existing ZKP-based blockchain voting systems differ in which proof obligation they emphasize. Some target ballot and administrator correctness, using a blockchain bulletin board to verify the tally without decrypting individual ballots [15], or combining homomorphic encryption, non-interactive ZKPs, and an oracle to let voters audit the administrator [16]. Others target anonymity and scalability, building on Semaphore for ranked-choice voting [5] or aggregating ballots off-chain with zk-SNARK batch proofs that keep on-chain state constant and offset voters' transaction costs [4]. Coercion resistance is pursued through fake-credential schemes [17] and, on a permissioned chain with primitives close to ours, deniable verification [18]. Our earlier protocol [19] separated an identity authority that blind-signs ballots from a voting authority that holds the homomorphic decryption key, achieving receipt-freeness when both are honest, and named Semaphore as a replacement for its trusted identity authority. The present system realizes that direction with a pluggable admission step: a Web3 wallet, a designated NFT, or World ID verification each feed the same Semaphore group and Groth16 ballot-validity circuit over EC-ElGamal ciphertexts, so no single off-chain component observes both a voter's identity and choice, while an organizer-sponsored relayer spares voters transaction fees.

3. System Design

In our system, the e-voting architecture is divided into on-chain and off-chain application services to leverage the complementary advantages of both environments. Security- and transparency-critical operations are performed on-chain, while computationally intensive and application-level operations are handled off-chain. This division reduces gas costs while preserving security, transparency, and voter privacy. The off-chain layer comprises the Frontend, Backend, and Tallier. The Frontend prepares identity data and ballots and retrieves public election information; the Backend coordinates the workflow among the other components. The Tallier generates the election key pair during setup and, upon receiving a tallying request from the Backend after the election has ended, aggregates and decrypts the accepted encrypted ballots to compute the final result. The on-chain layer comprises the Voting System and Semaphore contracts, which maintain verifiable election parameters, identity membership, and accepted encrypted ballots.

As shown in Fig. 1, the workflow is divided into four phases, each corresponding to a distinct responsibility and system state. Voting Setup establishes the candidate list, voting rules, and public encryption key required by the subsequent phases. Identity Verification and Registration confirms voter eligibility and registers an anonymous identity commitment before any ballot is submitted, thereby separating voter authorization from ballot selection. Vote Casting uses the verified candidate list and public encryption key, together with the voters registered identity, to prepare and submit an encrypted ballot for verification. After ballot collection has ended, Tallying aggregates and decrypts the accepted encrypted ballots to produce the final result.

3.1. Voting Setup

The Voting Setup phase establishes the election specific rules and cryptographic parameters that govern the subsequent voting process. To support deployment across different application domains, eligibility verification is implemented as an extensible policy selected by the organizer and maintained by the Backend. When registration begins, the backend applies the verifier associated with the policy specified for that election. Keeping this logic outside the core on-chain procedure allows additional eligibility requirements without redesigning ballot submission and verification.

Technically, the backend records the election definition and creates a dedicated Voting Pool. A distinct Semaphore group is also established for each Voting Pool and used for anonymous voter registration. The backend then serializes the candidate list in a deterministic format, computes its Keccak-256 hash, and records the hash in the Voting Pool while retaining the readable list off chain. Because the

backend is not trusted to return the list unchanged during voting, the frontend later recomputes its hash and compares it with the on-chain commitment before preparing a ballot. After creating the Voting Pool, the backend requests an election-specific key pair from the Tallier. The system uses asymmetric EC-ElGamal encryption over BabyJubJub, with a scalar private key retained by the Tallier and a public key represented as the curve point (pk_x,pk_y) published to the Voting Pool. This allows the frontend to encrypt ballots without access to the decryption key.

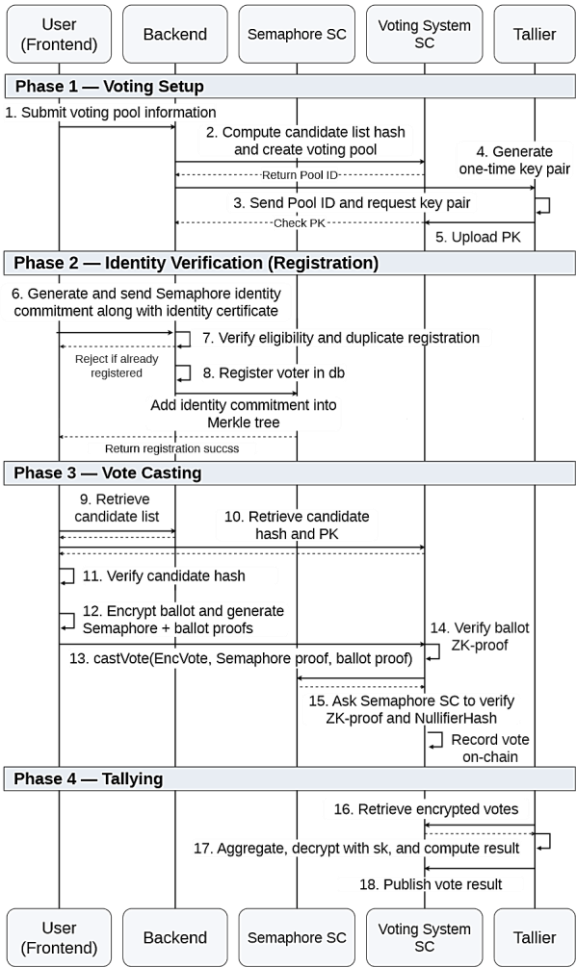


Fig. 1. Flowchart and sequence diagram of the proposed e-voting system architecture.

3.2. Identity Verification (Registration)

Ensuring that only eligible voters participate requires identity-related information to be processed by a responsible party. In conventional electronic voting systems, this responsibility is commonly delegated to a trusted third-party authority. In the proposed system, the Backend performs eligibility verification and manages registration records in its database. However, successful enrollment does not depend solely on the Backends assertion. The corresponding identity commitment must also be added to the elections Semaphore group on-chain, allowing the voter to verify the registration outcome independently through the public contract state. The Backend therefore performs the operational role of a registration authority, while its enrollment actions remain externally observable.

The Frontend generates a Semaphore identity consisting of a private identity secret and a public identity commitment. The secret remains on the voters device, while the commitment is submitted to the Backend with the information required by the selected eligibility mechanism. After verifying eligibility and checking for repeated registration, the Backend adds the commitment to the corresponding Semaphore group. Semaphore is a zero-knowledge membership protocol that represents registered commitments in a Merkle tree. This contract-verifiable group provides authorization independent of the Backends database, while Semaphores zero knowledge proof and nullifier allow anonymous, single use participation. No ballot information is involved in registration, thereby separating identity verification from ballot validation. Although the Backend processes eligibility data, the voters choice is protected before submission and remains unavailable to it. This separation prevents a single off-chain component from observing both the voters identity and plaintext ballot, reducing the plausibility of correlating voters with their choices or manipulating ballots.

3.3. Vote Casting

The Vote Casting phase is responsible for transforming an eligible voters choices into a verifiable encrypted ballot. Its purpose is to preserve the secrecy of those choices while ensuring that each accepted submission uses the published candidate list, satisfies the configured voting rules, and originates from a registered voter who has not previously voted in the same election. Before constructing the ballot, the Frontend retrieves the ordered candidate list from the Backend and recomputes its Keccak-256 hash using the same deterministic encoding applied during Voting Setup. The result is compared with the candidate hash stored in the on chain Voting Pool. This comparison is necessary because the Backend-provided list is not treated as authoritative: removing, replacing, or reordering a candidate would change the meaning of the corresponding positions in the ballot. The Frontend also reads the public encryption key, candidate count, and required number of selections from the Voting Pool, ensuring that ballot construction follows the election state.

After confirming this information, the Frontend represents the voters selections as a binary vector $m = (m_0,\ldots,m_{n-1})$, where $m_i = 1$ denotes a selected candidate and $m_i = 0$ denotes an unselected candidate. The vector follows the candidate order verified through the candidate-list hash. It is then extended to the fixed ballot capacity of 32 positions by setting every position beyond n to zero. This representation gives elections within the configured capacity a

consistent ciphertext format, while the on-chain candidate count identifies the valid positions. The value 32 is a configurable implementation capacity rather than a limitation of EC-ElGamal; supporting a larger election requires a corresponding ballot circuit and contract configuration. Each position in the extended vector is independently encrypted using exponential EC-ElGamal over the BabyJubJub curve. As established during Voting Setup, the Tallier selects a private scalar x and publishes the corresponding public key $PK = xG$, where G is a fixed generator of the BabyJubJub subgroup. For each ballot position, the Frontend samples a fresh nonzero random scalar r_i and computes

$$C_{1,i} = r_i G, C_{2,i} = m_i G + r_i PK \quad (1)$$

The term m_iG encodes the binary selection as a curve point, where $m_i = 0$ produces the identity point, whereas $m_i = 1$ produces G. Each encrypted entry therefore consists of two curve points, represented by the four coordinates $(C^x_{1,i}, C^y_{1,i}, C^x_{2,i}, C^y_{2,i})$. Fresh randomness ensures that identical selections do not produce identical ciphertexts. The BabyJubJub curve is adopted because its point operations can be expressed efficiently over the field used by the Circom circuit, allowing the encryption equations to be verified within a zero-knowledge proof.

Encryption hides the voters selections but prevents the Voting Pool from directly checking whether they follow the voting rules. Because the Frontend runs on the voters device, its local validation can be modified and cannot be trusted by the contract. The system therefore expresses ballot validity as a Circom arithmetic circuit whose algebraic constraints can be checked through a zero-knowledge proof. In this circuit, the selections and encryption randomness form the private witness, while the election public key and voting constraints are public inputs. The circuit verifies, in brief, that the selections are binary, their total number is correct, padded positions contain no votes, and the submitted ballot follows the specified EC-ElGamal encryption. This allows ballot validity to be established without revealing the selections or encryption randomness.

The system uses Groth16 [20] to convert satisfaction of these constraints into a compact proof that can be verified on-chain. A one-time, circuit-specific setup first generates a proving key and a verification key. A Groth16 verifier checks only the proof and its declared public signals; other data included in the transaction does not automatically affect verification. The ciphertext submitted to the Voting Pool must therefore be represented by a public signal so that the proof applies specifically to that ballot. The current ballot contains 32 encrypted positions, each represented by two curve points with two coordinates each, producing 32×2×2 = 128 ciphertext coordinates. Including all 128 coordinates as public signals would enlarge the verification key and increase on-chain verification costs. A compact representation of the ciphertext is therefore used instead.

To reduce this cost, the system represents the ciphertext using a hash-derived challenge γ and a compact fingerprint z. After arranging its 128 coordinates as $ct = (ct_0,\ldots,ct_{127})$, the Voting Pool computes: $\gamma = \text{Keccak256}(ct) \bmod p_{\text{SNARK}}$,

$$\gamma = \text{Keccak256}(ct) \bmod p_{\text{SNARK}}, \quad z = \sum_{j=0}^{127} c\, t_j \gamma^j \bmod p_{\text{SNARK}}, \quad (2)$$

where p_{SNARK} is the prime modulus of field used by Groth16 circuit. The value γ acts as a cipher text dependent evaluation point, and z is the evaluation of the ciphertext-coordinate polynomial at, compressing all coordinates into one field element using only field arithmetic. Keccak-256 is used to derive γ because it can be computed efficiently in Solidity and prevents the submitter from freely selecting an evaluation point to construct a collision. Using the Keccak-256 digest itself as the fingerprint would require the circuit to implement the hash function, introducing many bit-level constraints. Instead, the circuit computes z from the coordinates produced by its verified encryption process, while the Voting Pool computes z from the submitted ciphertext. Publishing γ and z as public signals therefore binds the proof to that ciphertext, except with negligible collision probability.

The Frontend combines the proving key with the private witness and public inputs to generate a proof $\pi = (A,B,C)$. Here, A and C are points in G_1, while B is a point in G_2. These are two groups of elliptic-curve points defined by BN254 for pairing-based cryptography. A point in G_1 contains two field coordinates, whereas a point in G_2 requires four field values because its coordinates are defined over an extension field. These points do not correspond to individual parts of the ballot; they collectively encode evidence that the circuit constraints are satisfied. The verifier receives the proof together with the six public signals

$$u = (u_1, \ldots, u_6) = (z, PK_x, PK_y, n, v, \gamma), \quad (3)$$

where p_{SNARK} is the prime number that defines the circuits finite field; all circuit additions and multiplications are performed modulo this value. The hash-derived value γ determines how the 128 coordinates are combined, while z is the resulting compact fingerprint. Since every coordinate contributes to z, changing the ciphertext will normally change its fingerprint. Deriving γ from the ciphertext through Keccak-256 also prevents the submitter from freely choosing an evaluation point that could make different ciphertexts produce the same z. Keccak256 can be computed efficiently by Solidity but would require many additional constraints if implemented inside the circuit. The Voting Pool therefore derives γ from the submitted ciphertext, while the circuit uses this public value to calculate z through simple field additions and multiplications. The circuit calculates z

from the encryption produced by the voters private selections and randomness, and the Voting Pool independently calculates it from the submitted ciphertext. If the proof is accepted using the contract-derived γ and z, the submitted ciphertext corresponds to the ballot validated by the circuit, except with negligible collision probability.

The Solidity verifier then performs a bilinear-pairing check involving A, B, C, VK_u, and the fixed points contained in the verification key. A bilinear pairing maps a point from G_1 and a point from G_2 into a common target group while preserving their scalar relationships. Since the circuit constraints are encoded into the proving and verification keys during setup, proof points generated from a valid witness satisfy this pairing relation. The verifier can therefore confirm circuit satisfaction without receiving the private witness. BN254 was selected for this proof because Ethereum provides precompiled contracts for its curve arithmetic and pairing operations, making Groth16 verification practical on-chain. Groth16 also produces a constant-size proof regardless of the circuit size and is directly supported by Circom and snarkjs, which provide the browser-side prover and Solidity verifier used in the implementation.

Because the proving and verification keys are circuit specific, changing the ballot capacity requires generating new circuit artifacts and deploying the corresponding verifier. The current circuit supports 32 positions, while the evaluated 256-position configuration still generated a proof within one minute.

3.4. Tallying

Once the Voting Pool closes, the Backend requests the Tallier to start counting, and the Tallier retrieves the encrypted ballots from VoteCast events. These on-chain events provide the authoritative set of accepted ballots. Although every ballot contains 32 encrypted positions, only the first n positions are processed, where n is the candidate count stored in the Voting Pool.

The Tallier aggregates the ciphertexts separately for each candidate before decryption. Let T_i be the number of votes received by candidate i, R_i the sum of the encryption randomness at that position, and $PK = xG_{BJ}$ the election public key. The aggregated ciphertext and its decryption can be expressed as

$$\begin{gathered}\left(\overline{C}_{1,i}, \overline{C}_{2,i}\right) = \left(R_i G_{BJ}, T_i G_{BJ} + R_i PK\right), \\ \overline{C}_{2,i} - x\overline{C}_{1,i} = T_i G_{BJ}.\end{gathered} \tag{4}$$

Because $PK = xG_{BJ}$, applying the private key x removes the accumulated random term and leaves a curve point representing the total vote count. The result is initially expressed as T_iG_{BJ}, rather than the integer T_i. Since the possible value is bounded by the number of accepted ballots, the Tallier uses the Baby-step Giant-step algorithm to recover the corresponding integer efficiently. Repeating this operation for every candidate produces the complete vote-count vector.

After tallying, the Tallier publishes the count vector and election result without disclosing the private key x. For candidate, Tallier computes one aggregate decryption factor $D_i = x\overline{C}_{1,i}$ for this candidate and provides a Chaum Pedersen proof [21] that it was produced using the private key corresponding to PK. To generate the proof, the Tallier selects a fresh random scalar w_i, computes $A_i = w_iG_{BJ}$ and $B_i = w_i\overline{C}_{1,i}$, derives a hashbased challenge c_i, and returns $s_i = w_i + c_ix$. The verifier accepts only if

$$s_i G_{BJ} = A_i + c_i PK, s_i \overline{C}_{1,i} = B_i + c_i D_i \tag{5}$$

The first equality binds the response to the election public key, while the second confirms that the same undisclosed private key was applied to the aggregate ciphertext. The fresh value w_i prevents the proof from revealing x. Anyone can then reconstruct the aggregate ciphertexts from the accepted VoteCast events, verify the proof, and confirm that $\overline{C}_{2,i} - D_i = T_iG_{BJ}$. The tally is therefore publicly verifiable without exposing the election private key or enabling the decryption of individual ballots.

4. Features & Discussion

4.1. Gasless Voting

To reduce the participation barrier, the proposed framework uses an organizer-operated relayer to sponsor voter transactions, rather than the ballot-box cost-offset mechanism of Emami [4] et al. The organizer submits blockchain transactions through a backend relayer and pays the associated transaction fees on behalf of voters. In addition, voters can verify that their ballots have been correctly recorded on the blockchain, improving transparency and end-to-end verifiability. Consequently, voters can participate in blockchain-based voting without paying gas fees, making the proposed framework more accessible for practical deployment.

4.2. Flexible Identity Verification

The proposed framework separates identity verification from ballot casting, allowing organizers to select different identity verification methods without modifying the underlying voting protocol. For example, third-party authentication is suitable for organizational voting and informal polls, while Web3 applications may rely on wallet ownership or designated NFTs to define eligible voters. For voting scenarios requiring strict one-person-one-vote enforcement, stronger verification methods, such as government-issued digital identities or proof-of-

personhood systems such as World ID, can be adopted.

4.3. Options for Disclosing Voting Result

Our system supports two result disclosure options. In formal elections or policy decision-making, results should remain confidential until the voting period ends to prevent bandwagon effects and strategic voting. During this period, encrypted ballots can be aggregated without revealing intermediate results. Alternatively, popularity contests or interactive voting may benefit from real-time result disclosure, as public rankings encourage participation, competition, and community engagement.

5. Conclusion

This paper presents a flexible blockchain-based anonymous e-voting framework using the Semaphore protocol. By separating identity verification from ballot casting, the framework supports multiple authentication methods without modifying the voting protocol. Configurable result disclosure and gasless voting improve deployment flexibility and usability. The proposed framework provides a practical and user-friendly foundation for blockchain-based anonymous voting across diverse real-world voting scenarios.

Acknowledgements

This work was supported by Academia Sinica (Taiwan) under Grant Number [AS-TP-113-M03].

References

[1]. S. Park, M. Specter, N. Narula, R. L. Rivest, Going from bad to worse: From internet voting to blockchain voting, *Journal of Cybersecurity*, Vol. 7, Issue 1, 2021, tyaa025.

[2]. U. Jafar, M. J. A. Aziz, Z. Shukur, Blockchain for electronic voting system – review and open research challenges, *Sensors*, Vol. 21, Issue 17, 2021, 5874.

[3]. R. Taş, Ö. Ö. Tanrıöver, A systematic review of challenges and opportunities of blockchain for e-voting, *Symmetry*, Vol. 12, Issue 8, 2020, 1328.

[4]. A. Emami, H. Yajam, M. A. Akhaee, R. Asghari, A scalable decentralized privacy-preserving e-voting system based on zero-knowledge off-chain computations, *Journal of Information Security and Applications*, Vol. 79, 2023, 103645.

[5]. C. Onur, A. Yurdakul, Electanon: A blockchain-based, anonymous, robust, and scalable ranked-choice voting protocol, *Distributed Ledger Technologies: Research and Practice*, Vol. 2, Issue 3, 2023, 18.

[6]. PSE Lab, Semaphore: A zero-knowledge protocol for anonymous interactions, https://semaphore.pse.dev/

[7]. F. Þ. Hjálmarsson, G. K. Hreiðarsson, M. Hamdaqa, G. Hjálmtýsson, Blockchain-based e-voting system, in *Proceedings of the 11th IEEE International Conference on Cloud Computing (CLOUD),* 2018, pp. 983-986.

[8]. N. Kshetri, J. Voas, Blockchain-enabled e-voting, *IEEE Software*, Vol. 35, Issue 4, 2018, pp. 95-99.

[9]. J. Huang, D. He, M. S. Obaidat, P. Vijayakumar, et al., The application of the blockchain technology in voting systems: A review, *ACM Computing Surveys*, Vol. 54, Issue 3, 2021, 57.

[10]. A. Benabdallah, A. Audras, L. Coudert, N. El Madhoun, et al., Analysis of blockchain solutions for e-voting: A systematic literature review, *IEEE Access*, Vol. 10, 2022, pp. 70746-70759.

[11]. M.-V. Vladucu, Z. Dong, J. Medina, R. Rojas-Cessa, E-voting meets blockchain: A survey, *IEEE Access*, Vol. 11, 2023, pp. 23293-23308.

[12]. X. Sun, F. R. Yu, P. Zhang, Z. Sun, et al., A survey on zero-knowledge proof in blockchain, *IEEE Network*, Vol. 35, Issue 4, 2021, pp. 198-205.

[13]. P. Sangraula, N. B. Adhikari, Zero knowledge proof on top of blockchain for anonymous and verifiable e-voting system, *Journal of The Institution of Engineers (India): Series B*, Vol. 106, Issue 6, 2025, pp. 1861-1872.

[14]. S. Majumder, S. Ray, D. Sadhukhan, M. Dasgupta, et al., ECC-Exonum-EVoting: A novel signature-based e-voting scheme using blockchain and zero knowledge property, *IEEE Open Journal of the Communications Society*, Vol. 5, 2024, pp. 583-598.

[15]. S. Panja, B. K. Roy, A secure end-to-end verifiable e-voting system using zero knowledge based blockchain, *Cryptology ePrint Archive*, 2018, https://eprint.iacr.org/2018/466

[16]. Y. Wu, S. Kasahara, Smart contract-based e-voting system using homomorphic encryption and zero-knowledge proof, *Lecture Notes in Computer Science*, Vol. 13907, 2023, pp. 67-83.

[17]. Z. Yin, B. Zhang, A. Nastenko, R. Oliynykov, et al., A scalable coercion-resistant voting scheme for blockchain decision-making, *IEEE Transactions on Dependable and Secure Computing*, Vol. 23, Issue 3, 2026, pp. 5090-5105.

[18]. Y.-N. Jung, S.-B. Cho, S.-H. Yun, H.-J. Seo, Anonymous e-voting on Hyperledger Fabric: Practical mitigation of the verification paradox and coercion resistance, *Applied Sciences*, Vol. 16, 2026, 6953.

[19]. H.-M. Sung, T. Chen, J.-Y. Lin, Y.-P. Hung, A blockchain-based receipt-free voting protocol suitable for large-scale participation, in *Proceedings of the Pacific Neighborhood Consortium Annual Conference and Joint Meetings (PNC)*, 2025, pp. 64-71.

[20]. J. Groth, On the size of pairing-based non-interactive arguments, *Lecture Notes in Computer Science*, Vol. 9666, 2016, pp. 305-326.

[21]. D. Chaum, T. P. Pedersen, Wallet databases with observers, *Lecture Notes in Computer Science*, Vol. 740, 1993, pp. 89-105.

(012)

Embedded Cryptocurrency Wallets: A Comparative User Study on Security and Usability

Dalibor Detko [1], Juraj Gergel [1], Martin Perešni [1] and Ivan Homoliak [1,2]
[1] Brno University of Technology, Faculty of Information Technology, Czech Republic
[2] Slovak University of Technology in Bratislava, Faculty of Informatics and Information Technologies, Slovakia
E-mail: {xdetkod00, xgerge02, iperesini, homoliak}@fit.vut.cz

Summary: Embedded cryptocurrency wallets replace seed-phrase management with familiar online authentication and provider-mediated key custody. This can reduce onboarding burden, but it relocates trust to identity accounts, Wallet-as-a-Service (WaaS) providers, recovery procedures, account policies, and supporting infrastructure. We compare six web-wallet configurations using a six-domain trust model and an evidence-labelled property matrix. Our timed study ($n = 4$) and separate preference study ($n = 14$) provide exploratory usability evidence. The timing results describe complete configurations rather than authentication mechanisms in isolation, while the preference results are interpreted subject to fixed order and prototype differences. Our findings support configuration-level assessment of embedded custody and explicit treatment of provider and recovery dependencies as authorization paths.

Keywords: Embedded wallets, Account abstraction, Passkeys, MPC, zkLogin, Usability, Operational security.

1. Introduction

Cryptocurrency wallets [6, 3, 5] and modern embedded wallets create and use keys within an application interface, authenticate users through email One-Time Passwords (OTPs), passkeys, social logins, or OpenID Connect (OIDC), and delegate key operations to provider Trusted Execution Environment (TEE) or Multi-Party Computation (MPC) infrastructure. Account Abstraction (AA) can add programmable validation, gas sponsorship, and recovery to conventional Externally Owned Accounts (EOAs) [1, 4]. Prior surveys classify authentication factors and threats [6, 3, 5], while human-centered studies show that operational behavior is influenced by convenience and users' mental models [8, 10]. The central question is therefore not whether a wallet is seedless, but which entities and factors can authorize a transfer or recovery.

2. Trust Domains and Method

We separate DApp intent presentation, authentication, WaaS key custody, account policy, recovery, and infrastructure into distinct trust domains (Fig. 1). Relevant adversaries include identity-account compromise, phishing, device theft, client malware, malicious interfaces, provider compromise, infrastructure manipulation, and recovery-path abuse [9, 2]. Properties are reported as observations, configuration inspection, documentary evidence, or not evaluated, rather than as a single security score.

We evaluated six web prototypes across two user studies. In Study 1, we timed login, transaction submission, and session re-authentication for Privy OTP EOA and Thirdweb passkey AA on Ethereum Sepolia with four participants on a standardized setup (OTP preceding passkey). In Study 2, we collected barriers, fears, a forced convenience-security choice, and preferred configurations across 14 participants guided through Sui zkLogin, ZeroDev passkey AA, Dynamic MPC, and a 2-of-3 software multisignature baseline. We evaluate each study separately to account for differences in configuration stacks and procedures.

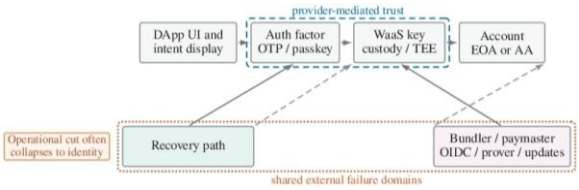


Fig. 1. Trust domains of an embedded wallet. Authorization can depend on shared identity and provider services.

3. Results

In Study 1, median login time was 34.5 s for OTP EOA and 58.0 s for passkey AA; median transaction time was 31.0 s and 44.0 s, respectively (Fig. 2). The timing results describe complete configurations rather than an authentication mechanism in isolation. Two participants preferred OTP for everyday use, whereas three judged passkeys more secure after the tasks.

In Study 2, passkey and Dynamic MPC received six and five first-choice selections, followed by multisignature (3) and zkLogin (0). Eleven of 14 participants selected convenience over security (Fig. 3). Multi-select barrier endorsements (Fig. 4) were led by lack of understanding (9) and security concerns (8), followed by fear of losing access (6) and

risk of mistakes (6). The most frequently selected fears were losing access (11) and hacking or phishing (9). These outcomes are descriptive because prototype order and task design differed.

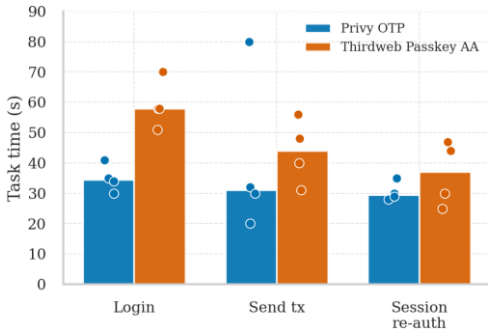


Fig. 2. Study 1 task times. Bars show medians and points show individual observations ($n = 4$).

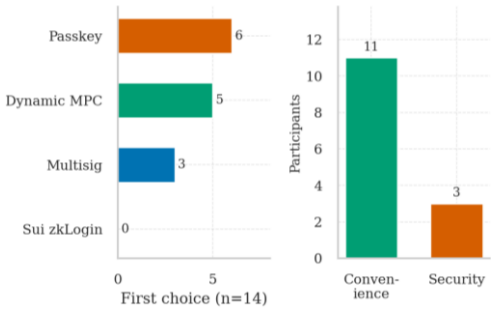


Fig. 3. Study 2 first-choice selections and forced convenience-security choice ($n = 14$).

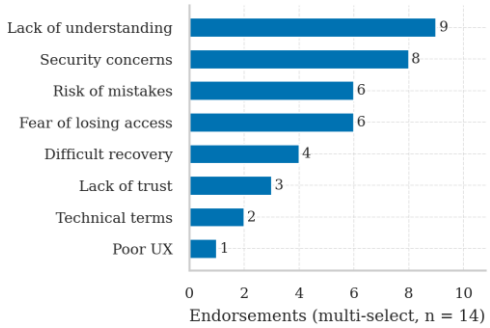


Fig. 4. Multi-select adoption barriers in Study 2 ($n = 14$). Counts are endorsements, not exclusive primary reasons.

4. Security Interpretation and Conclusion

OTP EOA can reduce to the mailbox or provider. Such risks could be mitigated through smart-contract OTP verification or password-authenticated key exchange such as Balthazar [7]. Cloud-synchronized passkeys can reintroduce an identity root, and the Dynamic configuration is gated by Web2 login and linked social recovery. The evaluated zkLogin configuration requires both OIDC and a client-held salt. At $p = 0.05$, independent 2-of-3 compromise is 0.00725, rising to 0.01718 under a 1 % common-mode failure ($c = 0.01$) vs. single-factor baseline (0.05), with break-even $c^* \approx 4.31$ %. Embedded-wallet evaluation should identify the practical root of trust and treat recovery and provider dependencies as part of the authorization policy.

Acknowledgment

This work was supported by the Science Grant Agency VEGA under project 1/0771/26, the Slovak Research and Development Agency under the Contract no. APVV-25-0804, and by the BUT internal project FIT-S-23-8151.

References

[1]. V. Buterin, A. Dietrichs, Y. Weiss, K. Gazso, et al., ERC-4337: Account abstraction using alt mempool, Ethereum Improvement Proposals, 2021, https://eips.ethereum.org/EIPS/eip-4337

[2]. P. Chatzigiannis, K. C. Wang, S. S. Arora, M. Minaei, A composability analysis framework for Web3 wallet recovery mechanisms, in *Proceedings of the IEEE Symposium on Security and Privacy (SP)*, 2025, pp. 1531-1546.

[3]. Y. Erinle, Y. Kethepalli, Y. Feng, J. Xu, SoK: Design, vulnerabilities, and security measures of cryptocurrency wallets, *Computer Networks*, Vol. 273, 2025, pp. 1-22.

[4]. I. Homoliak, D. Breitenbacher, O. Hujnak, P. Hartel, et al., SmartOTPs: An air-gapped 2-factor authentication for smart-contract wallets, in *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies (AFT)*, 2020, pp. 154-167.

[5]. I. Homoliak, M. Perešni, SoK: Cryptocurrency wallets: A security review and classification based on authentication factors, in *Proceedings of the IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, 2024, pp. 1-8.

[6]. S. Houy, P. Schmid, A. Bartel, Security aspects of cryptocurrency wallets: A systematic literature review, *ACM Computing Surveys*, Vol. 56, Issue 1, 2024, 4.

[7]. T. Krajči, S. Olekšák, I. Homoliak, Balthazar wallet: Making password authentication practical on Web3 via OPAQUE and privacy-preserving smart contracts, Cryptology ePrint Archive, 2026, https://eprint.iacr.org/2026/504

[8]. A. Voskobojnikov, O. Wiese, M. M. Koushki, V. Roth, et al., The U in crypto stands for usable: An empirical study of user experience with mobile cryptocurrency wallets, in *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI)*, 2021, 65.

[9]. K. Yan, X. Zhang, W. Diao, Stealing trust: Unraveling blind message attacks in Web3 authentication, in *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2024, pp. 555-569.

[10]. Y. Yu, T. Sharma, S. Das, Y. Wang, 'Don't put all your eggs in one basket': How cryptocurrency users choose and secure their wallets, in *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI)*, 2024, 353.

(013)

Wallet Security Is Compositional: A Failure-Domain-Aware Evaluation of Cryptocurrency Wallet Configurations

Tomáš Fišer [1], Martin Perešíni [1] and Ivan Homoliak [1,2]
[1] Brno University of Technology, Faculty of Information Technology, Czech Republic
[2] Slovak University of Technology in Bratislava, Faculty of Informatics and Information Technologies, Slovakia
{xfiser16, iperesini, homoliak}@fit.vut.cz

Summary: Cryptocurrency-wallet comparisons often assign a single security score to systems that implement distinct custody functions. This study evaluates custody configurations by relating the properties of key storage, signing, account policy, intent display, recovery, and infrastructure to their shared failure domains. We evaluate an executable testbed covering an Ethereum Virtual Machine (EVM) software-key path, a hardware-signer abstraction, Bitcoin Pay-to-Witness-Script-Hash (P2WSH) 2-of-3 multisignature, and a Safe smart account; each finding identifies its evidence source. The testbed establishes selected restoration, authorization, and replay invariants, but cannot establish client-interface usability, phishing resistance, or production-device security. We model threshold authorization under independent and common-mode compromise. With illustrative per-factor compromise probability 0.05, independent 2-of-3 authorization reduces compromise probability from 0.05 to 0.00725; a 1% common-mode event raises it to 0.01718, remaining advantageous over single-key baselines up to break-even rate $c^* \approx 4.31\%$. The resulting framework supports reproducible, configuration-specific security claims in place of product-level rankings.

Keywords: Cryptocurrency wallets, Key management, Multisignature, Failure domains, Operational security.

1. Introduction

Cryptocurrency custody comprises key storage, transaction signing, account authorization, transaction construction and intent display, recovery, and network infrastructure. Browser wallets, hardware signers, Bitcoin multisignature coordinators, and Ethereum smart accounts occupy different layers and may be combined within one configuration. A product-level score obscures these distinctions and may treat inapplicable features as security advantages. Existing surveys classify wallet threats and key-management defenses [1, 2], while authentication-focused research by Homoliak and Perešíni distinguishes factors verified on the blockchain from locally applied controls [3]. While the 2024 factor taxonomy [3] categorized factor verification locations, our six-layer custody model expands beyond factor authentication to cover full compositional custody layers – including key storage, signing execution, intent display, backup recovery, and software update failure domains. Human-centered studies show that operational behaviour cannot be inferred from protocol design alone [4, 5]. This study distinguishes executed invariants from findings supported by source inspection or documentation, and examines how shared failure domains affect threshold authorization. The scope of the study is focused on the application layer of security reference architecture [6].

2. Model and Threats

The six-layer model comprises a key store, signing process, account policy, transaction-construction and intent-display interface, recovery path, and network or update infrastructure. The threat model considers a host attacker controlling one general-purpose computer, a backup attacker obtaining one stored backup copy, a coordinator attacker modifying an unsigned request, and an infrastructure attacker controlling a remote service or update path. Physical extraction from a production hardware device is outside the experimental scope. A hardware signer or air-gapped component is host-safe (or host-isolated) if private keys remain confidential even when connected to a compromised host, provided display integrity is maintained on device. The evaluation considers authorization safety, key confidentiality, transaction-intent integrity, replay resistance, loss tolerance, recovery liveness, and failure domain independence. Each finding is labelled OBS for a retained executable observation, CODE for source inspection, DOC for a standard or peer-reviewed source, or NE when not evaluated. Inapplicable properties are marked NA. The labels are not aggregated into a security score.

3. Reference Testbed and Analytical Method

Our testbed implements four reference paths: a single-key EVM model, a hardware-signer abstraction, a Bitcoin 2-of-3 P2WSH wallet, and a Safe 2-of-3 smart account. Executions run against Bitcoin Core 25.0 regtest and an Anvil local test node (Foundry). The hardware path pairs a Trezor Model T3T1 (T3T1) emulator for setup. The testbed isolates underlying cryptographic and contract logic rather than executing commercial web or desktop client applications. The Safe path interacts directly with canonical version

1.3.0 contracts. Consequently, the testbed specifically evaluates protocol-level invariants rather than human usability, phishing resistance, or physical hardware tampering.

The protocol oracles verify deterministic address derivation, passphrase divergence, signature verification, rejection of modified messages, and threshold authorization acceptance. We evaluate threshold policies under both co-located (single host) and independent failure-domain deployments.

For n authorization factors with threshold t and independent per-factor compromise probability p, the probability of authorization compromise is

$$P_c(t,n,p) = \sum_{k=t}^{n} \binom{n}{k} p^k (1-p)^{n-k} \qquad (1)$$

If a common event compromises every factor with probability c, then $P_{c,\text{shared}} = c + (1-c)P_c$. Thebreak-evencommon-mode probability where 2-of-3 risk matches single-key baseline risk $p = 0.05$ is $c^* \approx 4.31\ \%$. The parameters are sensitivity inputs rather than estimates of incident prevalence.

4. Results

Fig. 1 distinguishes custody layers from their failure domains. Three keys controlled by one process constitute three cryptographic factors but one operational domain. Independent t-of-n authorization requires compromise of t keys, whereas co-location introduces the shared host as a one-event cut set.

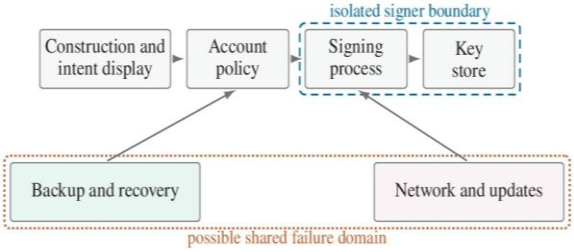


Fig. 1. Custody layers and potential shared failure domains.

The Bitcoin execution constructed a P2WSH 2-of-3 output, restored its address, rejected an incorrect passphrase, produced two signatures, and broadcast a transaction accepted by Bitcoin Core. The EVM execution restored an address, rejected an incorrect password and a modified message, and confirmed a transaction on Anvil. These are protocol rather than client interface observations.

Table 1 presents the conditional property matrix. Hardware isolation can reduce host exposure without changing a 1-of-1 account policy. A 2-of-3 policy tolerates one lost factor only with independent factors and backups. Transaction-intent integrity was not evaluated.

Table 1. Selected properties and evidence.

Property	EVM	HW	BTC	Safe
One key authorizes	Y^{CODE}	Y^{DOC}	N^{OBS}	N^{CODE}
One loss tolerated, independent	N	N	Y^{OBS}	Y^{CODE}
Host-safe, independent	N	Y^{DOC}	Y^{CODE}	Y^{CODE}
Host-safe, co-located	N	NE	N^{CODE}	N^{CODE}
Intent integrity	NE	NE	NE	NE

At $p = 0.05$, compromise probability is 0.05 for 1-of-1, 0.00725 for independent 2-of-3, and 0.001158 for independent 3-of-5 authorization. A common-mode event with probability 0.01 raises the 2-of-3 result to 0.01718. Fig. 2 shows this sensitivity.

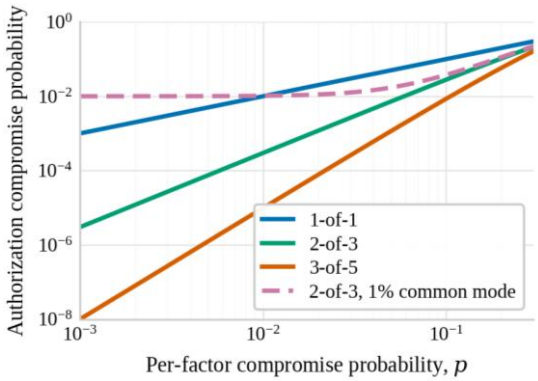


Fig. 2. Threshold sensitivity under explicit assumptions.

5. Relation to Prior Evidence

Prior research classifies wallet threats and authentication factors [1–3], while user studies show that protocol behaviour alone cannot establish operational usability [4, 5]. Application analysis documents configuration-specific risks [7]; EthClipper and blind-message attacks show that key isolation and valid signatures do not guarantee transaction-intent integrity [8, 9]. Recovery mechanisms can add authorization paths that require compositional assessment [10], smart contracts can isolate authorization via air-gapped authenticators [11] or password-authenticated key exchange [12], and zero-knowledge light-client verification mitigates RPC infrastructure dependencies [13]. Each finding is tied to a custody layer and explicit failure-domain assumption. The curves are sensitivity analyses, not attacker-frequency estimates.

Unlike work that designs a new authentication protocol, including prior research on Ethereum transaction authentication [14], this study provides a common unit of comparison for heterogeneous custody components. A client-interface test cannot establish resistance to physical extraction, a signature oracle cannot establish informed consent, and a 2-of-3 execution cannot establish independence for co-located keys. This separation makes claims

verifiable and prevents signer count from being mistaken for failure-domain count.

6. Conclusion

Wallet security depends on the complete custody configuration and the independence of its failure domains. The evaluation replaces product-level rankings with explicit properties, evidence provenance, and conditional analysis. The testbed confirms selected protocol invariants, while the analytical sensitivity model establishes common-mode failure bounds ($c^* \approx 4.31\%$). Future work should use pinned client and firmware versions, separate signers and backups, retained interface evidence for adversarial requests, and measurements of transaction and recovery costs. Usability claims require a participant study.

Acknowledgment

This work was supported by VEGA 1/0771/26, APVV-25-0804, and FIT-S-23-8151.

References

[1]. S. Houy, P. Schmid, A. Bartel, Security aspects of cryptocurrency wallets: A systematic literature review, *ACM Computing Surveys*, Vol. 56, Issue 1, 2024, 4.

[2]. Y. Erinle, Y. Kethepalli, Y. Feng, J. Xu, SoK: Design, vulnerabilities, and security measures of cryptocurrency wallets, *Computer Networks*, Vol. 273, 2025, 111364.

[3]. I. Homoliak, M. Perešíni, SoK: Cryptocurrency wallets: A security review and classification based on authentication factors, in *Proceedings of the IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, 2024, pp. 1-8.

[4]. A. Voskobojnikov, O. Wiese, M. M. Koushki, V. Roth, et al., The U in crypto stands for usable: An empirical study of user experience with mobile cryptocurrency wallets, in *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI)*, 2021, 65.

[5]. Y. Yu, T. Sharma, S. Das, Y. Wang, Don't put all your eggs in one basket: How cryptocurrency users choose and secure their wallets, in *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI)*, 2024, 353.

[6]. I. Homoliak, S. Venugopalan, D. Reijsbergen, Q. Hum, et al., The security reference architecture for blockchains: Toward a standardized model for studying vulnerabilities, threats, and defenses, *IEEE Communications Surveys & Tutorials*, Vol. 23, Issue 1, 2021, pp. 341-390.

[7]. M. Ikram, M. A. Kaafar, I. W. B. Sentana, An empirical analysis of security and privacy risks in Android cryptocurrency wallet apps, *Lecture Notes in Computer Science*, Vol. 13906, 2023, pp. 699-725.

[8]. N. Ivanov, Q. Yan, EthClipper: A clipboard meddling attack on hardware wallets with address verification evasion, in *Proceedings of the IEEE Conference on Communications and Network Security (CNS)*, 2021, pp. 191-199.

[9]. K. Yan, X. Zhang, W. Diao, Stealing trust: Unraveling blind message attacks in Web3 authentication, in *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2024, pp. 555-569.

[10]. P. Chatzigiannis, K. C. Wang, S. S. Arora, M. Minaei, A composability analysis framework for Web3 wallet recovery mechanisms, in *Proceedings of the IEEE Symposium on Security and Privacy (SP)*, 2025, pp. 1531-1546.

[11]. I. Homoliak, D. Breitenbacher, O. Hujnak, P. Hartel, et al., SmartOTPs: An air-gapped 2-factor authentication for smart-contract wallets, in *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies (AFT)*, 2020, pp. 154-167.

[12]. T. Krajči, S. Olekšák, I. Homoliak, Balthazar wallet: Making password authentication practical on Web3 via OPAQUE and privacy-preserving smart contracts, Cryptology ePrint Archive, 2026, https://eprint.iacr.org/2026/504

[13]. M. Perešíni, S. Olekšák, S. Slávka, I. Homoliak, Snarklet: Efficient mobile wallet synchronization with zk-SNARKs, in *Proceedings of the IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, 2025, pp. 1-7.

[14]. I. Dimitrov, W. Prinz, The evolution of private keys: Advancing Ethereum transaction authentication through zero-knowledge proofs, in *Proceedings of the 3rd Blockchain and Cryptocurrency Conference (B2C)*, 2024, pp. 6-9.

(015)

Transactional and Deposit-Like Digital Money: Comparing U.S. Payment Stablecoins and China's e-CNY

Saša Žiković [1,2], Maria Čuljak [3], Bojan Tomić [4] and Hoda Mansour [5]
[1] Institute Digital Assets, Bucharest Academy of Economic Studies, Bucharest, Romania
[2] Pan-European University, European Centre for Business Research, Prague, Czech Republic
[3] Deutsche Bank, Frankfurt am Main, Germany
[4] EFFECTUS University College for Law and Finance, Zagreb, Croatia
[5] University of Business and Technology, Jeddah, Kingdom of Saudi Arabia
E-mail: maria.culjak@db.com

Summary: This paper compares U.S. payment stablecoins and China's digital renminbi (e-CNY) through the distinction between transactional and deposit-like digital money. It argues that remuneration, deposit substitutability, institutional embedding and international monetary channels determine whether a digital claim remains mainly a payment instrument or becomes a competitor to bank deposits. The U.S. framework seeks to contain stablecoins within a non-interest-bearing payments function, while the reformed e-CNY combines public issuance with increasingly deposit-like and bank-integrated features. These design choices shape bank funding, credit intermediation, safe-asset demand and cross-border settlement. U.S. dollar stablecoins can extend an established dollar network and generate Treasury-bill demand; the e-CNY is better understood as infrastructure for expanding renminbi settlement capacity. This paper is an abbreviated and updated conference version of the authors' SSRN preprint [1].

Keywords: Digital money, Stablecoins, e-CNY, Bank intermediation, Deposit substitution.

1. Introduction

Digital money changes finance most profoundly when it changes the economic function of money. A digital balance used primarily for payment and settlement is transactional. A balance that combines liquidity, safety and an economic return, or is protected and managed like a deposit, becomes deposit-like. This distinction matters because commercial bank deposits are both payment balances and a core source of credit funding. Once digital money becomes a credible store of value, it can alter portfolio allocation, bank liability structures and credit intermediation.

The United States and China offer a purposive contrast. The U.S. approach permits privately issued, reserve-backed payment stablecoins while restricting direct issuer-paid interest or yield. The enacted provision is not yet effective: under the official implementation schedule it applies on 18 January 2027, or 120 days after final implementing regulations are issued if earlier [2]. China's e-CNY is public digital money increasingly embedded in the state banking architecture and reported as moving toward interest-bearing, protected balances [3]. The cases therefore isolate economic function from issuer type: private money may remain transactional, while public money may become deposit-like.

This conference paper is an abbreviated and updated version of the SSRN preprint [1]. It adopts a revised title, sharpens the comparative method, corrects the timing of the U.S. yield restriction, separates BIS foreign-exchange data from SWIFT payment-message data, aligns Table 1 with the four analytical dimensions, and treats Treasury-bill estimates as scenarios rather than a model.

Table 1. Comparative framework for digital money design.

Dimension	Transactional	Deposit-like
Remuneration	No direct holder return	Interest, yield or equivalent return
Deposit substitutability	Held mainly for payments	Safe, liquid balance-holding claim
Institutional embedding	Often outside bank balance sheets	Integrated into bank or public architecture
International channel	Extends payment networks	Affects settlement, safe assets and currency use

2. Literature and Contribution

Stablecoin and central bank digital currency research developed from different starting points but converges on deposit substitution. Stablecoin scholarship first emphasized peg stability, reserve quality, redemption credibility and the fragility of privately issued money [6]. As stablecoins became more integrated with conventional finance, attention shifted to their capacity to compete with deposit funding and transmit reserve demand into sovereign money markets.

The central bank digital currency literature began with access to public money, payment-system resilience and monetary sovereignty. It nevertheless confronted the same balance-sheet question: a safe central-bank claim can draw funds from deposits even without interest, while remuneration can strengthen substitution. Tiered remuneration, quantity limits and access rules have therefore been proposed to preserve payment utility while containing financial-stability risks [7, 9].

This paper brings the two literatures into one finance-based framework. It separates issuer identity from monetary function, distinguishes external disintermediation from internal balance-sheet reconfiguration, and adds monetary hierarchy. Identical design features can have different international effects depending on whether the instrument is denominated in an incumbent global currency or a secondary currency building settlement infrastructure.

3. Framework and Method

The analysis uses a comparative institutional design rather than a causal model. The United States and China were selected because they combine systemic monetary relevance with sharply contrasting responses to the same problem: whether digital money should remain a payment claim or acquire deposit-like balance-holding functions. Evidence is assessed along four dimensions: remuneration; deposit substitutability; institutional embedding; and the international monetary channel. The material consists of legislation, official publications, regulatory statements and descriptive market data. The comparison is asymmetric by design because the dollar is the incumbent global currency, whereas the renminbi has a more limited, infrastructure-led role.

Remuneration captures explicit interest, yield or equivalent return. Deposit substitutability concerns whether users can treat the instrument as a safe and liquid store of value. Institutional embedding distinguishes reserve-backed structures outside bank credit creation from balances integrated into commercial-bank or public monetary architecture. The international channel captures effects on cross-border settlement, safe-asset demand and currency use.

The framework does not assume automatic disintermediation. It identifies whether substitution moves funds outside banking or reconfigures balances within it. It also distinguishes aggregate from distributional effects because smaller banks that rely on retail deposits may face pressures not visible in system-wide lending totals.

4. U.S. Stablecoins and Functional Containment

The U.S. strategy can be described as functional containment. Payment stablecoins may expand digital dollar settlement, including across borders, but the restriction on issuer-paid interest or yield is intended to reduce their attractiveness as savings instruments. A non-remunerated stablecoin can still substitute for cash or transaction balances, but it is less competitive with an interest-bearing bank deposit. The objective is to permit payment innovation while limiting migration of structural funding from banks into narrow reserve portfolios.

The banking effect is likely to be distributional as well as aggregate. Large banks may replace deposits through wholesale funding, whereas smaller institutions dependent on retail deposits may face greater pressure. Stablecoins also connect payment growth to safe-asset demand because issuers typically hold liquid reserve assets. The paper therefore treats potential Treasury-bill demand as a scenario calculation: an assumed stablecoin market size multiplied by an assumed bill allocation. A USD 1-3 trillion market with a 40-50 % allocation implies USD 0.4-1.5 trillion in bill demand. These figures are arithmetic illustrations, not forecasts, behavioral estimates or an econometric model.

The international implications are significant. Even a non-remunerated stablecoin can circulate globally as a portable digital dollar claim, particularly where conventional access to dollar banking is costly. The reserve structure adds a second channel by linking digital-dollar growth to demand for short-term sovereign assets. Stablecoins can therefore extend dollar payment functionality while reinforcing the safe-asset base of the existing monetary system.

5. China’s e-CNY and Deposit-Like Integration

The e-CNY began closer to digital cash: sovereign, non-remunerated and oriented toward retail payment modernization. Reported reforms move it toward deposit-like functionality through interest-bearing balances, protection and integration with commercial-bank asset-liability management [3]. The key consequence is not necessarily external disintermediation. Because the instrument operates within a managed monetary and banking architecture, migration into e-CNY may reconfigure the location and control of balances within the system rather than transfer funding to private reserve-backed issuers.

Institutional embedding is therefore central. If e-CNY balances are protected and incorporated into commercial-bank balance-sheet management, the instrument can become attractive as a store of value without reproducing the private stablecoin reserve model. The policy problem shifts from preventing funds from leaving banking to governing how digital sovereign balances interact with bank liabilities, monetary control and credit allocation.

Internationally, the e-CNY starts from a different monetary position. It cannot rely on the dollar’s deep offshore credit markets and safe-asset base. Its more plausible role is infrastructural: improving the

usability and programmability of renminbi settlement in China-linked trade and institutional networks. In combination with payment infrastructure and cross-border digital-currency experiments, it may widen settlement capacity without implying near-term displacement of the dollar [10].

Table 2. Institutional comparison of the two cases.

Dimension	U.S. stablecoins	China e-CNY
Main strategy	Functional containment	Deposit-like integration
Deposit channel	Possible shift to private reserve assets	Reconfiguration within bank/state architecture
International effect	Digital dollar circulation and safe-asset demand	Renminbi settlement capacity

The comparison confirms that issuer identity alone is insufficient. A private instrument may remain transactional because it does not remunerate holders, while a public instrument may become deposit-like because it provides return, protection and institutional integration. Similar mechanisms therefore operate through different balance-sheet channels.

6. International Monetary Implications

Digital design interacts with pre-existing monetary hierarchy. The dollar was on one side of 89 % of over-the-counter foreign-exchange trades in the BIS April 2025 survey [4]. Because each trade contains two currencies, currency shares sum to 200 %; the observation should not be described simply as 89 % of global turnover. Separately, the dollar accounted for 49.25 % of the value of specified SWIFT global-payment messages in February 2026 [5]. SWIFT notes that the tracker covers designated message types and does not constitute complete market or industry statistics. These are distinct observations from different data systems.

The existing dollar network gives U.S. stablecoins a reinforcing effect. Portable digital dollar claims can scale through established demand for dollar liquidity, while reserve portfolios may add demand at the short end of the Treasury market. The e-CNY instead supports gradual network construction. Its value lies in reducing frictions in renminbi settlement and embedding official digital money in selected commercial and institutional corridors. Thus, the same broad technology can deepen an incumbent currency network or help build the operational capacity of a secondary currency.

7. Policy Implications

Regulation should focus on monetary function rather than labels. A prohibition on issuer-paid yield can reduce direct competition with deposits, but indirect rewards, liquidity and perceived safety may still make a claim deposit-like. Supervisors should assess remuneration together with reserve composition, redemption arrangements, distribution channels and concentration of holdings. They should also separate aggregate from distributional effects: modest system-wide changes can conceal material pressure on banks that depend heavily on retail deposits.

International analysis similarly requires caution. Stablecoin growth does not by itself prove a change in reserve-currency status, and e-CNY infrastructure does not establish near-term displacement of the dollar. The relevant channels are narrower and observable: portability of currency claims, demand for reserve assets, settlement costs, payment-network participation and institutional adoption.

8. Conclusions

The central analytical distinction is transactional versus deposit-like digital money, not simply private versus public issuance. Remuneration strengthens deposit substitutability, but protection, liquidity and institutional embedding also matter. The U.S. framework seeks to preserve stablecoins as payment instruments while allowing international digital-dollar circulation. China's reforms move the e-CNY toward a protected, bank-integrated form of official digital money. Their domestic effects differ because one can shift funding toward private reserve assets, while the other can reconfigure balances within a managed banking system.

Digital money also transmits monetary power internationally. Dollar stablecoins may reinforce an established currency hierarchy and safe-asset market; the e-CNY may expand renminbi settlement capacity through infrastructure. These conclusions are institutional rather than causal. Future research should test deposit migration across bank types, the sensitivity of short-term sovereign yields to stablecoin reserve demand, and settlement changes across cross-border digital-currency corridors.

References

[1]. S. Žiković, B. Tomić, M. Čuljak, H. Mansour, Transactional vs deposit-like digital money: Stablecoins, e-CNY and international monetary competition, *SSRN*, 2026, https://papers.ssrn.com/sol3/papers.cfm?abstract_id = 7049098.

[2]. Federal Deposit Insurance Corporation, GENIUS Act requirements and standards for FDIC-supervised permitted payment stablecoin issuers and insured depository institutions, *Federal Register*, Vol. 91, Issue 69, 2026, pp. 1-25.

[3]. China to enhance digital yuan management with deposit features starting 2026, *Xinhua / The State Council of the People's Republic of China*, 2025.

[4]. Triennial central bank survey: Foreign exchange turnover in April 2025, *Bank for International Settlements*, Basel, 2025.

[5]. Global currency tracker: Monthly reporting and statistics, *Society for Worldwide Interbank Financial Telecommunication*, 2026, pp. 1-8.

[6]. G. B. Gorton, J. Y. Zhang, Taming wildcat stablecoins, *University of Chicago Law Review*, Vol. 90, Issue 3, 2023, pp. 909-971.

[7]. U. Bindseil, Tiered CBDC and the financial system, ECB Working Paper No. 2351, *European Central Bank*, 2020, pp. 1-42.

[8]. R. Ahmed, I. Aldasoro, Stablecoins and safe asset prices, BIS Working Paper No. 1270, *Bank for International Settlements*, 2026.

[9]. T. Ahnert, K. Assenmacher, P. Hoffmann, A. Leonello, et al., The economics of central bank digital currency, ECB Working Paper No. 2713, *European Central Bank*, 2022.

[10]. Project mBridge reaches minimum viable product stage and invites further international participation, *Bank for International Settlements*, Basel, 2024.

(016)

Algebraic Analysis of Groth16 and Multi-Framework Empirical Benchmarking of zk-SNARK Constraint Pipelines

Michal Vaňo [1], Martin Perešni [1] and Ivan Homoliak [1,2]
[1] Brno University of Technology, Faculty of Information Technology, Czech Republic
[2] Slovak University of Technology in Bratislava, Faculty of Informatics and Information Technologies, Slovakia
E-mail: {xvanom00, iperesini, homoliak}@fit.vut.cz

Summary: Zero-Knowledge Succinct Non-Interactive Arguments of Knowledge (zk-SNARKs), particularly the Groth16 protocol, provide cryptographic foundations for privacy-preserving computations and succinct state verification in blockchain networks. Despite widespread adoption, selecting an optimal constraint compilation pipeline and framework remains challenging due to varying algebraic overheads. This work presents an algebraic analysis and structural deconstruction of Groth16, detailing the mathematical mechanics from Rank-1 Constraint Systems (R1CS) and Quadratic Arithmetic Programs (QAP) to pairing evaluations over elliptic curves. We analyze the algebraic roles of the five trusted setup trapdoors $(\tau, \alpha, \beta, \gamma, \delta)$ and prover blinding scalars (r, s), detailing how circuit polynomial binding, public-input isolation, and perfect zero-knowledge are maintained. Complementing the theoretical analysis, we introduce an empirical benchmarking framework that evaluates compilation, witness generation, prover latency, memory footprint, and verification performance across Circom/snarkjs, ZoKrates, and Arkworks for constraint scales up to $N = 2^{18}$. Empirical findings show that Rust-native Arkworks achieves up to 4.2× faster prover execution compared to JavaScript-bound snarkjs, while Groth16 maintains a constant 128-byte proof size and sub-2.5 ms verification latency regardless of circuit depth.

Keywords: Zero-knowledge proofs, Groth16, zk-SNARKs, R1CS, QAP, Cryptographic benchmarking, Blockchain security.

1. Introduction and Background

Zero-Knowledge Proofs (ZKPs) allow a prover to establish the validity of a computational statement to a verifier without revealing secret witness information. In blockchain architectures, zero-knowledge proofs enable privacy-preserving asset transactions, stealth address derivation [8], privacy-preserving cross-chain bridging [9], private on-chain financial derivatives [7], anonymous identity verification [11], efficient mobile wallet synchronization [4, 10], and single secret leader election protocols [1, 3, 6]. Among non-interactive zero-knowledge arguments, pairing-based zk-SNARKs, specifically the Groth16 protocol [2], achieve near-optimal verifier efficiency, requiring a constant proof size of 128 bytes (two elements in $\mathbb{G}_1$ and one in $\mathbb{G}_2$) and only three pairing evaluations for verification.

The execution workflow begins by compiling an arithmetic circuit $\mathcal{C}$ into a Rank-1 Constraint System (R1CS), defined by matrices $\mathbf{L}, \mathbf{R}, \mathbf{O} \in \mathbb{F}_p^{m\times(n+1)}$ and witness vector $\mathbf{a} = (a_0, \ldots, a_n)$ with $a_0 = 1$ satisfying $\mathbf{Oa} = (\mathbf{La}) \circ (\mathbf{Ra})$. Through Lagrange interpolation over domain $\{1, \ldots, m\}$, R1CS is converted into a Quadratic Arithmetic Program (QAP):

$$\sum_{i=0}^{n} a_i u_i(x) \cdot \sum_{i=0}^{n} a_i v_i(x) = \\ = \sum_{i=0}^{n} a_i w_i(x) + h(x)t(x), \qquad (1)$$

where $t(x) = \prod_{j=1}^{m} (x - j)$ is the target polynomial and $h(x) = \sum_{k=0}^{m-2} h_k x^k$ is the quotient polynomial.

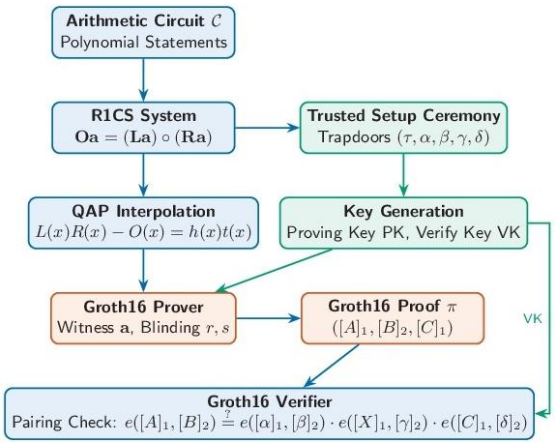


Fig. 1. Groth16 constraint compilation, setup, and pairing verification.

2. Algebraic Trapdoor Analysis

While standard Groth16 supports proof element randomization for zero-knowledge, an un-blinded QAP check $e(A, B) = e(C, G_2)$ would be vulnerable to arbitrary polynomial forgery. Groth16 resolves this vulnerability by introducing five secret toxic waste trapdoors $(\tau, \alpha, \beta, \gamma, \delta)$ during setup and two prover blinding scalars (r, s):

1. **Circuit Binding (α, β):** Multiplicative factors α and β shift the polynomials A and B, forcing a prover to incorporate circuit-bound polynomials $u_i(x)$ and $v_i(x)$;

2. **Input Separation (γ, δ):** Denominators γ and δ separate public inputs ($i \leq \ell$) from private witness inputs ($i > \ell$), enforcing that public statement inputs cannot be tampered with;
3. **Zero Knowledge (r, s):** Random scalars $r, s \in \mathbb{F}_p^*$ sampled by the prover randomize proof elements. Defining $[B]_1 = [\beta]_1 + \sum_{i=0}^{n} a_i v_i(\tau) G_1 + s[\delta]_1$, the proof tuple $([A]_1, [B]_2, [C]_1)$ is given by:

$$[A]_1 = [\alpha]_1 + \sum_{i=0}^{n} a_i u_i(\tau) G_1 + r[\delta]_1, \quad (2)$$

$$[B]_2 = [\beta]_2 + \sum_{i=0}^{n} a_i v_i(\tau) G_2 + s[\delta]_2, \quad (3)$$

$$[C]_1 = \sum_{i=\ell+1}^{n} a_i [\Psi_i^{\text{priv}}]_1 + \\ + \sum_{k=0}^{m-2} h_k [\frac{\tau^k t(\tau)}{\delta} G_1] + s[A]_1 + r[B]_2 - \\ - rs[\delta]_1 \quad (4)$$

Verification relies on a single bilinear pairing check:

$$e([A]_1, [B]_2) = e([\alpha]_1, [\beta]_2) \cdot \\ \cdot e([X]_1, [\gamma]_2) \cdot e([C]_1, [\delta]_2), \quad (5)$$

where $[X]_1 = \sum_{i=0}^{\ell} a_i [\Psi_i^{\text{pub}}]_1$ represents public inputs.

3. Empirical Evaluation

Complementing contract-validated ZK benchmarking frameworks [5], we constructed an extensible benchmarking suite featuring a React.js frontend and a FastAPI backend orchestrating execution runners for Circom/snarkjs, ZoKrates, and Arkworks (Rust). We evaluated performance across constraint counts $N \in [16,\ 262144]$.

As illustrated in Fig. 2, Arkworks achieves superior prover performance due to parallel multi-scalar multiplication (MSM) algorithms in Rust. At $N = 262144$ constraints, Arkworks completes proof generation in 32.8 s, whereas snarkjs requires 138.4 s. While Circom incurs a separate source-to-R1CS compilation phase in JavaScript before prover execution, Arkworks embeds circuit constraints directly into native Rust data structures, concentrating execution time in MSM proving. Fig. 3 shows memory footprint scaling, where Arkworks consumes 116 MB compared to 339 MB for snarkjs and 506 MB for ZoKrates. Across all frameworks, verification latency remains constant at under 2.5 ms with a fixed proof size of 128 bytes.

4. Conclusion

This paper presented an algebraic deconstruction of Groth16 trapdoors and a comprehensive empirical benchmark of zk-SNARK frameworks. Rust-native implementations provide substantial performance gains for heavy proving workloads, while Groth16 retains optimal properties for on-chain verification.

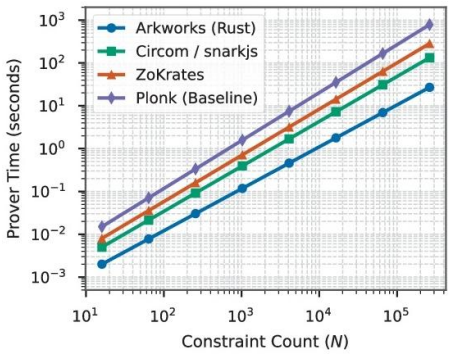


Fig. 2. Prover execution latency across constraint sizes (N).

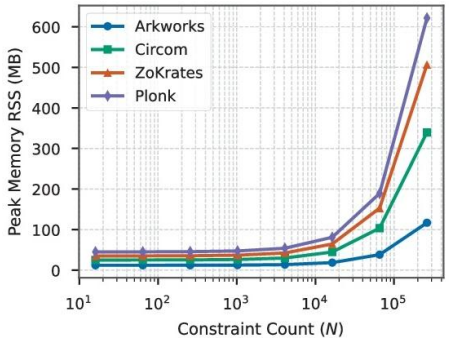


Fig. 3. Peak memory consumption (RSS in MB) during proving.

Acknowledgment

This work was supported by the Science Grant Agency VEGA under project 1/0771/26, the Slovak Research and Development Agency under the Contract no. APVV-25-0804, and by the BUT internal project FIT-S-23-8151.

References

[1]. T. Burianová, M. Perešni, I. Homoliak, Single secret leader election in Ethereum PoS: An empirical study of Whisk and homomorphic sortition under DoS on the leader and censorship, in *Proceedings of the IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, 2026, pp. 1-9.

[2]. J. Groth, On the size of pairing-based non-interactive arguments, *Lecture Notes in Computer Science*, Vol. 9665, 2016, pp. 305-326.

[3]. T. Hladky, M. Peresini, J. Mariani, I. Homoliak, et al., SSLE-DAG: A high-throughput proof-of-stake consensus protocol combining an adaptive DAG with a single secret leader election, *Cryptology ePrint Archive*, 2026, Paper 2026/647.

[4]. I. Homoliak, M. Perešíni, SoK: Cryptocurrency wallets – a security review and classification based on

authentication factors, in *Proceedings of the IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, 2024, pp. 1-8.

[5]. M. Hulek, M. Perešni, I. Homoliak, A unified framework for contract-validated benchmarking of zero-knowledge proving systems, in *Proceedings of the IEEE Conference on Blockchain Research & Applications for Innovative Networks and Services (BRAINS)*, 2026.

[6]. G. Kadianakis, J. Drake, dapplion, EIP-7441: Upgrade block proposer election to Whisk, Ethereum Improvement Proposals, 2023, https://eips.ethereum.org/EIPS/eip-7441

[7]. K. Koštál, R. Bitarovský, L. Mastilak, Hopinzero: zkSNARK-verified privacy for on-chain options with pooled liquidity, in *Proceedings of the IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, 2026, pp. 1-9.

[8]. K. Koštál, L. Častven, M. Géci, ZK-SNARK stealth addresses: Hash-commitment ownership proofs as an alternative to ECDH-based address derivation, in *Proceedings of the IEEE International Conference on Decentralized Applications and Infrastructure (DAPPS)*, 2026, pp. 1-8.

[9]. K. Koštál, A. Rybak, PPBRidge: A zero-knowledge privacy-preserving cross-chain bridge for EVM networks, in *Proceedings of the IEEE International Conference on Decentralized Applications and Infrastructure (DAPPS)*, 2026, pp. 1-8.

[10]. M. Perešíni, S. Olekšák, S. Slávka, I. Homoliak, SNARKlet: Efficient mobile wallet synchronization with zk-SNARKs, in *Proceedings of the IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, 2025, pp. 1-7.

[11]. A. Vožda, M. Perešni, J. Mariani, I. Homoliak, Proof-of-uniqueness: Sybil-resistant privacy-preserving decentralized identity through threshold-OPRF and zk-SNARK registry, in *Proceedings of the International Workshop on Cryptocurrencies and Blockchain Technology (CBT)*, 2026.

(018)

Pre-Mempool Blockchain Transaction Compliance: Regulatory Convergence toward a Pre-Broadcast Screening Mandate

Praveen Giri
QuantChainAnalysis UG, Schorndorfer Str. 5, 70374 Stuttgart, Germany
Tel.: +49 176 85913171
E-mail: contact@quantchainanalysis.com

Summary: Existing blockchain compliance and forensic tools operate exclusively on confirmed, settled transactions, a design constraint that leaves post-hoc detection unable to prevent harm once a transfer becomes cryptographically irreversible. This paper argues that three independent regulatory frameworks, the EU's Regulation (EU) 2023/1113 on information accompanying transfers of crypto-assets, the Financial Action Task Force's revised Recommendation 16, and the United States' proposed GENIUS Act rulemaking, have moved toward a shared technical requirement: screening and, where warranted, blocking a transaction before it is broadcast, not after it settles. We present QARS (Quantum Amplitude Risk Scoring), a patent-pending wallet risk-scoring engine (German patent application DE 10 2026 001 732.7) designed to operate at this pre-broadcast layer, and report a live, dated evaluation against a small, manually verified corpus of real exploit and clean-wallet addresses.

Keywords: Pre-mempool screening, Blockchain compliance, AML/CFT, Quantum amplitude risk scoring, Regulatory convergence, FATF Travel Rule, MiCA.

1. Introduction

On 21 February 2025, the Lazarus Group drained approximately 1.46 billion USD from the Bybit exchange [1]. Established forensic analytics tools correctly identified the malicious wallets, but only after settlement, once the funds were already cryptographically irreversible. This is not a shortcoming of any particular vendor; it is an architectural property of how post-settlement analytics work. A tool that only examines confirmed transactions can, by construction, never intervene before value transfer completes. This paper examines the regulatory pressure now building against that architectural gap, and describes an engineering approach designed to operate before settlement.

2. Regulatory Convergence

Three regulatory instruments have moved toward transfer-time risk control, though through different legal mechanisms and with different degrees of prescriptiveness. Regulation (EU) 2023/1113, applicable since December 2024, requires crypto-asset service providers to maintain internal policies, procedures and controls for the implementation of restrictive measures, in practice sanctions screening, on transfers of crypto-assets [2]. This is a distinct obligation from MiCA Article 68, which governs CASP internal governance arrangements rather than transfer-level sanctions screening, and the two should not be conflated. FATF Recommendation 16 was substantively revised at the FATF's June 2025 Plenary and retitled Payment Transparency, strengthening data-transparency and traceability requirements across the payment chain, including crypto-asset transfers [3]. FATF's own explanatory note to the revised standard clarifies that Recommendation 16 does not itself mandate real-time sanctions screening or a specific method for combining risk indicators; the modality of screening and any indicator-combination approach is left to national regulation and industry practice. QARS's interference-based combination of risk dimensions, described in Section 3, is accordingly presented as an engineering design choice suited to this discretionary space, not as a direct requirement of Recommendation 16. The proposed GENIUS Act rulemaking in the United States, targeting a January 2027 effective date, requires permitted stablecoin issuers to maintain blocking, freezing, and rejection capabilities together with forensic recordkeeping suitable for legal proceedings; treating this specifically as a pre-broadcast capability is an engineering interpretation adopted in this work, not an explicit statutory implementation requirement. Taken together, these three instruments motivate this paper's central premise: that transaction screening architectures built only for post-settlement analysis are increasingly misaligned with where regulatory obligations are converging, even where the instruments differ in how prescriptive they are about timing and method.

3. QARS Architecture

This paper builds on the author's companion SSRN preprint [5], which sets out the QARS mathematical formalism in full and reports a preliminary evaluation on the same corpus described in Section 4; the present contribution's distinct focus is the regulatory-convergence argument set out in Section 2, presented here at extended-abstract length for this venue.

QARS screens wallet activity across thirteen risk dimensions, including mixer proximity, sanctions-list clustering, value and velocity anomalies, travel-rule completeness, and history-relative transaction magnitude. Rather than a linear weighted sum, QARS combines these dimensions with a classical, Born-rule-inspired amplitude interference formalism, an engineering design choice within the discretionary space described in Section 2: each dimension is treated as a probability amplitude, and a coupling matrix encodes how strongly co-occurring risk signals reinforce one another; it executes entirely on classical hardware and is not a quantum computation. This computation runs on every screening request and is disclosed on every report as a named escalation signal, while the live gate decision (pass, enhanced due diligence, or block) is currently derived from a separate, independently auditable weighted-sum score, full integration remaining an item of ongoing work. Every gate decision is recorded in a cryptographically hash-chained forensic ledger intended to support evidentiary use [4].

4. Preliminary Evaluation

We evaluated QARS against a small corpus: thirteen wallet addresses drawn from nine publicly documented exploit events between 2022 and 2026 (including the Bybit, Ronin, and Harmony incidents), and four high-volume production wallets operated by major exchanges and decentralized finance protocols, used as a clean-traffic comparison set. Every address was manually cross-referenced by the author against public incident records, including blockchain forensics post-mortems and, where available, law-enforcement filings, and confirmed against on-chain data; no address was synthetically generated or inferred. Verification here refers to this manual cross-referencing process performed by the author against primary and public secondary sources, not to an external or third-party audit.

Evaluated live against the production pipeline as deployed on 7 August 2026, QARS correctly escalated 11 of the 13 malicious addresses (84.6 % recall) and correctly passed 2 of the 4 clean addresses (50.0 % specificity). Against a bare weighted-sum ablation, the composite escalation layer shows a statistically significant improvement (exact McNemar $p = 0.0010$), though the pooled association between prediction and ground truth does not yet reach significance at this sample size (Fisher's exact $p = 0.22$). The two undetected malicious addresses and the one clean address affected by a deliberate, policy-driven escalation rule are itemized in the companion technical paper. We report this as a preliminary, small-sample, dated finding, not a generalizable performance claim; a corpus of this size cannot establish statistical significance on its own, and expanding it is the immediate next step described below.

5. Conclusion

Three regulatory instruments across different jurisdictions are each moving, through different mechanisms and degrees of prescriptiveness, toward transfer-time rather than post-settlement risk control. We described QARS, an engine designed for this layer, and reported preliminary, dated evidence from a small verified corpus that the underlying detection approach is viable. The open question this work leaves is not whether any single framework mandates pre-broadcast screening in identical terms (they do not) but whether current post-settlement compliance architectures can adapt in time to satisfy whichever combination of these obligations applies before enforcement deadlines arrive.

6. Future Work

Immediate priorities are expanding the labeled evaluation corpus by an order of magnitude, extending calibration to additional blockchain networks beyond the Ethereum-family chains evaluated to date, and documenting a dynamic Bayesian reweighting mechanism for new exploit post-mortems.

This paper's relationship to the author's other QARS publications is as follows: the SSRN preprint [5] presents the full mathematical formalism together with a preliminary evaluation on the same small corpus reported here; this extended abstract presents a focused regulatory-convergence argument alongside that evaluation, at a length suited to this venue; and the full mathematical framework, an expanded empirical evaluation, and production engineering findings are detailed in a companion technical manuscript currently under peer review.

References

[1]. North Korean hackers behind supersized heist (citing Elliptic research), Axios, 2025, https://www.axios.com/2025/02/25/north-korean-hackers-behind-supersized-heist-codebook

[2]. Regulation (EU) 2023/1113 of the European Parliament and of the Council of 31 May 2023 on information accompanying transfers of funds and certain crypto-assets, Article 23, *European Union*, 2023.

[3]. International standards on combating money laundering and the financing of terrorism and proliferation, Recommendation 16, Payment transparency, *Financial Action Task Force*, 2025.

[4]. P. Giri, German Patent Application DE 10 2026 001 732.7, *QuantChainAnalysis UG*, Germany, 2026.

[5]. P. Giri, Pre-mempool blockchain transaction compliance: Toward a regulatory-convergent risk scoring framework, *SSRN*, 2026, https://ssrn.com/abstract=6996883

ISBN 978-84-09-93453-9

