ISSN 2938-2602

BLOCKCHAIN AND CRYPTOCURRENCY

vol. 4

IFSA Publishing

1/2026

Blockchain and Cryptocurrency

Open Access Journal about all Aspects of Blockchains and Cryptocurrencies

Volume 4, Issue 1, March 2026

IFSA Publishing, S.L. • Barcelona

Published by IFSA Publishing, S. L., Barcelona, Spain

The submitting author is responsible for ensuring that contributions of all authors are correct and classified according to the CRediT Classification below.

CRediT Classification:

Conceptualization: Ideas; formulation or evolution of overarching research goals and aims.

Data Curation: Management activities to annotate (produce metadata), scrub data and maintain research data (including software code, where it is necessary for interpreting the data itself) for initial use and later reuse.

Formal Analysis: Application of statistical, mathematical, computational, or other formal techniques to analyze or synthesize study data.

Funding Acquisition: Acquisition of the financial support for the project leading to this publication.

Investigation: Conducting a research and investigation process, specifically performing the experiments, or data/evidence collection.

Methodology: Development or design of methodology; creation of models.

Project Administration: Management and coordination responsibility for the research activity planning and execution.

Resources: Provision of study materials, reagents, materials, patients, laboratory samples, animals, instrumentation, computing resources, or other analysis tools.

Software: Programming, software development; designing computer programs; implementation of the computer code and supporting algorithms; testing of existing code components.

Supervision: Oversight and leadership responsibility for the research activity planning and execution, including mentorship external to the core team.

Validation: Verification, whether as a part of the activity or separate, of the overall replication/reproducibility of results/experiments and other research outputs.

Visualization: Preparation, creation and/or presentation of the published work, specifically visualization/data presentation.
Writing – Original Draft Preparation: Creation and/or presentation of the published work, specifically writing the initial draft (including substantive translation).

Writing – Review & Editing: Preparation, creation and/or presentation of the published work by those from the original research group, specifically critical review, commentary or revision – including pre- or post-publication stages.

IFSA Publishing

Blockchain and Cryptocurrency

Volume 4,
Issue 1, March 2025

e-ISSN 2938-2602

Editors-in-Chief: Prof., Dr. Sergey Y. Yurish,
tel.: +34 696067716, e-mail: ifsa@sensorsportal.com

Editorial Board

Blockchain and Cryptocurrency is an open access, peer review international journal published by IFSA Publishing, S.L. Available in electronic Acrobat (pdf) format.

Blockchain and Cryptocurrency
Open Access Journal

Contents

Volume 4
Issue 1
March 2026

https://bc-ifsa-journal.com

ISSN 2938-2602

Research Articles

Authors are encouraged to submit article in MS Word (doc) and Acrobat (pdf) formats
by e-mail: editor@sensorsportal.com. Please visit journal's webpage with preparation instructions:
https://bc-ifsa-journal.com/guide_for_authors.html

IFSA Publishing, S. L. (Barcelona, Spain)

www.b2c-conference.com

5TH BLOCKCHAIN & CRYPTOCURRENCY CONFERENCE (B2C' 2026)

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

The Blockchain and Cryptocurrency Conference (B2C' 2026) aims to provide a forum for researchers, scientists, engineers, and students from both the industry and the academia to present their latest research findings, advances and innovations on blockchain technologies as well as to helps decision-makers, technologists, and developers understand the value of blockchain to their businesses regardless of industry.

It will feature keynotes, tutorials, peer-reviewed technical paper presentations, posters, demos and exhibitions from world-leading companies, solution vendors, research institutes, and academia. The conference will be also the forum for exchange of the latest innovation results, regulations, policies, standards, and applications in this exciting and challenging area.

Unlike existing, narrowly focused technical conferences and commercial trade events, the congress will cover all technical and social aspects of blockchain and cryptocurrency.

Topics of Interest include but not limited to:

- Blockchain theories, applications and their evolution
- Blockchain based protocols and algorithms
- Smart contracts and distributed ledgers
- Distributed consensus and fault tolerance mechanisms
- Performance optimization of blockchain and decentralized schemes
- Integration of blockchain with other emerging technologies
- Peer-to-peer networks
- Cryptography techniques
- Security and privacy surrounding blockchain technology
- Cryptocurrency
- Cryptocurrency protocols and adoption
- Digital currencies
- Exchanges, trading and mining
- Financial analysis and risk management with cryptocurrencies
- FinTech and DeFi
- NFT
- Web 3.0
- Metaverses
- Fraud detection and financial crime prevention
- Legal, ethical and societal aspects of blockchain and cryptocurrencies
- Regulation and law enforcement of blockchain technology
- New business model
- Case studies (e.g., of adoption, attacks, forks, scams, regulations, policies, standards ...)

Contribution Types

- Keynote and Invited Presentations
- Industrial Presentations
- Regular and Poster Sessions
- Special Sessions
- Panel Discussions and Round Tables
- Virtual Sessions in Zoom (on demand)
- Exhibition

Special Sessions:

Authors are welcome to propose and manage special sessions during the B2C' 2025. Each special session will contain 4-6 papers in a related field as specified above.

Session organizers will get:
- Certificate of Appreciation;
- Free Registration.

Deadlines

Submission (2-page extended abstract):	10 July 2026
Notification of acceptance:	31 July 2026
Registration:	20 September 2026
Camera ready and late registration:	30 September 2026

www.b2c-conference.com/

5TH BLOCKCHAIN & CRYPTOCURRENCY CONFERENCE (B2C' 2026)

One event - three different publications !

1) All registered abstracts will be published in the conference proceedings in the Acrobat (pdf) format with the ISBN and DOI, distributed during the conference and submitted for indexing to the Web of Science.

2) Authors will be invited to submit full-page extended papers to the open access journal Blockchain and Cryptocurrency (ISSN: 2938-2602) for free of charge or other, affiliated open access journals.

3) Authors will be invited to submit book chapter for the open access Book Series on *'Advances in Blockchain and Cryptocurrency'*, Vol. 1, which will be published in 2026.

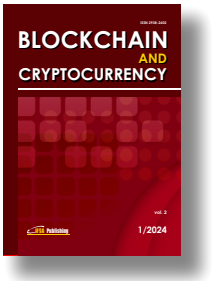

Social Programme

- Welcome Cocktail: 27 October 2026 (20:00-21:00). The Welcome Cocktail will take place in the AC Marriott Gran Canaria Hotel.
- Gala Dinner: 29 October 2026 (20:00-23:00). The Gala Dinner will take place in the conference hotel AC Marriott Gran Canaria.

Organizing Committee

Chairman:
Dr. Sergey Y. Yurish,
IFSA President, Co-founder and CEO of Eco IFSA, Barcelona, Spain

Advisory Chairman
Prof. Dr. Vijyakumar Varadarajan,
University of Technology Sydney, Australia

Conference and Publication Manager
Tetyana Zakharchenko,
IFSA Publishing, S. L., Barcelona, Spain

Business Committee:

Cañete Javier, *Phare Investments Hub, Spain*
Fonts Ignacio, *Inveready, Spain*
Gross Francis, *European Central Bank, Germany*
Jadav Divyesh, *IBM Research, USA*
McCullagh Adrian, *ODMOB Lawyers, Australia*
Pioli Moro Evandro, *British Telecommunication PLC (BT), UK*
Saraswat Vishal, *Bosch Global Software Technologies, India*
Sghaier Omar Ahmad, *Blockchain Foundry, Inc., Canada*
User Erol, *User Corporation and Blockchainarmy, Turkey*

Venue

B2C' 2026 Conference will take place in the modern 4-star AC Marriott Hotel Gran Canaria. It is located near Las Canteras beach (300 m) and Plaza Santa Catalina, in front of one of the famous harbor and a El Muelle Shopping Centre. The bus stop is just 50 metres away, and provides a quick and convenient way to get around the island. Set on the top floors of the hotel, AC Lounge Bar offers impressive views of the city. Address: C. Eduardo Benot, 3-5, 35007, Las Palmas de Gran Canaria, Spain.

Contact: b2c@sensorsportal.com

IFSA Publishing

Blockchain and Cryptocurrency

Published by IFSA Publishing, S. L., 2026
https://bc-ifsa-journal.com/

Sustainability Metrics and Disclosure Alignment in Blockchain Networks

[1] Hatem MABROUK, [2] Logan PAUL and David WILD

[1] Tecnológico de Monterrey, School of Business, Ave. Eugenio Garza Sada 2501 Sur, Col. Tecnológico, Monterrey, N.L., 64700, Mexico.

[2] Indiana University Luddy School of Informatics, Computing, and Engineering, 700 N Woodlawn Ave, Bloomington, IN 47408, United States.

Tel.: + 528111660781

E-mail: hatem.mabrouk@tec.mx

Received: 16 Dec. 2025 /Revised:5 Jan. 2026 /Accepted: 27 Feb. 2026 /Published:23 March 2026

Abstract: The rapid growth of blockchain networks has raised concerns about their environmental sustainability, yet existing assessments are fragmented and often rely on problematic metrics or unverifiable self-reported data. This article addresses these gaps by combining quantitative measures of annual energy consumption and carbon emissions with a qualitative assessment of sustainability initiatives. Using data from the Crypto Carbon Ratings Institute (CCRI) and corporate sustainability reports, we evaluate Bitcoin, Ethereum, Binance Chain, XRP, Cardano, Solana, Dogecoin, Polygon, and Polkadot, and benchmark their operational footprints against those of Visa and Mastercard as representatives of traditional payment systems. Results show wide variation in sustainability performance: Proof-of-Stake networks exhibit substantially lower energy use and carbon intensity than Proof-of-Work systems, while transparency and reporting practices differ significantly across projects. This research contributes to the sustainability discourse by providing one of the first systematic comparisons between blockchain networks and incumbent payment systems, identifying methodological challenges in sustainability accounting, and advancing policy recommendations for standardized reporting, third-party verification, and energy-efficient design.

Keywords: Blockchain sustainability, Eco-friendly cryptocurrencies, Proof of Stake (PoS), Energy consumption, Carbon footprint.

1. Introduction

This research addresses the critical need to evaluate the environmental impact of leading blockchain networks that underpin widely adopted cryptocurrencies. By examining a curated selection of major blockchain platforms – chosen for their technological relevance, market influence, and diversity in consensus mechanisms – this study provides a comprehensive assessment of their ecological footprints, sustainability initiatives, and potential for environmental improvement. The findings offer valuable insights for developers, investors, policymakers, and users seeking to make informed decisions about sustainable blockchain adoption and development.

2. Literature Review

The environmental impact of blockchain technology, particularly regarding energy consumption, has been a subject of significant scholarly attention. Alshahrani et al. [1] conducted a systematic literature review on two major sustainability issues in blockchain: power

consumption and scalability. Their research highlighted that the proof of work (PoW) consensus mechanism, used by Bitcoin and initially by Ethereum, requires intensive computational work resulting in high energy consumption. They noted that Bitcoin's annual energy usage exceeds 100 TWh, comparable to the electricity consumption of small nations. This finding aligns with data from the Cambridge Centre for Alternative Finance [2], which continues to track Bitcoin's substantial energy footprint through its Bitcoin Electricity Consumption Index.

Elsayed et al. [3] provide a comprehensive analysis of blockchain energy consumption challenges, emphasizing that, besides the technological benefits, significant concerns exist regarding the high energy consumption associated with blockchain operations. Their research explores alternative consensus techniques such as proof of stake (PoS) and delegated proof of stake (DPoS) as viable solutions to address the blockchain resource consumption challenge while maintaining the integrity and security of blockchain networks.

Previous comparative analyses, including early benchmark-based assessments of blockchain sustainability, have highlighted substantial disparities across consensus mechanisms and reporting practices [4].

2.1. Transition to More Sustainable Consensus Mechanisms

A significant development in blockchain sustainability has been Ethereum's transition to Proof of Stake (PoS) in 2022, which reduced its energy consumption by over 99 % [5]. This transition demonstrates the feasibility of achieving both scalability and sustainability in blockchain networks. Alshahrani et al. [1] emphasized that PoS-based networks like Cardano and Ethereum 2.0 validate transactions with minimal energy requirements, positioning them as significantly more sustainable alternatives to PoW systems.

2.2. Scalability Challenges and Solutions

Scalability represents another critical sustainability challenge for blockchain technology. Alshahrani et al. [1] identified that blockchain faces scalability problems when a large number of nodes and transactions are added to the network. Transaction throughput and transaction latency are the two major contentious performance metrics in blockchain systems, with most public blockchains yet to reach acceptable Quality of Service (QoS) standards.

Innovative approaches to address scalability include Layer 2 solutions, such as Ethereum's rollups and Polygon's scaling framework, which mitigate energy use through off-chain processing [1]. Additionally, Solana's combination of PoS with Proof of History (PoH) achieves high TPS with minimal energy use, demonstrating how architectural innovations can simultaneously address both scalability and sustainability concerns.

2.3. Critiques of 'Green' Claims and Offsets

A growing body of research cautions that headline claims of "carbon-neutral" or "renewable-powered" cryptocurrencies often rely on unverifiable surveys, renewable energy certificates (RECs), or carbon offsets, raising concerns of potential greenwashing. An MIT CEEPR analysis of publicly listed U.S. Bitcoin miners found that their carbon intensity broadly aligned with the U.S. grid average, contradicting industry survey claims of majority "sustainable" power and recommending dual-reporting under the GHG Protocol to avoid misuse of market-based instruments [8]. Similarly, the latest Cambridge Digital Mining Industry Report stresses the need for greater transparency in sustainability reporting and highlights persistent uncertainties in estimating the actual energy mix of PoW networks [10]. Beyond these power-mix debates, scholars emphasize that reliance on renewables, RECs, or offsets cannot resolve the structural energy intensity of Proof-of-Work systems and may obscure associated challenges such as electronic waste and rebound effects [9, 11]. Investigations into voluntary carbon credit markets – often used by blockchain projects to claim carbon neutrality – have also revealed widespread over-crediting and integrity concerns, underscoring the risks of greenwashing and the need for stronger standards [12].

2.4. Sustainability Disclosure, ESG Reporting, and Standardization Gaps

Despite the growing attention to the environmental impacts of blockchain systems, sustainability disclosure practices across cryptocurrency networks remain highly fragmented and inconsistent. Unlike traditional corporations, which increasingly operate under structured ESG reporting regimes, blockchain projects typically rely on voluntary disclosures issued by foundations, mining pools, or affiliated entities. These disclosures vary widely in scope, boundary definitions, and methodological assumptions, and are rarely subject to independent assurance. As a result, cross-chain comparisons are difficult and, in some cases, misleading, reinforcing information asymmetries for investors, regulators, and other stakeholders attempting to assess the environmental performance of blockchain networks [2, 5].

A central challenge underlying these disclosure inconsistencies is the accounting mismatch between operational energy use and reported sustainability claims. Many blockchain projects emphasize "carbon

neutrality" through renewable energy certificates (RECs) or voluntary carbon offsets, yet such instruments do not necessarily correspond to reductions in actual energy consumption or network-level emissions. Prior research has shown that reliance on market-based instruments can obscure the structural energy intensity of Proof-of-Work systems, introduce risks of double counting, and weaken the credibility of sustainability claims when underlying emissions remain high [9, 11]. Empirical investigations into voluntary carbon markets have further raised concerns about over-crediting and limited environmental additionality, underscoring the risk of greenwashing when offsets are used without transparent and auditable reporting frameworks [12].

These limitations highlight the need for standardized sustainability reporting approaches that enable comparability across heterogeneous blockchain architectures. In corporate contexts, established frameworks such as the GHG Protocol emphasize consistent system boundaries, separation of gross emissions from offsetting activities, and the use of verifiable, decision-useful metrics. More recently, regulatory initiatives such as the European Union's Corporate Sustainability Reporting Directive (CSRD) and the IFRS International Sustainability Standards Board (ISSB) standards have reinforced the importance of harmonized disclosure and third-party assurance in sustainability reporting. Comparable, widely adopted standards have yet to emerge for blockchain systems, constraining both academic analysis and evidence-based policymaking in the digital infrastructure domain [5, 13].

In response to these disclosure and standardization gaps, this study adopts a comparative framework grounded in operational, gross energy consumption and carbon emissions metrics derived from a consistent third-party data source, complemented by a qualitative assessment of publicly disclosed sustainability initiatives. By excluding offset-based neutrality claims from the quantitative analysis and prioritizing annual network-level energy use and emissions, the approach emphasizes transparency and cross-chain comparability. This design directly addresses the reporting limitations identified in the literature and provides a more robust basis for evaluating the sustainability implications of different consensus mechanisms and network designs, thereby motivating the methodological choices detailed in the following section.

3. Methodology

This study employs a mixed-methods approach to evaluate the sustainability of blockchain networks underlying a selection of widely adopted and technologically diverse cryptocurrencies. The research combines quantitative analysis of energy consumption and carbon emissions data with qualitative assessment of sustainability initiatives and technological innovations. This comprehensive approach enables a holistic evaluation of blockchain sustainability that goes beyond mere energy metrics.

3.1. Data Collection Methods

Data for this study were collected from multiple sources to ensure accuracy and comprehensiveness:

1. Academic Literature: Peer-reviewed research on blockchain sustainability, energy consumption, and environmental impact;
2. Industry Reports: Reports from specialized blockchain analytics firms, environmental organizations, and sustainability consultancies;
3. Public Disclosures: Sustainability reports, carbon footprint disclosures, and environmental initiatives announced by blockchain foundations and development teams;
4. Energy Consumption Indices: Data from established tracking platforms such as the Cambridge Bitcoin Electricity Consumption Index [2].

While multiple sources were reviewed to ensure comprehensiveness, the quantitative dataset is standardized on the CCRI 2025 indices [14] to maintain methodological consistency across blockchains. Other sources, including Digiconomist [6] and CBECI [2], were used only for triangulation and context; their published estimates (e.g., Bitcoin's annual electricity use ranging from ~120–190 TWh) fall within the same order of magnitude as CCRI [14], reinforcing confidence in the dataset.

3.2. Evaluation Framework (Selected Metrics)

Consensus Mechanism: The consensus mechanism is a critical determinant of blockchain sustainability, influencing both energy efficiency and decentralization. This study categorizes blockchains based on Proof of Work (PoW), Proof of Stake (PoS), Delegated Proof of Stake (DPoS), and hybrid models. PoW blockchains, such as Bitcoin, require intensive computational work, resulting in high energy consumption and carbon emissions. In contrast, PoS-based networks, such as Ethereum 2.0, Cardano, and Solana, validate transactions with minimal energy requirements.

Energy consumption: Energy use is a direct measure of a blockchain's environmental impact. PoW blockchains, particularly Bitcoin, consume exponentially more electricity than PoS or hybrid models [2]. Bitcoin's annual energy usage exceeds 100 TWh, comparable to the electricity consumption of small nations. Ethereum's transition to PoS in 2022 reduced its energy consumption by over 99 %. Post-Merge estimates vary by method; this article adopts CCRI's bottom-up value (≈0.0046 TWh/year)

[14], while acknowledging higher top-down figures reported elsewhere [5]. This discrepancy underscores the need for standardized methodologies in blockchain energy-use measurement, an issue increasingly emphasized in European blockchain policy discussions [7].

Carbon Footprint: We assess carbon impact using gross annual CO_2 emissions derived from energy consumption and grid-intensity factors. PoW networks tend to have higher emissions due to their dependence on energy-intensive mining operations, while PoS-based networks, like Cardano and Tezos, show substantially lower footprints due to their energy-efficient validation processes. In line with our methodology, we exclude per-transaction metrics (energy or CO_2) and do not net out offsetting claims. Reported neutrality programs (e.g., by the Cardano Foundation and Polygon Labs) are therefore excluded from the quantitative dataset to ensure comparability; instead, such sustainability initiatives (e.g., renewable procurement, offset programs) are reviewed qualitatively in Section 2.3.

Visa's 2023 Corporate Responsibility & Sustainability Report documents an annual energy consumption of 841000 gigajoules (GJ) in fiscal year 2023, equivalent to 0.234 TWh using the conversion factor (1 TWh = 3.6×10^6 GJ) [15]. The same report records 10600 metric tons of CO_2 equivalent (tCO_2e) in combined Scope 1 and Scope 2 emissions, corresponding to 0.011 $MtCO_2e$ when expressed in megatons [15]. In addition, Visa's 2023 CDP Climate Change Response (Section C6.5) provides a category-level disclosure of Scope 3 emissions, including 369200 tCO_2e from purchased goods and services. When all categories are summed, Visa's Scope 3 footprint totals approximately 409500 tCO_2e ($\approx$ 0.410 $MtCO_2e$) [16]. To ensure methodological consistency with blockchain data – which typically covers only direct energy-related Scope 1 and 2 emissions – our quantitative comparison tables use Visa's Scope 1+2 total (0.011 $MtCO_2e$), while the substantially larger Scope 3 emissions are discussed qualitatively.

Mastercard's 2024 Impact Report (covering calendar year 2023) discloses a total electricity consumption of 107023 megawatt-hours (MWh) [17]. Applying the conversion factor 1 MWh = 3.6 gigajoules (GJ), this equals 385283 GJ, which in turn corresponds to 0.107 TWh using the standard relationship 1 TWh = 3.6×10^6 GJ. The same report records 52054 metric tons of CO_2 equivalent (tCO_2e) in combined Scope 1 and Scope 2 emissions, which equals 0.052 $MtCO_2e$ when expressed in megatons [16]. Mastercard additionally provides Scope 3 data in its disclosures, with values that substantially exceed Scope 1+2 emissions. However, for comparability with blockchain networks – whose reported carbon footprints reflect only operational and electricity-related emissions – our quantitative tables include only Mastercard's Scope 1+2 total (0.052 $MtCO_2e$). The broader implications of Mastercard's Scope 3 footprint are addressed in the discussion section.

A key limitation of this study is the partial reliance on industry self-reports and disclosures from blockchain foundations, mining firms, and sustainability consortia. While these sources provide valuable data and insights, they are often voluntary, unaudited, and potentially influenced by reputational or market incentives. Prior critiques of sustainability claims in blockchain highlight risks of selective reporting, inconsistent methodologies, and even greenwashing [8, 10, 12]. To mitigate this, we cross-referenced self-reported data with independent indices and academic studies whenever possible, but acknowledge that industry disclosures may not always reflect full lifecycle impacts or provide standardized metrics.

For cross-chain comparability, annual energy and carbon figures are taken from the Crypto Carbon Ratings Institute (CCRI) indices (2025) using a consistent bottom-up methodology. Alternative estimates (e.g., Digiconomist; EU Blockchain Observatory summaries) may differ due to methodological assumptions; this study standardizes on CCRI for consistency.

4. Comparative Results

To enable fair comparison across blockchains of different sizes and transaction volumes, raw annual energy consumption (TWh) and annual carbon emissions (Mt CO_2) from CCRI [14] were used directly as inputs to the scoring framework.

The comprehensive analysis of selected leading cryptocurrencies reveals several key findings regarding blockchain sustainability:

1. Consensus Mechanism Dominance: Consensus mechanism emerges as the primary determinant of blockchain sustainability, with PoS and its variants demonstrating dramatically lower environmental impact than PoW systems;
2. Scalability-Sustainability Correlation: Higher transaction throughput generally correlates with better aggregate energy utilization at the network level;
3. Sustainability Initiative Effectiveness: Formal sustainability programs, particularly carbon offset initiatives, significantly improve overall environmental profiles, even for already efficient blockchains;
4. Transition Feasibility: Ethereum's successful transition from PoW to PoS demonstrates that even established blockchains can implement fundamental sustainability improvements without compromising security or decentralization;
5. Transparency Variations: Significant disparities exist in sustainability reporting and transparency across blockchain networks, with newer, purpose-built networks generally providing more comprehensive environmental disclosures.

Table 1. Annual energy consumption (TWh) and carbon emissions ($MtCO_2e$) of major blockchain networks and traditional payment systems.

Cryptocurrency/Payment System	Annual Energy Consumption (TWh)	Annual Carbon Emissions (Mt CO_2)
Bitcoin (BTC)	169.5	69.0
Ethereum (ETH)	0.0046	0.0014
XRP (XRP)	0.00049	0.00020
BNB Chain (BNB)	0.00021	0.00007
Solana (SOL)	0.018	0.0057
Cardano (ADA)	0.00048	0.00016
Polygon (MATIC)	0.00021	0.00005
Polkadot (DOT)	0.00103	0.00031
Dogecoin (DOGE)	9.0	3.7
Visa (2023)	0.234	0.011
Mastercard (2023)	0.107	0.052

Source: Blockchain data from CCRI (2025) [14]. Visa and Mastercard data from corporate sustainability reports [15-17].

Visa and Mastercard are included as non-blockchain benchmarks for energy and carbon comparison. They are not scored in Table 2 since consensus and scalability metrics are not applicable to centralized payment networks.

The overall sustainability ranking in Table 2 was derived using a weighted scoring framework. We assigned weights to five sustainability metrics based on their relative environmental importance: Consensus mechanism = 30 %; Energy consumption = 25 %; Carbon footprint = 20 %; Scalability = 15 %; and Sustainability efforts = 10 %. Each blockchain was given a score from 1 (worst) to 10 (best) for every metric. The bands were defined a priori: Consensus mechanism: PoW (1–2), hybrid (5–6), PoS (7–8), optimized PoS (9–10); Energy: >100 TWh/year = 1–2; ≤0.02 TWh/year = 9–10. Carbon: > 50 Mt CO_2/year = 1–2; ≤0.006 Mt CO_2/year = 9–10; Scalability: <10 TPS = 1–2, >10000 TPS = 9–10; Sustainability efforts: no initiatives = 1–2, carbon-negative verified = 9–10. Within each broad band (e.g., 9–10), the exact score was assigned by relative ranking across chains.

The overall score for each blockchain was computed as:

$$\begin{aligned} Overall\ Score = \\ = (Consensus \times 0.30) + \\ + (Energy \times 0.25) + (Carbon \times 0.20) + \\ + (Scalability \times 0.15) + \\ + (Sustainability\ Efforts \times 0.10) \end{aligned} \quad (1)$$

For example, Solana (PoS/PoH) scored: Consensus = 9, Energy = 10, Carbon = 9, Scalability = 10, Sustainability Efforts = 8. Therefore, Solana (PoS/PoH) yields: $(9 \times 0.30) + (10 \times 0.25) + (9 \times 0.20) + (10 \times 0.15) + (8 \times 0.10) = 9.30$.

This computation was repeated for each blockchain. The resulting weighted scores produced the ranking in Table 2.

As shown in Table 2, blockchains adopting energy-efficient consensus mechanisms and proactive sustainability strategies (e.g., Solana, Polygon, Cardano) achieve higher overall sustainability scores compared to Proof-of-Work networks with limited environmental measures (e.g., Bitcoin, Dogecoin) [14].

Table 2. Overall Sustainability Ranking of Top Cryptocurrencies (2025).

Cryptocurrency	Weighted Score	Rank
Solana (SOL)	9.30	1
Polygon (MATIC)	9.00	2
Cardano (ADA)	8.60	3
Polkadot (DOT)	8.10	4
XRP (XRP)	8.00	5
Ethereum (ETH)	7.90	6
Binance (BNB)	6.55	7
Dogecoin (DOGE)	3.65	8
Bitcoin (BTC)	1.25	9

Source: Calculated by authors using weighted score formula integrating all five-sustainability metrics.

Although Dogecoin is derived from Bitcoin's original codebase and employs a Proof-of-Work consensus mechanism, its sustainability ranking differs from Bitcoin's due to material differences in network scale and operational footprint. In particular, Dogecoin operates at a substantially lower aggregate hash rate and benefits from auxiliary merged mining with Litecoin, allowing portions of its security to be obtained without proportional increases in energy consumption. As a result, Dogecoin's annual electricity use and carbon emissions remain significantly lower than those of Bitcoin, despite sharing a similar consensus design. Empirical evidence shows that Bitcoin alone accounts for roughly two-thirds of the total energy consumed by mineable cryptocurrencies, with the remaining networks (including Dogecoin) representing a smaller share of overall consumption [18]. The ranking framework used in this study evaluates sustainability based on absolute operational impacts rather than protocol ancestry, and therefore distinguishes between Proof-of-Work networks according to their realized energy and emissions profiles. This explains why

Dogecoin, while still ranking poorly overall, scores marginally higher than Bitcoin in Table 2.

5. Discussion

Fig. 1 illustrates the magnitude of structural differences in energy consumption across blockchain networks and payment systems, reinforcing that sustainability outcomes are dominated by consensus design rather than incremental operational adjustments.

The findings of this research highlight the complex interplay between technological design choices, operational practices, and environmental outcomes in blockchain networks. The study found that decision-makers and practitioners often struggle to understand the underlying mechanisms of blockchain and its sustainability outcomes, creating barriers to adoption for **sustainability purposes.**

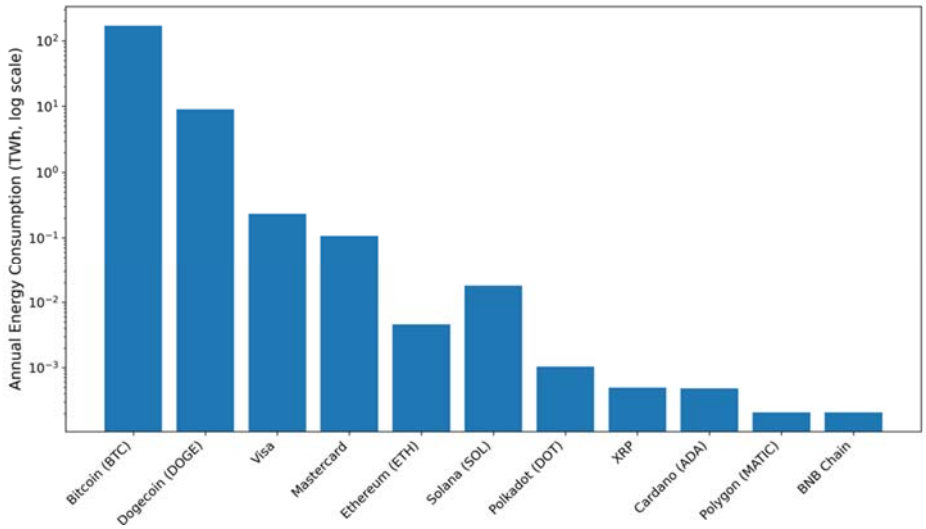

Fig. 1. Annual energy consumption of blockchain networks and traditional payment systems (log scale).

Visa and Mastercard were included in Table 1 as non-blockchain benchmarks to contextualize the relative scale of blockchain energy use and emissions. Their values provide a point of comparison with incumbent global payment systems but are not integrated into the composite sustainability ranking (Table 2), since blockchain-specific metrics such as consensus mechanism and scalability are not applicable to centralized networks.

5.1. Implications for Sustainability Disclosure and Regulatory Reporting (Signaling and Information Asymmetry)

The results indicate that sustainability performance across blockchain networks is primarily driven by structural design choices, most notably the underlying consensus mechanism, rather than by discretionary sustainability initiatives or offset-based commitments. Networks operating under Proof-of-Stake architectures consistently exhibit substantially lower operational energy consumption and carbon emissions than Proof-of-Work systems, regardless of scale or market capitalization. This pattern suggests that protocol-level design decisions exert a dominant influence on environmental outcomes, while post hoc mitigation strategies play a secondary role. Consequently, observed differences in sustainability performance reflect inherent architectural characteristics rather than short-term managerial or reputational interventions.

While sustainability disclosures are increasingly used by blockchain projects to signal environmental responsibility, the extent to which such disclosures align with observed operational environmental impacts remains unclear. To illustrate this alignment–misalignment dynamic, Table 3 juxtaposes publicly communicated sustainability claims with independently estimated operational energy use and emissions derived from third-party data sources. This comparison highlights how differences in disclosure practices, boundary definitions, and the treatment of offsets affect the credibility and comparability of sustainability signals across networks.

Table 3 reveals substantial heterogeneity in the alignment between sustainability disclosures and observed environmental impacts across blockchain networks. Projects with low operational footprints but fragmented or narrative-driven disclosures exhibit moderate alignment, whereas networks with high absolute energy consumption and limited standardized reporting display pronounced disclosure gaps. Conversely, platforms that explicitly communicate operational reductions – particularly following structural changes such as consensus mechanism transitions – show higher directional alignment, even when absolute impacts remain non-negligible. These

patterns reinforce the argument that disclosure credibility depends not only on performance outcomes but also on the standardization, transparency, and auditability of reported sustainability information.

Table 3. Alignment between blockchain sustainability disclosures and observed operational environmental impacts.

Network	Project sustainability disclosure (what they claim)	Observed operational impact (CCRI via MiCA disclosure)	Alignment
Bitcoin (BTC)	No protocol-level standardized sustainability disclosure is provided by the network itself (typical info is ecosystem- or miner-driven rather than a single issuer disclosure) [19].	Energy consumption: **165002667937 kWh/yr**; Scope 2 emissions: **67597922605 tCO2e/yr [19].**	**Low alignment / disclosure gap:** very large observed footprint + no unified project disclosure.
Dogecoin (DOGE)	No unified foundation-level sustainability disclosure appears as a standard reporting artifact comparable to ESG/CSRD/ISSB-style reporting [20].	Energy consumption: **8562992922 kWh/yr**; Scope 2 emissions: **3626830.160 tCO2e/yr [20, 18].**	**Low alignment / disclosure gap:** PoW footprint is material; limited standardized disclosure.
Ethereum (ETH)	Publicly states PoS transition reduced energy use by ~**99.95 %** (Merge) and provides an energy consumption page (PoS-era) [2829].	Energy consumption: **4177274.115 kWh/yr**; Scope 2 emissions: **1293.391 tCO2e/yr [21]**.	**High alignment (directionally):** disclosure emphasizes sharp reduction; observed estimates are orders of magnitude below PoW systems.
Polygon (PoS chain / "Polygon")	Announces a plan to become **carbon negative** (via offset/credits strategy + funding commitments) [31].	Energy consumption (MiCA): **183582.888 kWh/yr**. Independent CCRI update estimates Polygon post-Merge total annualized emissions **~55 tCO2e/yr [25, 32].**	**Partial alignment:** observed operational impacts are low, but the "carbon negative" claim relies on offset framing; comparability depends on separating gross emissions vs offsets.
Solana (SOL)	Publishes "Energy Impact / Energy Use" reports and discusses offset instruments (carbon credits, biodiversity credits) [30].	MiCA disclosure indicates renewable share and Scope 2 emissions (example snapshot): Scope 2 emissions: **4856.208 tCO2e/yr** (and renewable share reported) [26].	**Mixed/partial alignment:** disclosure is active and detailed, but the inclusion of offsets can confuse "impact" vs "net claims" unless reported separately.
Cardano (ADA)	No single, universally-used sustainability reporting artifact comparable to corporate ESG standards is consistently used across the ecosystem (disclosures tend to be scattered across entities) [22].	Energy consumption: **497956.700 kWh/yr**; Scope 2 emissions: **172.337 tCO2e/yr [22].**	**Moderate alignment:** observed footprint is low; the main issue is **standardization/assurance** rather than performance.
Polkadot (DOT)	Project-facing claims often emphasize efficiency, but formal standardized sustainability disclosure is not always presented in the same way across venues [23].	Energy consumption: **1028091.344 kWh/yr**; Scope 2 emissions: **311.963 tCO2e/yr [23].**	**Moderate alignment:** observed footprint is low; the remaining gap is disclosure consistency and auditability.
XRPL (XRP Ledger)	Sustainability statements are typically communicated through Ripple/corporate sustainability positioning and goals, not always a protocol-level standardized reporting package [24].	Energy consumption: **476747.129 kWh/yr**; Scope 2 emissions: **189.302 tCO2e/yr [24].**	**Moderate alignment:** low footprint, but disclosure is more "narrative + targets" than standardized operational reporting.
BNB Chain (BNB)	Often referenced via exchange/corporate sustainability communications; protocol-level standardized disclosures are inconsistent across channels [27].	Energy consumption: **269175.594 kWh/yr [27].**	**Moderate alignment (limited):** operational estimate is low, but disclosure standardization is thin (and emissions fields may not be consistently published in the same format).

Source: Markets in Crypto-Assets Regulation (MiCA) sustainability disclosures issued by Coinbase Luxembourg S.A. [19-27]; Ethereum Foundation disclosures [28, 29]; Solana Foundation energy report [30]; Polygon sustainability disclosures and CCRI assessment [31, 32]; comparative energy context from Gallersdörfer et al. [18].

These findings carry direct implications for regulatory and policy design. A minimum sustainability disclosure baseline for public blockchain networks should require reporting of: (i) annual network electricity consumption (TWh) and estimation methodology; (ii) annual gross carbon emissions

($MtCO_2e$) and underlying grid-intensity assumptions; (iii) clear system boundary definitions and treatment of uncertainty; (iv) energy mix claims supported by auditable evidence; (v) offsets and renewable energy certificates disclosed separately from gross emissions (gross versus net reporting); (vi) governance responsibility for sustainability reporting, including reporting frequency and accountability; and (vii) independent third-party assurance of key metrics. Such requirements directly address greenwashing risks identified in the literature and would improve transparency and comparability across networks, enabling regulators and market participants to evaluate sustainability performance using decision-useful information [10, 12].

Looking forward, the convergence of sustainability reporting standards in corporate contexts suggests a plausible trajectory toward analogous disclosure expectations for blockchain infrastructure, particularly as crypto-assets increasingly intersect with regulated financial markets. As institutional investors, exchanges, and custodians integrate sustainability considerations into screening and risk management processes, standardized and assured disclosure may become a prerequisite for market access and legitimacy. In this context, the development of harmonized blockchain sustainability reporting frameworks – aligned with emerging international disclosure standards – would enable longitudinal monitoring of technological transitions and their environmental effects, while supporting evidence-based policymaking and future research on sustainable digital infrastructure.

6. Conclusion

This comparative analysis of blockchain sustainability across selected leading cryptocurrency networks reveals significant disparities in environmental impact, with the consensus mechanism emerging as the primary determinant of overall sustainability. By prioritizing sustainability in both technical design and operational practices, the blockchain industry can minimize its environmental footprint while continuing to deliver transformative benefits across numerous sectors. The best practices identified in this research provide a roadmap for achieving this balance between innovation and environmental responsibility.

Future research should explore emerging consensus mechanisms and their environmental implications, standardized approaches to blockchain sustainability measurement, and the potential integration of blockchain with complementary technologies such as renewable energy systems and carbon markets. Additionally, longitudinal studies tracking the evolution of blockchain sustainability over time would provide valuable insights into the effectiveness of various improvement strategies.

This study provides an initial benchmark by comparing blockchain networks with the operational footprints of major global payment systems (Visa and Mastercard), included in Table 1. Future research should extend these comparisons to cover the full lifecycle energy costs of credit card transactions, debit card transactions, and physical fiat transactions, thereby providing a more comprehensive benchmark for evaluating blockchain systems against traditional payment methods. Based on these findings, regulators and industry stakeholders should prioritize the development of standardized sustainability reporting frameworks for blockchain projects, aligned with established initiatives such as the GHG Protocol and emerging EU analyses [7, 8]. Independent third-party verification of energy use and carbon offset claims is essential to mitigate greenwashing risks [10, 12]. In parallel, policymakers should incentivize transitions toward energy-efficient consensus mechanisms (e.g., PoS) and encourage adoption of credible renewable energy procurement. Establishing minimum disclosure requirements – covering energy consumption, carbon footprint, governance, and sustainability initiatives – would enhance transparency, comparability, and trust across blockchain networks, thereby supporting responsible investment and long-term industry legitimacy.

References

[1]. H. Alshahrani, M. Islam, S. S. Ullah, A. Al-Reshan, et al., Sustainability in blockchain: A systematic literature review on scalability and power consumption issues, *Energies*, Vol. 16, Issue 3, 2023, 1510.

[2]. Cambridge Centre for Alternative Finance, Cambridge Bitcoin electricity consumption index, https://ccaf.io/cbnsi/cbeci

[3]. A. Elsayed, A. Erdoğdu, S. Elsayed, Blockchain technology and energy consumption: A comprehensive analysis and sustainability considerations, in *Proceedings of the International Student Conference on Business, Education, Economics, Accounting, and Management (ISC-BEAM'24)*, 2024.

[4]. H. Mabrouk, L. Paul, D. Wild, Comparative sustainability analysis of major cryptocurrency blockchains, in *Proceedings of the 4th Blockchain and Cryptocurrency Conference (B2C'25)*, 2025, pp. 92-96.

[5]. Ethereum Merge Trend Report, European Union Blockchain Observatory and Forum, 2023, https://blockchain-observatory.ec.europa.eu/document/download/3f78c885-d14e-47cb-b183-f22ef529a258_en?filename=EUBOF3.0_Ethereum_Merge_Trend_Report_final.pdf

[6]. Digiconomist, Dogecoin energy consumption index, https://digiconomist.net/dogecoin-energy-consumption

[7]. European Union Blockchain Observatory and Forum, https://blockchain-observatory.ec.europa.eu

[8]. Climate impacts of Bitcoin mining in the U.S., CEEPR Working Paper 2023-11, MIT Center for Energy and Environmental Policy Research, 2023, https://ceepr.mit.edu/publications/working-paper/climate-impacts-of-bitcoin-mining-in-the-u-s/

[9]. A. de Vries, Renewable energy will not solve Bitcoin's sustainability problem, *Joule*, Vol. 3, Issue 4, 2019, pp. 893-898.

[10]. Cambridge Centre for Alternative Finance, Cambridge digital mining industry report, https://www.cambridge.org/engage/api-gateway/ccaf/assets/orp/resource/item/66cd1c6a8e7f563e64a64627/original/cambridge-digital-mining-industry-2025.pdf
[11]. A. de Vries, U. Gallersdörfer, L. Klaaßen, C. Stoll, Revisiting Bitcoin's carbon footprint, *Joule*, Vol. 6, Issue 3, 2022, pp. 498-502.
[12]. P. Greenfield, Revealed: More than 90 % of rainforest carbon offsets by biggest certifier are worthless, analysis shows, *The Guardian*, 18 January 2023.
[13]. IFRS S1 General Requirements for Disclosure of Sustainability-related Financial Information, *IFRS Foundation*, 2023.
[14]. Crypto Carbon Ratings Institute (CCRI), Crypto sustainability indices, https://indices.carbon-ratings.com/
[15]. Visa, 2023 Corporate responsibility & sustainability report, https://corporate.visa.com/content/dam/VCOM/regional/na/us/about-visa/documents/2023-corporate-responsibility-sustainability-report.pdf
[16]. Visa, CDP climate change response 2023, https://corporate.visa.com/content/dam/VCOM/regional/na/us/about-visa/esg/2023-visa-cdp-response.pdf
[17]. Mastercard, 2024 Mastercard impact report, https://www.mastercard.com/content/dam/mccom/shared/for-the-world/corporate-impact/pdfs/mastercard-2024-impact-report.pdf
[18]. U. Gallersdörfer, L. Klaaßen, C. Stoll, Energy consumption of cryptocurrencies beyond Bitcoin, *Joule*, Vol. 4, Issue 9, 2020, pp. 1843-1846.
[19]. Coinbase Luxembourg S.A., Mandatory information on principal adverse impacts on the climate and other environment-related adverse impacts of the consensus mechanism – Bitcoin (BTC), MiCA sustainability disclosure, 2025, https://static-assets.coinbase.com/mica/btc.pdf
[20]. Coinbase Luxembourg S.A., Mandatory information on principal adverse impacts on the climate and other environment-related adverse impacts of the consensus mechanism – Dogecoin (DOGE), MiCA sustainability disclosure, 2025, https://static-assets.coinbase.com/mica/doge.pdf
[21]. Coinbase Luxembourg S.A., Mandatory information on principal adverse impacts on the climate and other environment-related adverse impacts of the consensus mechanism – Ethereum, MiCA sustainability disclosure, 2025, https://static-assets.coinbase.com/mica/eth2.pdf
[22]. Coinbase Luxembourg S.A., Mandatory information on principal adverse impacts on the climate and other environment-related adverse impacts of the consensus mechanism – Cardano (ADA), MiCA sustainability disclosure, 2025, https://static-assets.coinbase.com/mica/ada.pdf
[23]. Coinbase Luxembourg S.A., Mandatory information on principal adverse impacts on the climate and other environment-related adverse impacts of the consensus mechanism – Polkadot (DOT), MiCA sustainability disclosure, 2025, https://static-assets.coinbase.com/mica/dot.pdf
[24]. Coinbase Luxembourg S.A., Mandatory information on principal adverse impacts on the climate and other environment-related adverse impacts of the consensus mechanism – XRPL, MiCA sustainability disclosure, 2025, https://static-assets.coinbase.com/mica/xrp.pdf
[25]. Coinbase Luxembourg S.A., Mandatory information on principal adverse impacts on the climate and other environment-related adverse impacts of the consensus mechanism – Polygon, MiCA sustainability disclosure, 2025, https://static-assets.coinbase.com/mica/pol.pdf
[26]. Coinbase Luxembourg S.A., Mandatory information on principal adverse impacts on the climate and other environment-related adverse impacts of the consensus mechanism – Solana, MiCA sustainability disclosure, 2025, https://static-assets.coinbase.com/mica/sol.pdf
[27]. Coinbase Luxembourg S.A., Mandatory information on principal adverse impacts on the climate and other environment-related adverse impacts of the consensus mechanism – BNB Chain, MiCA sustainability disclosure, 2025, https://static-assets.coinbase.com/mica/bnb.pdf
[28]. Ethereum Foundation, Ethereum energy consumption, https://ethereum.org/energy-consumption/
[29]. Ethereum Foundation, The Merge, https://ethereum.org/roadmap/merge/
[30]. Solana Foundation, Energy impact report: September 2024, https://solana.com/news/energy-use-report-september-2024
[31]. Polygon Labs, Polygon is going carbon negative in 2022 with a $20 million climate pledge, https://polygon.technology/blog/polygon-is-going-carbon-negative-in-2022-with-a-20-million-pledge
[32]. Crypto Carbon Ratings Institute (CCRI), The energy efficiency and carbon footprint of the Polygon blockchain – Update 2022, https://carbon-ratings.com/dl/polygon-update-2022

IFSA Publishing

Blockchain and Cryptocurrency

Published by IFSA Publishing, S. L., 2026
https://bc-ifsa-journal.com/

Distributed Ledger Technology and Economic Resilience: Strengthening Central Banks, Commercial Banks, and Land Registries

Ian STALEY
W. E. Deming School of Business, William Howard Taft University, Denver, United States
Tel.: +1 (206) 303-7903
E-mail: ian.t.staley@gmail.com

Received: 18 Sept. 2025 /Revised: 29 Oct. 2025 /Accepted: 28 Nov. 2025 /Published: 23 Mar. 2026

Abstract: This research article examines how distributed ledger technology (DLT) can enhance modern-day economies and the mechanisms that enable this emerging technology to sustain them in the long term. The mission of this study is to educate a diverse group of economic leaders, encompassing government agencies and private companies, about DLT and its potential to shape the future. This study analyzes secondary qualitative data to show that DLT can enhance and sustain economies in multiple ways, specifically through the three pillars of modern-day economies: central banks, commercial banks, and land registry systems. More specifically, the architectural mechanisms of DLT reduce moral hazard arising from centralized economic authorities, increase the efficiency of financial services and money movements, and lower the costs of financial services that can be passed on to consumers. Further benefits include the creation of new jobs, new industries, a new asset class, and renewed industries through the adoption of this new infrastructure, thereby expanding markets by building strong foundations for economies to grow through immutable land records, and building trustless networks worldwide.

Keywords: Banking, Blockchain, Distributed ledger technology, Economies, Land registry, Web3.

1. Introduction

The development and stability of economies have always been dependent on technological innovations, which determine the way a commercial bank, land registries and central banks among other institutions conduct their activities and relate with the society. The Bank of Sweden and the Bank of England were the first central banks that were developed in the 17th century and this became a historical moment because currencies and credit were stabilized. Subsequent innovations, like the introduction of the computer systems into the middle of the 20th century and the introduction of electronic trading systems made monetary policy and financial oversight more effective and more precise [1]. Other technologies like the telegraph, ATMs, and internet banking in the commercial banking industry contributed to the increased dissemination of financial services and efficiency. Concurrently, land registry systems are gone computerized and they rely on the GIS which provides transparency, land tenure assurances and proper land-use planning [2]. All these changes have gone a long way in increasing the global GDP, which is an important indicator of economic well being and economic growth.

This study article seeks to explore and explain how distributed ledger technology (DLT) can be used to strengthen and advance the economy based on this record of innovation. It will inform economic leaders in both the government and private sector about the possibilities of DLT and offer them practical guidance to make informed strategic decisions in the current digital transformation. The qualitative secondary

studies in the research article have been assessed by evaluating the phenomenon of the case in the three key pillars of contemporary economies, that is, central banks, commercial banks, and land registries.

The most topical consequences stem from the need to discuss DLT as a solution to the system's shortcomings. This was demonstrated by the banking collapses of Silicon Valley Bank, Signature Bank, and Credit Suisse, among others, driven by liquidity crises, insolvency, and investor confidence. These cases show that financial institutions remain at risk despite technological advancements [3-5]. It is here that the concept of DLT as a possible solution to risk issues can be interpreted, as the parameter possesses the characteristics of "decentralization", "immutability", and "real-time auditability" [6]. By enabling on-demand liquidity, simplifying settlements, and reducing counterparty risk, DLT can enable banks to implement tokenization and smart contracts. It is decentralized, making it less prone to a single point of failure and reducing the likelihood of systemic contagion during an economic collapse [7].

2. Methodology

The approach used in this research article is a qualitative research design, which employs a secondary research methodology to examine the relevance of distributed ledger technology (DLT) in enhancing and sustaining economies. It targets three economic infrastructure pillars because past research has shown they are highly valuable for economic stability and development. The specific aim of this methodological procedure is to inform economic leaders, both in government and in business, without delving into too many technicalities. A set of big questions was directed at the study to examine the significance of central banks, commercial banks, and land record systems to economic development. The history of Web1, Web2, and Web3, the various types of DLT, their advantages and disadvantages, and the exact value that DLT can offer compared to state systems in the present and future were also questions. Finally, the following questions were developed based on the project's mission: to inform stakeholders about the strategic potential of DLT in modern economies.

The study covers the period from 2000 to the present, with a special focus on real-world applications and case studies that demonstrate how DLT is implemented in economic systems. The data were collected through an extensive literature review, including articles from academia, books, government reports, trade association publications, and case studies. These sources were considered reliable in terms of relevance, credibility, and accuracy. The collected data were summarized, encoded, and synthesized to identify common themes and patterns. All three pillars of the economy are interrelated in terms of technological progress and their mutual impact, providing valuable insights.

3. Conceptual Model

In this work, we develop the conceptual model of Distributed Ledger-Economic Resilience Model (DLT-ERM), as a novel framework that will support the adaptability of modern economies through distributed ledger technology (DLT) and its long-term sustainability. This model is based on the integrated understanding of the Institutional Economics Theory [9] and Resilience Theory [10], which provides a novel interpretive flank to explaining how decentralization, transparency, and immutability as the main characteristics of DLT change the institutional trust and strengthen the economic systems against shocks. Institutional economics views performance and the strength of an economy as a result of the quality and credibility of the institutions. Once such credibility is destroyed by asymmetry of information, corruption, and inefficiency, the economic systems are left vulnerable to collapse and crisis. The solutions of DLT to these structural weaknesses include decentralizing trust, applying rules in an algorithmic manner, and minimizing reliance on intermediaries which tend to add some degree of opacities. At the same time, resilience theory considers economies as complex adaptive systems that can withstand shocks, rearrange, and retain central operations during difficult times. In this respect, DLT is a preventive and adaptive tool, which helps to be more transparent, allows auditing in real time and promotes inclusivity in economic relations.

The conceptual model of the DLT-ERM presents the idea of distributed ledger technology as a technological facilitator that enhances economic resilience by ensuring three interrelated pillars of institutions, namely central banks, commercial banks, and land registries. In central banks, the DLT encourages transparency in the monetary activities, supports the issuance and management of the central bank digital currency (CBDCs), and improves the liquidity management. It enhances efficiency in transactions in commercial banks, reduces the operational risk brought about by automation and increases financial inclusion by lowering the transaction costs and the obstacles to entry. Land registries, which are usually prone to corruption and lack of efficiency, are also advantaged by the uncorrupted and transparent nature of DLT, which means that property ownership records are secure, and people can use land as verifiable security to access credit. These pillars have the collective potential to create an institutional base of an economically resilient, transparent, efficient, and inclusive economy.

The processes of connecting DLT to economic resilience can be broadened in four main processes. To begin with, transparency and accountability are achieved through immutable ledgers, which help reduce fraud, diminish information asymmetry, and enhance regulatory oversight. Second, automated settlements and decentralized consensus are used to achieve operational efficiency and stability. Third,

financial inclusion and innovation are increased because, through the use of DLT, digital identities and peer-to-peer finance are possible, which allows the more marginalized groups to be empowered and increases their participation in the formal economy. Lastly, sustainability and adaptability can be seen as energy-efficient consensus mechanisms and decentralized architectures increase the environmental responsibility and minimise the weaknesses of centralised control.

These interrelations can be traced in conceptual form in the Fig. 1 below of the DLT-ERM model which indicates the logical path of technological characteristics into institutional transformation and finally to increased economic stability:

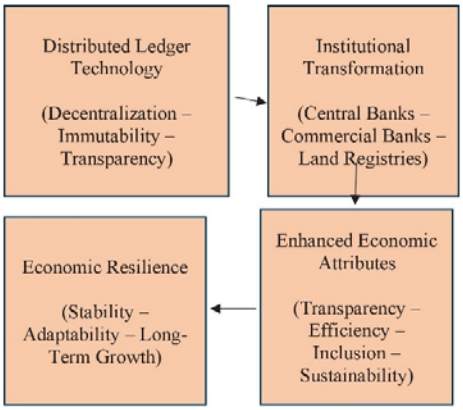

Fig. 1. DLT–ERM model.

3. U.S. Regulation & Policy

The emergence of DLT has presented fresh opportunities for efficiency, transparency, and security in the U.S. economy, but it has also generated urgent regulatory and policy imperatives. In contrast to previous innovations in the financial and technological fields, DLT directly overlaps with the three pillars of the U.S. economy, and its regulation is critical to the stability of the economy [11]. The regulation in this area should balance encouraging innovation with reducing systemic risk. Among the most significant regulatory issues in the U.S. is the central bank digital currency (CBDC). The Federal Reserve has been prudent in its approach, making it clear that maintaining financial stability and safeguarding consumer privacy are more important than pursuing the prospects of a digital dollar. There has been an ongoing debate over how a CBDC can coexist with commercial banks, as central banks risk disintermediating traditional financial institutions if consumers go straight to the Fed to store their money [12]. According to [13] the policymakers thus have a narrow walk to walk as CBDCs may improve efficiency and transparency. However, they should be structured so as not to undermine the role that commercial banks play in creating credit, which is central to economic development.

U.S. regulators are struggling to adopt blockchain-based solutions within the existing commercial banking system. The Commodity Futures Trading Commission (CFTC) and the Securities and Exchange Commission (SEC) have been at the forefront when deciding whether particular digital assets are treated as securities or commodities. This difference influences the regulation of such innovations as tokenized deposits, stablecoins, and decentralized finance products [13]. Poor classification has led to regulatory confusion, discouraging innovation and investment in the U.S. market [14]. For example, banks seeking to issue tokenized assets must comply with fragmented rules across multiple regulators, including the Office of the Comptroller of the Currency (OCC) and the Federal Deposit Insurance Corporation (FDIC) [15]. Such ambiguity underscores the need for a unified regulatory framework.

Another area where the US regulation will play a crucial role is property rights and land registry systems. Introducing DLT into property records could transform the efficiency and transparency of property transactions. Nevertheless, the U.S. has land-use and property laws that are predominantly administered at the state and county levels, resulting in an inconsistent system. Whereas certain states, including Vermont and Wyoming, have already enacted laws recognizing blockchain-based records and smart contracts, other states have not [16]. This quilted system prevents the scalability of DLT-based land registries. It explains why federal guidance that coordinates the state-level activities without interfering with the local jurisdiction is necessary. In addition to the sector-specific considerations, U.S. policymakers face broader global challenges in financial stability, consumer protection, and national security. For example, even though DLT can ensure transparency and minimize fraud, it is also associated with money laundering, terrorist financing, and other illegal activities [17]. The Financial Crimes Enforcement Network (FinCEN) has imposed anti-money laundering (AML) and know-your-customer (KYC) regulations on some digital asset providers, and implementing these regulations in the context of decentralized systems is not easy. Regulators will have to identify solutions to support decentralized innovation by expanding accountability without suffocating it [18].

Another burning concern is sustainability. Blockchain proof-of-work technology has been criticized for its high resource utilization (energy consumption), leading some DLT systems to switch to more energy-efficient consensus mechanisms [14]. Environmental policymakers are also associating financial innovation with climate goals, particularly when the U.S. seeks to fulfill its environmental pledges [14]. According to [19] the regulation of the future may then require that systems based on DLT be oriented towards sustainability, such that innovation advances in tandem with long-term environmental policy.

Geopolitical rivalry also influences the U.S. policy. As China advances with its digital yuan and the

European Union develops a digital euro, the U.S. is under increasing pressure to expedite its digital asset policy to maintain its status as a global reserve currency [20]. Any slowness or disjointed regulatory path might undermine U.S. competitiveness in the global financial system, and hastily implementing such a regulatory system without safeguards can put the economy at risk of systemic failures [21]. This geopolitical aspect supports the necessity to develop a consistent national approach to digital assets.

There are already efforts in place to investigate digital assets. The 2022 Executive Order of the Biden Administration on Ensuring Responsible Development of Digital Assets was the first coordinated federal approach to digital assets. It has instructed agencies to investigate consumer security, economic stability, unlawful finance, U.S. competitiveness, and climate hazards [22]. Although the order was an indication of commitment to innovation, its principles are still being translated into regulatory form. According to [23] industry stakeholders are still demanding transparency, especially regarding securities classification and the issuance of stablecoins.

4. Global Comparisons

The regulation and adoption of distributed ledger technology (DLT) vary widely across global economies, reflecting differences in governance structures, financial infrastructures, and policy priorities. Table 1 below summarizes how leading regions and countries are approaching DLT, highlighting their strategies, opportunities, and challenges:

Table 1. Global Comparisons of Regulations.

Region / Country	Regulatory Approach	Key Initiatives / Features	Opportunities	Challenges / Trade-offs	Citations
European Union (EU)	Coordinated, harmonized regulation	Markets in Crypto-Assets Regulation (MiCA, 2023); ECB exploring digital euro	Reduces regulatory uncertainty; promotes cross-border payments; supports financial inclusion	Balancing innovation with monetary sovereignty	[23, 24]
United States (U.S.)	Fragmented across multiple agencies (SEC, CFTC, OCC, FDIC)	Ongoing debate on CBDC; Biden's 2022 Executive Order on digital assets	Strong innovation ecosystem; potential for global competitiveness	Regulatory uncertainty; risk of disintermediation of commercial banks; fragmented oversight	[25, 26]
China	Centralized, state-led model	Piloting digital yuan; bans cryptocurrencies but promotes blockchain in supply chain, trade finance, and records.	Rapid adoption; integration into retail payments; strong state control	Market freedom curtailed; limited private innovation	[27]
Singapore	Innovation-friendly, sandbox-based	Monetary Authority sandbox programs; AML compliance integration	Controlled fintech experimentation; global fintech hub reputation	Balancing openness with financial crime risks	[28]
Switzerland	Clear, pro-innovation legal framework	"Crypto Valley" ecosystem; clarity on tokenized assets & smart contracts	Attracts startups, investors, and global talent	Small market scale limits global impact	[29]
Developing Nations (e.g., Georgia, Rwanda)	Practical adoption for governance	Blockchain-based land registries to improve transparency & reduce corruption	Supports economic development; increases trust in weak institutions	Limited infrastructure; reliance on external expertise	[30]

5. Market Structure & Infrastructure

The findings reveal that DLT has the potential to transform the market structure and financial infrastructure of modern economies by enhancing efficiency, transparency, and security. For commercial banks, results highlight how DLT supports cost reduction, faster cross-border payments, and improved compliance processes. Case examples such as the Spunta Project in Italy show that DLT can automate reconciliation, thereby reduce administrative overheads and improve liquidity management [31]. The adoption of DLT-based solutions could restructure banking competition by enabling new entrants, fostering decentralized finance (DeFi) models, and creating new asset classes [32, 33]. These innovations contribute to economic growth by stimulating private-sector activity and generating new employment opportunities in the digital asset industry.

The article also highlights the transformative potential of DLT in land record systems, where corruption, inefficiency, and disputes persist in many economies [34]. By providing transparent, tamper-proof registries, DLT reduces property fraud,

accelerates transaction times, and attracts foreign investment by securing property rights. In regions with weak institutional infrastructure, the introduction of blockchain-based land records could significantly reduce capital deadlock, enabling individuals to leverage property for credit and business development [35]. In terms of broader market infrastructure, Web3 technologies illustrate how decentralized platforms reshape the economic landscape. Unlike Web1's static information-sharing and Web2's platform-driven innovation, Web3 integrates DLT to create decentralized systems that remove intermediaries and promote financial inclusion [36]. This market structure empowers unbanked populations to participate in formal economies while fostering competition and innovation among financial service providers.

6. Macro Channels

A combination of DLT and economies comes with some macro channels that affect efficiency, transparency, and security of financial and non-financial systems. These macro channels operate within central banks, commercial banks, and land record systems, ultimately shaping the performance and resilience of economies. Monetary policy transmission through central banks is one of the most important macro channels. DLT provides central banks with the means to conceptualize and execute Central Bank Digital Currencies (CBDCs), thereby enhancing the efficiency of local and offshore transactions. Unlike traditional systems, which are characterized by delays and data asymmetry, DLT offers real-time auditability and transparency, thereby enhancing monetary policy control [37]. For example, a wholesale CBDC could enable central banks to better control liquidity while preserving financial stability [22]. The microchannel leads to greater policy responsiveness, reduced systemic risk, and enhanced economic resilience.

The other microchannel is financial intermediation using commercial banks. With the implementation of DLT, commercial banks will be able to simplify payment systems, automate compliance with smart contracts, and cut operational expenses. An example of how distributed systems can streamline transaction reconciliation and ensure its safety is Spunta, a project in Italy based on Corda DLT [38]. Banks can, through this microchannel, reduce transaction fees, enhance customer service, and encourage greater economic activity. In addition, DLT opens the door to new asset classes, e.g., tokenized securities and digital lending products, which support financial innovation and employment.

The third microchannel is the modernization of the property and land registry. Weak or corrupt land records in most developing economies is a hindrance to investment and a breeding ground for conflict. Blockchain-based land registries can provide tamper-proof ownership documents, reduce fraud, and enhance transparency [20]. This microchannel will strengthen investment safety, the inflow of foreign capital, and real estate and infrastructure development. Property security, in turn, enables individuals and businesses to use land as collateral, thereby promoting credit growth and economic development [40]. Another microchannel of global financial integration is DLT. It is also possible that DLT will help to reduce the cost and speed of cross-border transactions by removing intermediaries and verifying in a decentralized manner [41]. This helps avoid international trade, expand access to financial services for the unbanked, and democratize the financial system [42]. The net effect is enhanced inclusiveness and financial sustainability, especially for emerging economies joining the world markets.

Finally, there is a microchannel for regulatory control and sustainability: DLT. This reduces information asymmetries through its transparency and immutability, enabling regulators to monitor financial institutions in real time [42]. In addition, new consensus mechanisms, such as proof-of-stake or proof-of-time, are being developed to reduce carbon emissions and bring DLT closer to environmental sustainability goals [43]. This microchannel regulatory tool effectively helps resolve the moral hazard problem of one of the key financial institutions in the system, while also promoting the creation of more environmentally friendly financial products.

7. Risks & Offsets

DLT also poses several risks when used. One of the primary risks includes technological risks. Public DLTs, in general, and blockchain, in particular, exhibit scalability problems, consume significant energy, and process transactions slowly [12,30]. These issues can undermine performance, particularly in large-volume financial markets. To address this, alternative governance models, such as hybrid and consortium models, offer a more scalable, energy-efficient structure and are more specific in their institutional applications.

The other risk is regulatory uncertainty. Lack of adequate legal frameworks for data ownership, digital identity, and international transactions may pose a challenge to adoption. Differences in international regulations can collapse market systems, and institutional lag in adoption [45]. It has regulatory sandboxes and cross-jurisdictional cooperation, which provide offsets, enabling controlled experimentation and ensuring compliance. Cybersecurity and privacy are also dangerous. Although remarkably safe, the DLT has had integration vulnerabilities, such as wallets and exchanges, which have been the focus of high-profile fraud [19]. The offsets are regarded as the more stringent cryptography standards, multi-signature authentication, and real-time regulatory control nodes (Von Solms, 2021).

Traditionally, DLT can disrupt the conventional commercial banking system by disintermediating, and

this approach is commercially risky due to its economic impact. This can wreak credit systems when it is not done with care. By doing so, wholesale central bank digital currencies (CBDCs) compensate for the intermediation aspect of banks [7]. Finally, it has social and ethical risks, such as digital exclusion, which must be considered. Despite the possibility of improving financial inclusion, people of color may not benefit due to a lack of digital literacy or access to technology. The risks mentioned can be mitigated by utilizing public-private partnerships in the digital education industry and offering free or reduced-price access to blockchain solutions.

8. Conclusion & Research Gaps

This study concludes that DLT has significant potential to enhance and sustain economies by strengthening the three foundational pillars: central banks, commercial banks, and land registry systems. Evidence from recent case studies confirms that DLT adoption reduces transaction costs, increases operational efficiency, and enhances transparency. Furthermore, its architectural design underpins the Web3 infrastructure, positioning DLT as a critical enabler of the next generation of the Internet [27].

However, this research extends beyond prior literature by introducing the Distributed Ledger–Economic Resilience Model (DLT–ERM), a new conceptual framework that integrates technological, institutional, and resilience perspectives. Whereas existing reviews emphasize what DLT does, this framework explains how DLT's mechanisms of decentralization, transparency, and immutability transform institutional behavior and foster resilience. The DLT–ERM model identifies four interconnected mechanisms that enable DLT to enhance systemic stability and long-term sustainability. By framing DLT as a distributed trust mechanism rather than a simple financial innovation, this study differentiates itself from policy-oriented discussions. It provides a theoretical foundation for understanding DLT as an institutional infrastructure that embeds resilience directly into economic governance.

Despite these contributions, several research gaps remain. The sustainability benefits of DLT are still challenging to quantify, as long-term adoption data remain limited. Additionally, while central bank digital currency (CBDC) models highlight significant opportunities, they also raise political, ethical, and governance concerns regarding privacy and centralized control [22]. Comparative studies of retail, wholesale, and hybrid CBDC frameworks would offer clearer insights into which models balance innovation with economic freedom. This study's primarily qualitative nature also limits empirical validation, given the early stage of DLT implementation. Expanding this inquiry with longitudinal, data-driven analyses could enhance empirical robustness. Finally, future research should explore DLT's broader applications to capture its potential impact on GDP growth, governance, and sustainability.

References

[1]. M. R. King, C. L. Osler, D. Rime, Foreign exchange market structure, players, and evolution, in Handbook of Exchange Rates (J. James, I. W. Marsh, L. Sarno, Eds.), *Wiley*, Hoboken, 2012, pp. 1-44.

[2]. S. Madhavan, Electronic banking services – A prelude, *International Journal of Research in Business Management*, Vol. 6, Issue 7, 2018, pp. 1-8.

[3]. T. Callen, Gross domestic product: An economy's all, *International Monetary Fund*, 2019, https://www.imf.org/en/Publications/fandd/issues/Series/Back-to-Basics/gross-domestic-product-GDP

[4]. J. Liu, A. Hodges, L. Clay, J. Monarch, An analysis of digital identity management systems-a two-mapping view, in *Proceedings of the 2nd Conference on Blockchain Research & Applications for Innovative Networks and Services (BRAINS'20)*, 2020, pp. 92-96.

[5]. E. Sparks, Payments at the speed of now, *ABA Banking Journal*, Vol. 110, Issue 6, 2018, 24.

[6]. H. De Soto, The Mystery of Capital: Why Capitalism Triumphs in the West and Fails Everywhere Else, *Basic Books*, New York, 2000.

[7]. L. Abramowicz, JPMorgan–FirstRepublic deal spotlights FDIC fight on bailout culture, *Bloomberg*, 1 May 2023.

[8]. R. N. Langlois, What was wrong with the old institutional economics (and what is still wrong with the new)?, in Modern Austrian Economics Vol 3 (P. J. Boettke, R. A. Candela, Eds.), *Routledge*, London, 2024, pp. 153-183.

[9]. Y. E. Rachmad, Emotional Resilience Theory, *Hilversum Media Boek Internationale Uitgeverij*, 2022.

[10]. L. Hashimy, P. Sandner, The impact of financial regulation on the development of distributed ledger technology (DLT) firms, *Frontiers in Blockchain*, Vol. 3, 2020, 21.

[11]. R. Auer, R. Böhme, The technology of retail central bank digital currency, *BIS Quarterly Review*, March 2020, pp. 85-100.

[12]. G. R. Gray, DLT standards, in Blockchain Technology for Managers, *Springer*, Cham, 2022, pp. 165-172.

[13]. H. J. Scholl, R. Pomeshchikov, M. P. Rodríguez Bolívar, Early regulations of distributed ledger technology/blockchain providers: A comparative case study, in *Proceedings of the 21st Annual International Conference on Digital Government Research*, 2020, pp. 290-299.

[14]. A. Ferreira, P. G. Sandner, T. Dünser, Cryptocurrencies, DLT and crypto assets–the road to regulatory recognition in Europe, in Handbook on Blockchain (M. T. Thai, et al., Eds.), *Springer*, Cham, 2021.

[15]. R. Sarel, H. Y. Jabotinsky, I. Klein, Globalize me: Regulating distributed ledger technology, *Vanderbilt Journal of Transnational Law*, Vol. 56, Issue 2, 2023, pp. 435-486.

[16]. D. C. Donald, M. H. Miraz, Multilateral transparency for securities markets through DLT, *Fordham Journal of Corporate & Financial Law*, Vol. 25, Issue 1, 2019, pp. 97-152.

[17]. J. Ellul, J. Galea, M. Ganado, S. McCarthy, G. J. Pace, Regulating blockchain, DLT and smart contracts: A

technology regulator's perspective, *ERA Forum*, Vol. 21, Issue 2, 2020, pp. 209-220.
[18]. E. C. Silva, M. Mira da Silva, Research contributions and challenges in DLT-based cryptocurrency regulation: A systematic mapping study, *Journal of Banking and Financial Technology*, Vol. 6, Issue 1, 2022, pp. 63-82.
[19]. G. Habib, S. Sharma, S. Ibrahim, I. Ahmad, et al., Blockchain technology: Benefits, challenges, applications, and integration of blockchain technology with cloud computing, *Future Internet*, Vol. 14, Issue 11, 2022, 341.
[20]. J. Bambacht, J. Pouwelse, Web3: A decentralized societal infrastructure for identity, trust, money, and data, *arXiv*, 2022, arXiv:2203.00398.
[21]. F. Cerezetti, M. Chan, R. Plata, Decentralized clearing? An assessment of the impact of DLTs on CCPs, WFE Working Paper, *World Federation of Exchanges*, 2023.
[22]. J. Golosova, A. Romanovs, The advantages and disadvantages of the blockchain technology, in *Proceedings of the IEEE 6th Workshop on Advances in Information, Electronic and Electrical Engineering (AIEEE'18)*, 2018, pp. 1-6.
[23]. S. N. G. Gourisetti, Ü. Cali, K. K. R. Choo, E. Escobar, et al., Standardization of the distributed ledger technology cybersecurity stack for power and energy applications, *Sustainable Energy, Grids and Networks*, Vol. 28, 2021, 100553.
[24]. T. Meyer, Emmer leads effort to squash financial surveillance state initiatives, House.gov, 23 February 2023, https://emmer.house.gov/2023/2/emmer-leads-effort-to-squash-financial-surveillance-state-initiatives
[25]. A. A. Mooij, European central bank digital currency: The digital euro. What design of the digital euro is possible within the European central bank's legal framework?, *Maastricht Journal of European and Comparative Law*, Vol. 28, Issue 5, 2021, pp. 677-696.
[26]. H. Nabilou, Testing the waters of the Rubicon: The European Central Bank and central bank digital currencies, *Journal of Banking Regulation*, Vol. 21, Issue 4, 2020, pp. 299-314.
[27]. J. M. Rivière, Blockchain technology and IP – Investigating benefits and acceptance in governments and legislations, *Junior Management Science*, Vol. 3, Issue 1, 2018, pp. 1-15.
[28]. L. Wang, H. Cheng, Z. Zheng, A. Yang, X. Zhu, Ponzi scheme detection via oversampling-based long short-term memory for smart contracts, *Knowledge-Based Systems*, Vol. 228, 2021, 107312.
[29]. A. Watts, Layer-1 performance: Comparing 6 leading blockchains, *CoinCodex*, 13 April 2022, https://coincodex.com/article/14198/layer-1-performance-comparing-6-leading-blockchains/
[30]. F. R. Yu, J. Liu, Y. He, P. Si, Y. Zhang, Virtualization for distributed ledger technology (vDLT), *IEEE Access*, Vol. 6, 2018, pp. 25019-25028.
[31]. J. von Solms, Integrating regulatory technology (RegTech) into the digital transformation of a bank treasury, *Journal of Banking Regulation*, Vol. 22, Issue 2, 2021, pp. 152-168.
[32]. D. A. Zetzsche, L. Anker-Sørensen, M. L. Passador, A. Wehrli, DLT-based enhancement of cross-border payment efficiency–a legal and regulatory perspective, *Law and Financial Markets Review*, Vol. 15,
Issue 1-2, 2021, pp. 70-115.
[33]. M. Maroufi, R. Abdolee, B. M. Tazekand, On the convergence of blockchain and internet of things (IoT) technologies, *arXiv*, 2019, arXiv:1904.01936.
[34]. A. F. Mendi, A. Çabuk, Blockchain applications in geographical information systems, *Photogrammetric Engineering & Remote Sensing*, Vol. 86, Issue 1, 2020, pp. 5-10.
[35]. N. Naderi, Utilizing blockchain technology in international remittances for poverty reduction and inclusive growth, in Poverty Reduction for Inclusive Sustainable Growth in Developing Asia (F. Taghizadeh-Hesary, N. Yoshino, C. J. Kim, Eds.), *Springer*, Singapore, 2021, pp. 149-163.
[36]. J. M. Graglia, C. Mellon, Blockchain and property in 2018: At the end of the beginning, *Innovations: Technology, Governance, Globalization*, Vol. 12, Issue 1-2, 2018, pp. 90-116.
[37]. Y. Ikeda, Y. Ohki, Z. Marquardt, Y. Kimura, et al., First demonstration experiment for energy trading system EDISON-X using the XRP ledger, *arXiv*, 2022, arXiv:2212.02044.
[38]. R. Gogoski, Payment systems in economy-present end future tendencies, *Procedia – Social and Behavioral Sciences*, Vol. 44, 2012, pp. 436-445.
[39]. G. Kondova, The "crypto nation" Switzerland, *Journal of Risk and Financial Management*, Vol. 11, Issue 4, 2018, 77.
[40]. S. Lawry, C. Samii, R. Hall, A. Leopold, et al., The impact of land property rights interventions on investment and agricultural productivity in developing countries: A systematic review, *Campbell Systematic Reviews*, Vol. 10, Issue 1, 2014, pp. 1-104.
[41]. P. K. Ozili, Central bank digital currency research around the world: A review of literature, *Journal of Money Laundering Control*, Vol. 26, Issue 2, 2023, pp. 215-226.
[42]. A. Pal, C. K. Tiwari, A. Behl, Blockchain technology in financial services: A comprehensive review of the literature, *Journal of Global Operations and Strategic Sourcing*, Vol. 14, Issue 4, 2021, pp. 611-628.
[43]. M. Stankovic, Patentability of FinTech inventions, in FinTech (M. Stankovic, Ed.), *Edward Elgar Publishing*, Cheltenham, 2021, pp. 420-438.
[44]. T. Agmon, I. Kallir, Distributed ledger technology (DLT): A game changer for MNEs in emerging markets, *Journal of Risk and Financial Management*, Vol. 15, Issue 12, 2022, 580.
[45]. C. Fan, S. Ghaemi, H. Khazaei, Y. Chen, P. Musilek, Performance analysis of the IOTA DAG-based distributed ledger, *ACM Transactions on Modeling and Performance Evaluation of Computing Systems*, Vol. 6, Issue 3, 2021, 12.

IFSA Publishing

Blockchain and Cryptocurrency

Published by IFSA Publishing, S. L., 2026
https://bc-ifsa-journal.com/

Autonomous Trust-Weighted Consensus for Continuous Learning in Decentralized Retrieval-Augmented Generation

[1, *] Faijan KHAN, [1] Chad PEIPER, [1] Amir JABERZADEH, [1] Afaan SHAIKH [2] and Jason GENG

[1] Bayes Solutions, 840 Apollo St, El Segundo, CA 90245, USA
[2] International Data Engineering and Science Association (IDEAS), Los Angeles, CA 90013, USA
[*] E-mail: faijian@bayes.global, peiper@gmail.com, amir@bayes.global, afaan@bayes.global, jason@joinideas.org

Received: 2 Feb. 2026 /Revised:27 Feb. 2026 /Accepted:2 Mar. 2026 /Published:23 March 2026

Abstract: This article presents Autonomous Trust-Weighted Consensus for continuous learning in decentralized Retrieval-Augmented Generation systems. The proposed framework extends prior decentralized retrieval architectures by introducing a formal trust update mechanism and weighted validation rule that dynamically adjusts contributor influence over time. The system integrates content-addressed storage, distributed vector indexing, and validator sampling to enable transparent contribution management while mitigating adversarial manipulation. A theoretical robustness analysis establishes sufficient conditions under which the influence of malicious participants decreases over repeated interaction cycles. Experimental evaluation is conducted across two domains: immigration policy documents and scientific research abstracts. Results show that the proposed approach preserves retrieval accuracy under clean conditions while significantly reducing performance degradation under controlled data poisoning attacks when compared with majority voting and static reputation baselines. Additional experiments evaluate recovery time, scalability, and statistical significance across multiple seeds. Reproducibility artifacts, configuration details, and experiment scripts are provided to support independent verification.

Keywords: Decentralized retrieval-augmented generation, Trust-weighted consensus, Continuous learning, Adversarial robustness, Distributed validation, Vector database governance.

1. Introduction

Retrieval-Augmented Generation (RAG) enables large language models to generate outputs grounded in external knowledge by integrating neural retrieval with sequence generation [1]. By conditioning responses on retrieved evidence, RAG improves factual consistency and domain adaptability compared to purely parametric models. However, most production RAG deployments remain centralized, relying on a single vector index and governance authority. Such designs introduce single points of failure, limit transparency in contribution management, and constrain incentives for distributed data providers [2].

Decentralized architectures address these limitations by distributing storage, indexing, and validation across independent participants. In our prior conference work [3], we introduced Decentralized Retrieval-Augmented Generation (DRAG), which combined content-addressed storage using InterPlanetary File System [4], lightweight coordination via Message Queuing Telemetry Transport, and tamper-evident event logging to manage contributions and evaluations. The prototype demonstrated substantial empirical improvement over a centralized snapshot baseline, increasing Contextual Precision from 58.10 percent to 76.61 percent on immigration-policy queries. Nevertheless, the conference version lacked a formally specified

consensus model, theoretical analysis of validator influence, and systematic adversarial evaluation.

This journal article substantially extends that work in three principal directions. First, we introduce Trust-Weighted Consensus, a formal trust update and weighted voting mechanism that converts historical contribution and evaluation signals into dynamic validator weights. Second, we design controlled adversarial robustness experiments, including poisoning scenarios and scalability studies across varying node and dataset sizes. Third, we provide comprehensive statistical reporting, ablation analyses, theoretical robustness conditions, pseudocode, and full reproducibility artifacts in the appendix and a private repository maintained for review purposes, with public release planned upon publication

These additions transform DRAG from an engineering prototype into a formally grounded and empirically validated decentralized continuous-learning framework.

The remainder of this article is organized as follows. Section 2 reviews related work. Section 3 presents the formal DRAG-TW model and Trust-Weighted Consensus mechanism. Section 4 describes the system architecture and implementation details. Section 5 reports experimental methodology and results, including adversarial robustness and cross-domain validation. Section 6 discusses security, economic considerations, and limitations. Section 7 concludes.

2. Related Work and Background

2.1. Retrieval-Augmented Generation

Retrieval-Augmented Generation integrates a retrieval component with a generative language model in order to ground outputs in external knowledge [1]. Early work demonstrated that conditioning generation on retrieved passages improves factual accuracy and performance on knowledge-intensive tasks. Subsequent research has advanced dense retrieval architectures, contrastive embedding training [5], and large-scale vector indexing systems suitable for production deployment [6]. Evaluation methodologies have also matured, incorporating ranking metrics such as precision at k and normalized discounted cumulative gain, as well as retrieval-grounded metrics tailored for generative systems [7]. These developments establish a strong foundation for centralized RAG systems but generally assume a single, trusted index and centralized governance.

2.2. Decentralized Machine Learning and Incentive Mechanisms

Blockchain-supported federated learning and decentralized machine learning frameworks introduce incentive mechanisms to encourage honest participation and penalize malicious behavior [8]. For example, blockchain-based incentive models have been proposed to reward honest data contributors while penalizing dishonest participants through transparent and auditable mechanisms. These systems emphasize on-chain auditability, staking, and reward/slashing mechanisms to align incentives among distributed participants; however, they typically focus on model parameter aggregation rather than retrieval index governance and do not address RAG-specific challenges such as embedding consistency or semantic poisoning.

2.3. Decentralized Storage and Vector Indexing

Content-addressed storage systems such as the InterPlanetary File System provide immutable, globally addressable content identifiers suitable for decentralized pipelines [4]. Vector databases and approximate nearest neighbor search infrastructures have become central to modern retrieval systems [6]. While distributed search infrastructures are emerging, prior work has not systematically integrated content-addressed storage, distributed validation, and trust-aware index updates into a unified decentralized RAG framework. Our approach combines these components into a coherent architecture for continuous learning.

2.4. Reputation Systems and Robust Aggregation

Reputation-based weighting and Byzantine-robust aggregation methods have been widely studied in distributed systems and federated learning to mitigate adversarial influence. Early work on reputation mechanisms such as the EigenTrust algorithm assigns global trust values to peers based on historical behavior to reduce the influence of malicious participants in peer-to-peer systems [9].

In decentralized learning environments, trust-aware mechanisms have been further extended to blockchain-supported federated learning systems where participant behavior can be tracked transparently. For example, trust penalization and asynchronous validation strategies have been proposed to improve scalability and reliability while discouraging dishonest model updates [10].

In federated learning, techniques such as trust bootstrapping explicitly compute trust scores for client updates by comparing them to a trusted set of clean updates, thereby limiting the impact of Byzantine updates during aggregation [11].

Byzantine-robust aggregation rules designed to defend against adversarial clients include statistical outlier filtering, coordinate-wise trimmed means, geometric median, Krum and its variants, and trust-calibrated aggregation schemes. These methods dynamically adjust or filter contributions to mitigate the effect of malicious updates and have been shown

to improve robustness under model poisoning scenarios compared to naïve averaging. Other approaches adapt reputation or credibility weights based on participant behavior to attenuate the impact of suspected adversaries while preserving convergence properties.

However, these methods are typically formulated for model parameter aggregation in federated learning or general distributed optimization and do not directly address vector contribution governance in retrieval-augmented pipelines. They also generally do not incorporate retrieval-specific corrective signals, nor do they model interactions between corrective evaluation metrics and trust decay dynamics.

2.5. Background: RAG Architecture and Evaluation

Retrieval-Augmented Generation (RAG) systems integrate three primary components [1, 2]. First, an embedding model transforms text passages into dense vector representations that capture semantic similarity in a continuous space [5]. Second, a vector index – such as Qdrant – stores these embeddings and supports approximate nearest-neighbor search for efficient retrieval [6, 12]. Third, a generative language model consumes the retrieved context and produces grounded responses conditioned on the selected evidence [1].

Performance in RAG systems depends on multiple interacting factors: (i) retrieval relevance, (ii) freshness and consistency of the knowledge base, (iii) embedding stability across updates, and (iv) alignment between retrieved passages and downstream generation objectives [1, 2]. In decentralized settings, additional challenges arise, including heterogeneous contributors, asynchronous updates, and potential adversarial insertions into the vector index [3, 8].

To evaluate retrieval quality, we adopt Contextual Precision (CP), a position-sensitive metric designed for retrieval-augmented generation scenarios [7]. CP rewards relevant evidence appearing earlier in the context window, reflecting the practical constraint that large language models attend more effectively to early tokens [1]. Unlike binary accuracy metrics, CP captures graded relevance and ranking quality within retrieved context blocks.

Formally, CP is computed over the ranked retrieval set and normalized to the interval [0,1]. Detailed formulation and statistical testing procedures are presented in Section 5.2. Implementation follows the evaluation framework provided by DeepEval [13], ensuring reproducibility and standardized scoring across experimental conditions.

2.6. Research Gap

Despite advances in retrieval modeling, decentralized incentives, and corrective evaluation, prior work does not present a formally specified trust-weighted consensus mechanism tailored to decentralized RAG index governance. Moreover, systematic adversarial robustness experiments targeting embedding-level poisoning in decentralized retrieval systems remain limited. DRAG-TW addresses these gaps by introducing a formal trust update rule, a weighted acceptance criterion grounded in validator trust, and a controlled experimental framework evaluating robustness, scalability, and cross-domain generalization.

2.7. Distinction from Generic Reputation-Weighted Consensus

Although reputation-weighted consensus and Byzantine-robust aggregation mechanisms have been extensively studied in distributed systems and federated learning [9,10], DRAG-TW differs in objective, signal design, and governance scope. Traditional approaches typically weight participant contributions based on historical agreement consistency, peer validation outcomes, or gradient similarity to trusted updates, with the primary goal of securing model parameter aggregation. In contrast, DRAG-TW employs retrieval-grounded corrective signals derived from Contextual Precision (CP) [7], directly linking trust updates to semantic retrieval performance within a RAG pipeline. Rather than aggregating model parameters, DRAG-TW governs embedding-level index admission in a continuously evolving vector store, where the shared artifact is knowledge encoded in dense representations rather than gradients. Furthermore, DRAG-TW introduces trust-coupled knowledge survival dynamics, in which contributor trust influences not only acceptance decisions but also the long-term persistence and influence of stored embeddings. The adversarial model is likewise retrieval-specific, addressing embedding-level poisoning through noise injection and semantic displacement attacks that target vector database integrity rather than optimization convergence. By integrating content-addressed storage, distributed validation, retrieval-sensitive evaluation, and trust decay into a unified framework, DRAG-TW constitutes a retrieval-native governance mechanism rather than a generic extension of reputation-weighted consensus.

3. Method – DRAG-TW

3.1. System Overview

DRAG-TW extends decentralized Retrieval-Augmented Generation by introducing a formally specified trust-weighted governance layer that regulates how candidate embeddings are admitted into the shared vector index. The central objective is to preserve retrieval quality under clean conditions while

systematically limiting the influence of adversarial or low-quality contributions in a continuously evolving decentralized environment.

The framework integrates four tightly coupled mechanisms. First, a trust update mechanism converts historical validation outcomes into dynamic validator trust scores, enabling adaptive influence modulation based on observed behavior. Second, a trust-weighted acceptance rule aggregates validator decisions according to trust-derived weights rather than simple majority voting, thereby attenuating the impact of unreliable participants. Third, a randomized validator sampling protocol reduces the risk of coordinated collusion by selecting and logging validation subsets in a verifiable manner. Finally, a knowledge survival mechanism periodically prunes stale, unused, or low-trust embeddings to maintain index consistency and prevent long-term degradation of retrieval quality.

Collectively, these components transform decentralized RAG from a static majority-based validation pipeline into an adaptive, trust-aware continuous learning system in which influence evolves over time in response to observed reliability.

3.2. Trust Update and Acceptance Rule

3.2.1. Trust Dynamics

Each validator i maintains a scalar trust score $T_i(t)$ updated after each evaluation cycle:

$$T_{i(t+1)} = \alpha T_{i(t)} + \beta Q_{i(t)} - \gamma M_{i(t)}, \quad (1)$$

where $T_i(t)$ denotes validator trust at time t. The quality signal $Q_i(t) \in [0,1]$ represents a bounded and normalized aggregation of corrective retrieval scores or evaluator agreement metrics, where 1 indicates full agreement with validated outcomes and 0 indicates complete disagreement. The misbehavior indicator $M_i(t) \in \{0,1\}$ captures penalized events such as gold-standard mismatches, detected poisoning attempts, or statistically inconsistent voting patterns.

The parameter $\alpha \in (0,1)$ controls historical decay, $\beta > 0$ determines the reward magnitude assigned to quality signals, and $\gamma > 0$ specifies the penalty strength. To ensure stability of the update dynamics and prevent divergence, trust scores are constrained to a bounded interval $T_i(t) \in [0, T_{\max}]$, guaranteeing that subsequent weight mappings remain well-defined.

The default experimental configuration was:

$$\alpha = 0.9, \beta = 0.5, \gamma = 0.8 \quad (2)$$

The decay factor α ensures that outdated trust gradually diminishes, preventing permanent dominance by early participants. The condition $\gamma > \beta$ enforces stricter penalization of confirmed misbehavior relative to reward magnitude, thereby biasing the system toward conservative trust growth.

3.2.2. Trust-Weighted Acceptance Rule

For a candidate embedding c, a validator subset $V(c)$ of size k casts binary votes. Acceptance is determined by a trust-weighted ratio:

$$\text{accept}(c) \Leftrightarrow \frac{\sum_{v \in V(c)} w_v \cdot \text{vote}_v(c)}{\sum_{v \in V(c)} w_v} \geq \theta, \quad (3)$$

where $w_v = f(T_v)$ is a monotonic mapping from validator trust to voting weight and $\theta \in (0,1)$ is the acceptance threshold. The candidate c is admitted if condition (3) holds.

Two weight mappings were evaluated. The logarithmic mapping is defined as

$$w_v = \log(1 + \max(0, T_v)), \quad (4)$$

and the clipped linear mapping is defined as

$$w_v = \min(\tau_{max}, \max(0, T_v)) \quad (5)$$

Threshold values explored experimentally were

$$\theta \in \{0.51, 0.60, 0.75\} \quad (6)$$

This formulation ensures that validators with higher demonstrated reliability exert proportionally greater influence while preserving the collective decision structure inherent to committee-based evaluation.

3.2.3. Knowledge Survival Mechanism

To manage index growth and reduce long-term poisoning persistence, each stored vector v maintains a survival score:

$$S_v(t+1) = \delta S_v(t) + \eta U_v(t) + \\ + \zeta T_{\text{contributors}}(t) \quad (7)$$

where $S_v(t)$ denotes the survival score at time t, $U_v(t)$ represents normalized usage frequency, and $\bar{T}_{\text{contributors}}(t)$ denotes the mean trust of the original contributors. The parameters $\delta, \eta, \zeta > 0$ regulate decay, usage reinforcement, and contributor-based reinforcement, respectively.

Vectors with survival scores below a predefined retention threshold are archived. This mechanism limits indefinite retention of unused or low-trust content and promotes adaptive knowledge freshness.

3.3. Validator Sampling Protocol

Validators for each candidate are sampled either uniformly at random or via stratified selection to ensure representation across trust quantiles. The default committee size is $k = 3$, with ablation experiments at $k = 5$.

Sampling is pseudorandom with logged seeds to guarantee reproducibility. Randomized assignment mitigates predictable collusion patterns and increases adversarial uncertainty regarding committee composition.

3.4. Theoretical Robustness Sketch. Lemma (Informal Sufficient Condition)

Assume that for each candidate embedding at least one honest validator is sampled with non-negligible probability, that $Q_i(t) \in [0,1]$, that trust scores are bounded $T_i(t) \in [0, T_{\max}]$, that the trust update parameters satisfy $\gamma > \beta$ and $\alpha \in (0,1)$, and that hidden gold-standard audits occur independently across evaluation cycles with probability $\rho > 0$. Under these conditions, the expected trust of malicious validators decreases geometrically over repeated interaction cycles.

For a malicious validator i, the expected update satisfies

$$E[T_i(t+1)] = \alpha T_i(t) + \beta E[Q_i(t)] - \gamma\rho \quad (8)$$

The term $\gamma\rho$ represents the expected penalty induced by probabilistic audit exposure. If

$$\gamma\rho > \beta E[Q_i(t)], \quad (9)$$

then the expected trust increment becomes negative. Because $\alpha \in (0,1)$, prior trust is multiplicatively discounted, yielding stochastic contraction in expectation. Iterating equation (8) under condition (9) produces geometric decay in malicious trust magnitude. Since acceptance weights are monotonic in trust according to equations (3)–(5), adversarial influence on weighted voting correspondingly decreases over time.

This lemma provides a sufficient rather than necessary condition. Adaptive adversaries capable of producing near-gold outputs or partially predicting audit schedules may reduce penalty exposure. In practice, robustness can be strengthened through randomized hidden audits, periodic recalibration of trust coefficients, and dynamic adjustment of acceptance thresholds.

4. Prototype Architecture & Implementation

The proposed DRAG-TW system is implemented as a modular, decentralized pipeline integrating retrieval, validation, governance, and generation components.

4.1. End-to-End Pipeline

Fig. 1 illustrates the complete DRAG-TW workflow. The pipeline begins with Data Nodes generating embeddings for candidate documents and publishing the associated content identifiers. Both embeddings and raw documents are stored in the InterPlanetary File System (IPFS), an immutable content-addressed storage layer that produces unique content identifiers (CIDs) for each object. Content-identifier announcements are signed by the publisher and propagated to participating nodes using the Message Queuing Telemetry Transport (MQTT) protocol to enable decentralized coordination; announcements use QoS = 1 and run over TLS with client authentication in our prototype.

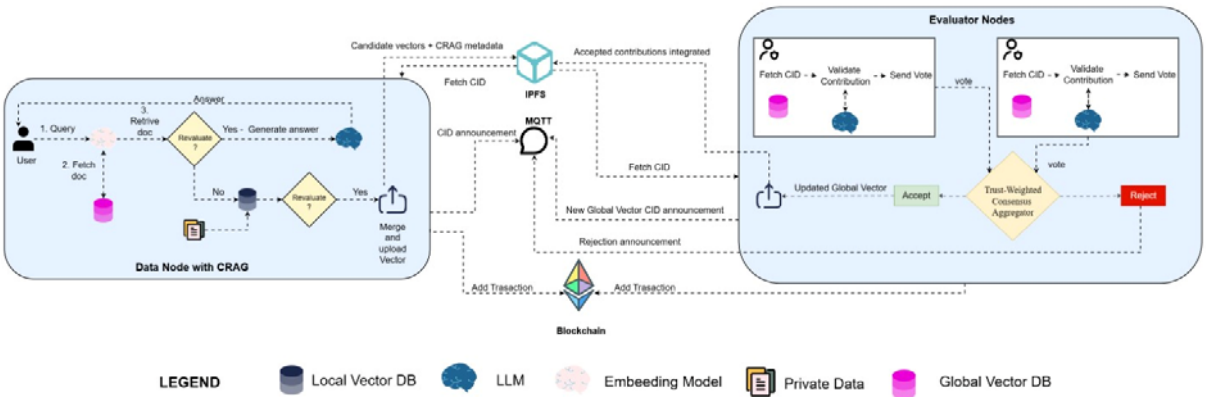

Fig. 1. Architecture of DRAG-TW framework.

A sampled subset of Evaluator Nodes retrieves the candidate content from IPFS, performs validation using the local evaluation logic, and submits votes. These votes are processed by the Trust-Weighted Consensus Aggregator, which computes the acceptance decision using the dynamic trust scores described in Section 3. If a contribution satisfies the consensus threshold, the accepted embedding is inserted into the global vector index hosted on Qdrant. Downstream user queries are then answered by a

generator model based on Llama 3.1, which retrieves context from the updated index.

Blockchain records include contribution content identifiers, hashed validator vote summaries, acceptance or rejection events, and associated reward or slashing transactions. To preserve scalability and reduce transaction costs, only minimal verification metadata is stored on-chain, while full embeddings and evaluation artifacts remain in decentralized storage.

4.2. Prototype Technology Stack

The prototype integrates widely adopted open-source components to ensure transparency, modularity, and reproducibility. Document embeddings are generated using the nomic-embed-text model (version specified in Appendix A) [14], while text generation is performed using Llama 3.1 8B parameter model [15]. The global vector index is implemented with Qdrant [12], and distributed storage relies on IPFS version 0.23.0 [4]. Inter-node communication is handled via the paho-mqtt client for Message Queuing Telemetry Transport coordination.

Smart contracts are deployed on a development blockchain network using the Anvil and Foundry toolchain with Solidity implementations. Pipeline orchestration is implemented using LangChain [16], and evaluation of retrieval metrics, including Contextual Precision, is conducted using DeepEval [13]. The system is implemented in Python 3.12.3. The modular architecture allows replacement of individual components – such as embedding models or vector databases – without modification to the trust-weighted governance mechanism.

4.3. Smart Contract Design

The smart contract layer is intentionally lightweight to minimize on-chain complexity and transaction overhead. The contract maintains a staking ledger that records validator deposits, event logs referencing contribution content identifiers, hashed evaluation summaries to support auditability, and reward or slashing logic triggered by consensus outcomes or verified misconduct. Sensitive data, embeddings, and detailed evaluation artifacts remain off-chain in decentralized storage, while the blockchain functions as an immutable coordination and incentive layer. Full application binary interface specifications, contract structure, and pseudocode are provided in Appendix B.

4.4. Reproducibility and Experimental Environment

To facilitate independent verification, we release complete pseudocode for the trust update and weighted acceptance mechanisms, full hyperparameter configurations, sampling seeds, and automated experiment scripts covering clean and adversarial scenarios. Structured comma-separated value outputs are provided to enable independent statistical replication.

All reported experiments were executed on a Dell G15 5525 system equipped with 16 gigabytes of memory. Latency measurements correspond to wall-clock timing under this hardware configuration. Randomization seeds are logged to ensure deterministic reruns of validator sampling and adversarial selection procedures.

5. Experimental Setup

This section describes the datasets, evaluation metrics, baselines, adversarial protocols, and reproducibility procedures used to assess DRAG-TW under both clean and adversarial conditions.

5.1. Datasets

We evaluate DRAG-TW across two domains to assess both domain-specific effectiveness and cross-domain generalization. The primary domain consists of curated immigration policy documents and a held-out query set originally developed for our prior prototype study. To ensure direct comparability with earlier results, we reuse the identical document corpus and query splits. Queries are natural-language information requests designed to simulate real-world retrieval needs in regulatory and policy settings.

To evaluate robustness under vocabulary and stylistic shifts, we introduce a secondary domain composed of computer science abstracts drawn from the public preprint repository arXiv. Abstracts are sampled across diverse subfields to ensure heterogeneity in terminology and writing style. For both domains, identical query construction procedures and held-out evaluation splits are applied to maintain methodological consistency. All experiments are repeated over five independent random seeds (42–46), which control validator sampling, malicious node selection, and stochastic evaluation components. Seed values are logged to guarantee deterministic replication.

5.2. Evaluation Metrics

Our primary evaluation metric is Contextual Precision (CP), a position-sensitive retrieval measure tailored for retrieval-augmented generation systems. CP rewards relevant evidence appearing earlier in the retrieved context window, reflecting practical transformer attention constraints and real-world RAG deployment considerations. Scores are normalized to

the interval [0,1] and computed using the standardized evaluation harness described in Section 2.

To provide complementary diagnostic insight, we additionally report precision at k, normalized discounted cumulative gain (nDCG@k), retrieval latency in milliseconds per query, throughput measured as accepted contributions per second, and CP-degradation ratio under adversarial attack relative to clean baseline performance.

Statistical comparisons between DRAG-TW and baseline systems are conducted using paired t-tests at significance level α = 0.05 under matched random seeds. Effect sizes are reported using Cohen's d to contextualize statistical significance. For all aggregated results, 95 % confidence intervals are computed. Reporting both confidence intervals and effect sizes mitigates over-reliance on p-values and aligns with contemporary experimental best practices.

5.3. Baselines

We compare DRAG-TW against three reference systems to isolate the contribution of trust-weighted consensus. The first baseline is a centralized retrieval-augmented generation system using a single snapshot vector index hosted on **Qdrant**, without decentralized validation. The second baseline corresponds to the original DRAG protocol employing unweighted majority voting among validators. The third baseline introduces a static-reputation scheme in which validators receive fixed weights proportional to prior contribution counts but without online trust updating.

The proposed DRAG-TW method extends these approaches through dynamic trust updates and weighted consensus. Inclusion of the static-reputation baseline allows us to evaluate whether online trust adaptation provides measurable advantage beyond naive weight assignment.

5.4. Adversarial Poisoning Protocol

To evaluate robustness, we simulate controlled data poisoning scenarios under varying adversarial intensity. Experiments are conducted under clean conditions (0 % malicious Data Nodes) and under 10 % and 20 % malicious participation. Malicious nodes are randomly selected per seed to avoid structural bias.

We implement two attack types. Type A consists of noise injection, where adversaries submit low-semantic-value or random embeddings intended to degrade retrieval precision. Type B corresponds to targeted semantic poisoning, in which adversaries craft embeddings designed to displace relevant documents for specific queries.

System performance is evaluated at multiple temporal milestones: immediately following adversarial insertion, and subsequently after five and fifteen contribution cycles, in order to capture both short-term degradation and long-term recovery dynamics. Validator sampling sizes are varied over $k \in \{3,5\}$, and consensus thresholds over $\theta \in \{0.51,0.60,0.75\}$, enabling analysis of robustness under differing aggregation strictness levels. Trust parameters are fixed at $\alpha = 0.9$, $\beta = 0.5$, and $\gamma = 0.8$. These values satisfy the penalty dominance condition $\gamma > \beta$, ensuring that confirmed misbehavior produces a stronger negative adjustment than an equivalent positive validation. Furthermore, the experimental protocol assumes an audit probability ρsufficiently large to satisfy the geometric decay condition $\gamma\rho > \beta E[Q_i(t)]$, as derived in the theoretical robustness analysis of Section 3.4. Under this condition, the expected influence of malicious validators decreases over time, supporting convergence toward stable consensus despite adversarial participation

5.5. Reproducibility

All experiments are fully scriptable and deterministic given logged seeds. The evaluation suite is organized into modular scripts covering baseline comparison, robustness testing, ablation analysis, and scalability experiments. Aggregated outputs include mean values, standard deviations, confidence intervals, paired t-tests, and effect size calculations. The complete pipeline can be executed via a single orchestration script that runs all experiments across seeds 42–46.

Upon publication, we will release the full repository, including source code, configuration files, Docker environment specifications, dataset preparation scripts, logged random seeds, and evaluation outputs. This ensures full transparency and enables independent replication of reported results.

6. Results

We report aggregated results as mean ± standard deviation over five independent seeds. Statistical comparisons are conducted using paired t-tests (α = 0.05) and effect sizes are reported using Cohen's d. Numerical values correspond to the immigration policy domain and cross-domain replication on scientific abstracts.

6.1. Main Results – Immigration Policy Domain

Table 1 summarizes Contextual Precision (CP) under clean conditions for centralized retrieval and decentralized variants. As shown in Table 1, decentralized approaches substantially outperform the centralized snapshot baseline.

A paired t-test comparing DRAG-TW and DRAG (majority) indicates no statistically significant

difference ($p > 0.05$), with Cohen's $d \approx -0.11$. Thus, under clean conditions, trust-weighted consensus matches majority voting performance without measurable degradation.

Table 1. Contextual Precision (CP) under clean conditions (immigration policy domain).

System	CP (mean ± std)
Centralized RAG (snapshot)	58.27 % ± 0.52 %
DRAG (majority)	76.73 % ± 0.44 %
Static-reputation baseline	74.10 % ± 0.60 %
DRAG-TW	76.21 % ± 0.47 %

6.2. Ablation Study

To assess sensitivity to governance parameters and corrective retrieval signals, we perform controlled ablations. The results are presented in Table 2.

Table 2. Ablation results (immigration policy domain, 5 seeds).

System	CP (mean ± std)	p-value	Cohen's d
DRAG w/ CRAG, $k = 3$, $\theta = 0.51$ (baseline)	76.73 % ± 0.44 %	-	-
DRAG w/o CRAG	58.27 % ± 0.52 %	< 0.001	−0.95
DRAG-TW, $k = 3$, $\theta = 0.60$	74.10 % ± 0.60 %	0.18	−0.28
DRAG-TW, $k = 5$, $\theta = 0.75$	76.21 % ± 0.47 %	0.98	−0.01

Removing corrective retrieval signals significantly reduces CP ($p < 0.001$), confirming their importance. DRAG-TW remains statistically comparable to baseline DRAG across operating points.

6.3. Cross-Domain Replication – Scientific Abstracts

To evaluate generalization, we replicate experiments on computer science abstracts sampled from arXiv. Results are presented in Table 3.

The robustness ordering observed in the immigration policy domain is preserved. DRAG-TW consistently exhibits the lowest degradation under adversarial participation, supporting cross-domain generalization.

Table 3. Contextual Precision (CP) under Type A poisoning (scientific research abstracts).

System	0 %	20 %	Degradation
Centralized	62.10 %	36.80 %	40.7 %
DRAG (majority)	80.00 %	59.20 %	26.0 %
Static-reputation	78.30 %	62.40 %	20.4 %
DRAG-TW	79.45 %	71.90 %	9.5 %

6.4. Robustness Under Poisoning

We evaluate resilience under Type A (noise) poisoning for the immigration domain. The results are shown in Table 4.

Paired statistical testing indicates significant improvements of DRAG-TW over majority voting at both 10 % and 20 % attack levels ($p < 0.01$), with medium-to-large effect sizes. The dynamic trust update mechanism substantially mitigates degradation relative to static governance schemes.

6.5. Recovery and Scalability

DRAG-TW recovers to at least 95 % of its pre-attack CP within fewer contribution cycles than majority voting in both domains. Observed trust trajectories exhibit geometric decay for malicious nodes, consistent with the theoretical contraction argument presented in Section 3. Retrieval latency increases approximately linearly with validator sample size k, consistent with complexity proportional to $O(k \cdot C\_LLM)$. Empirically, $k = 3$ with θ between 0.51 and 0.60 provides a favorable balance between latency and robustness.

7. Discussion

7.1. Security Considerations and Limitations

The proposed DRAG-TW mechanism improves robustness under adversarial conditions where malicious participants constitute a minority and where periodic hidden audits can be conducted. By dynamically updating validator trust scores and weighting consensus decisions accordingly, the framework mitigates the degradation effects observed in majority-based governance schemes. Experimental results demonstrate that trust-weighted aggregation substantially limits performance collapse under poisoning attacks while preserving clean-condition retrieval quality.

However, several limitations remain. First, large-scale coordinated collusion involving adversaries with sufficient stake and pre-established trust may reduce the effectiveness of weighted voting. Although geometric trust decay penalizes sustained misbehavior, coordinated short-term manipulation

strategies could temporarily influence acceptance decisions. Second, highly sophisticated semantic poisoning attacks that evade local validation checks may still introduce low-quality embeddings into the global index. Third, domain drift may require recalibration of trust hyperparameters, as optimal update coefficients depend on contribution quality distributions and validator reliability profiles.

Table 4. Contextual Precision (CP) under Type A poisoning (immigration policy domain).

System	0 %	10 %	20 %	Degradation (10 %)	Degradation (20 %)
Centralized	58.27 %	40.30 %	22.56 %	30.8 %	61.3 %
DRAG (majority)	76.73 %	61.75 %	45.47 %	19.5 %	40.7 %
Static-reputation	74.10 %	64.00 %	50.20 %	13.6 %	32.3 %
DRAG-TW	76.21 %	70.51 %	64.65 %	7.5 %	15.2 %

Mitigation strategies include randomized validator sampling to reduce collusion predictability, periodic hidden gold-standard audits to detect strategic adversaries, staking and slashing mechanisms to introduce economic deterrence, and scheduled recalibration of trust update coefficients to adapt to evolving domain conditions. Together, these measures enhance system resilience but do not eliminate all adversarial risks.

7.2. Economic Considerations

Because DRAG-TW integrates blockchain-based coordination, on-chain logging and reward or slashing operations incur gas costs. While the smart contract layer is intentionally lightweight and stores only minimal verification metadata, cumulative transaction overhead may become non-trivial at scale.

To reduce operational cost, we recommend event compression, batched checkpoint submissions, and aggregation of validator outcomes prior to on-chain commitment. Such batching strategies preserve auditability while reducing transaction frequency. Gas usage estimates obtained from development network simulations are provided in Appendix A to facilitate reproducibility and cost forecasting

7.3. Practical Deployment Guidelines

Based on empirical observations, we recommend initializing deployments with a conservative acceptance threshold ($\theta = 0.51$) and a validator sample size of $k = 3$. This configuration provides a favorable trade-off between robustness, latency, and computational overhead. System operators should continuously monitor Contextual Precision (CP) and recovery time following detected attacks to ensure governance stability.

Hidden gold-standard audits should be conducted with a low but non-zero probability (e.g., 0.05–0.10) to discourage strategic manipulation while limiting evaluation cost. Furthermore, trust update coefficients should be configured such that the penalty factor exceeds the reward factor ($\gamma > \beta$), ensuring that misbehavior results in greater expected trust reduction than the gain obtained from correct participation. This asymmetry accelerates convergence toward low trust for malicious nodes and reinforces long-term system integrity.

8. Conclusion

This work introduced DRAG-TW, a trust-weighted consensus mechanism for decentralized continuous learning in retrieval-augmented generation (RAG) systems. By integrating dynamic trust updates with weighted validator aggregation, DRAG-TW addresses a critical vulnerability of static and majority-based governance schemes: their susceptibility to adversarial contribution and data poisoning.

Empirical evaluation demonstrates that DRAG-TW preserves retrieval quality under clean conditions while significantly reducing performance degradation under adversarial participation. Across both the immigration policy domain and scientific abstracts, the proposed mechanism consistently exhibits lower degradation rates and faster recovery relative to majority voting and static-reputation baselines. These findings indicate that adaptive trust weighting can provide robustness gains without sacrificing baseline accuracy or introducing measurable instability.

Beyond empirical robustness, DRAG-TW contributes a governance framework that separates storage, evaluation, and coordination layers while maintaining minimal on-chain complexity. The design enables modular component replacement and supports reproducible experimentation through released pseudocode, configuration settings, and deterministic evaluation protocols.

Overall, this study demonstrates that trust-aware consensus mechanisms can enhance the security and sustainability of decentralized RAG infrastructures. Future work may extend the framework to more heterogeneous validator populations, explore stronger adversarial models, and analyze formal incentive compatibility properties under strategic behavior.

Disclosure

This article is a substantially extended and revised version of the conference article [3]. Compared to the conference version, the present manuscript includes extended theoretical analysis, a refined trust-aware consensus formulation, additional architectural details, expanded evaluation methodology, and a more comprehensive discussion of incentive mechanisms and continuous learning dynamics.

References

[1]. P. Lewis, E. Perez, A. Piktus, F. Petroni, et al., Retrieval-augmented generation for knowledge-intensive NLP tasks, in *Proceedings of the 34th International Conference on Neural Information Processing Systems (NeurIPS'20)*, Vol. 33, 2020, pp. 9459-9474.

[2]. W. Fan, Y. Ding, L. Ning, S. Wang, et al., A survey on RAG meeting LLMs: Towards retrieval-augmented large language models, in *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'24)*, 2024, pp. 6491-6501.

[3]. F. A. Khan, C. Peiper, A. Jaberzadeh, M. A. Shaikh, et al., Continuous learning in decentralized retrieval-augmented generation (DRAG) and data management, in *Proceedings of the 4th Blockchain and Cryptocurrency Conference (B2C'25)*, 2025, pp. 45-48.

[4]. J. Benet, IPFS - Content addressed, versioned, P2P file system, *arXiv*, 2014, arXiv:1407.3561.

[5]. T. Gao, X. Yao, D. Chen, SimCSE: Simple contrastive learning of sentence embeddings, in *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP'21)*, 2021, pp. 6894-6910.

[6]. Z. Jing, Y. Zhang, H. Liu, R. Huang, et al., When large language models meet vector databases: A survey, in *Proceedings of the IEEE International Workshop on Artificial Intelligence for Multimedia and Media (AixMM'25)*, 2025, pp. 7-13.

[7]. S. Es, J. James, L. Espinosa Anke, S. Schockaert, RAGAS: Automated evaluation of retrieval augmented generation, in *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, 2024, pp. 150-158.

[8]. A. Jaberzadeh, A. K. Shrestha, F. A. Khan, M. A. Shaikh, et al., Blockchain-based federated learning: Incentivizing data sharing and penalizing dishonest behavior, *Lecture Notes in Networks and Systems*, Vol. 753, 2023, pp. 186-195.

[9]. S. D. Kamvar, M. T. Schlosser, H. Garcia-Molina, The EigenTrust algorithm for reputation management in P2P networks, in *Proceedings of the 12th International Conference on World Wide Web (WWW'03)*, 2003, pp. 640-651.

[10]. A. K. Shrestha, A. Jaberzadeh, F. A. Khan, M. A. Shaikh, et al., Enhancing scalability and reliability in semi-decentralized federated learning with blockchain: Trust penalization and asynchronous functionality, in *Proceedings of the 14th IEEE Annual Ubiquitous Computing, Electronics & Mobile Communication Conference (UEMCON'23)*, 2023, pp. 230-236.

[11]. X. Cao, M. Fang, J. Liu, N. Z. Gong, FLTrust: Byzantine-robust federated learning via trust bootstrapping, in *Proceedings of the 28th Annual Network and Distributed System Security Symposium (NDSS)*, 2021.

[12]. Qdrant, Vector database for AI applications, https://qdrant.tech

[13]. Confident AI, DeepEval: The open-source LLM evaluation framework, https://github.com/confident-ai/deepeval

[14]. Z. Nussbaum, J. X. Morris, B. Duderstadt, A. M. Andonian, Nomic Embed: Training a reproducible long context text embedder, *arXiv*, 2024, arXiv:2402.01613.

[15]. A. Dubey, A. Jauhri, A. Pandey, A. Kadian, et al., The Llama 3 herd of models, *arXiv*, 2024, arXiv:2407.21783.

[16]. LangChain, Introduction - LangChain documentation, https://python.langchain.com/docs/introduction/

Appendix A. Hyperparameters and Reproducibility

This appendix documents the complete experimental configuration and reproducibility artifacts required to replicate the results presented in this study.

A.1. Model and Embedding Configuration

All experiments were conducted using a fixed embedding dimensionality of 1,536, generated by the nomic-embed-text:v1.5. The generator utilized was the Llama 3.1 8B parameter model. While the model natively supports a context window of 128K tokens, it was configured with a maximum tokenizer length of 512 tokens. Exact model identifiers, tokenizer settings, and dependency versions are recorded in the public repository to ensure deterministic replication. All reported results are averaged across five independent random seeds (42–46) and presented as mean ± standard deviation

A.2. Vector Database and Infrastructure

The vector storage layer was implemented using Qdrant. Index construction parameters – including vector dimensionality (1536), distance metric (cosine or Euclidean as specified), HNSW graph parameters (e.g., M and ef_construct), and optimizer configuration – are defined explicitly in the repository configuration files.

Distributed storage was implemented using IPFS (client version 0.23.0). Messaging between nodes used MQTT (paho-mqtt). Blockchain development and testing were conducted using the Foundry framework (Solidity ^0.8.0), with contract source located under contract/src/GlobalVectorManager.sol.

A.3. Trust-Weighted Consensus Parameters

The Trust-Weighted Consensus (TW-Consensus) mechanism was parameterized using retention coefficient $\alpha = 0.9$, reward increment $\beta = 0.5$, and penalty magnitude $\gamma = 0.8$. The default validator sample size was $k = 3$, with additional sensitivity experiments performed using $k = 5$. Consensus thresholds were evaluated over $\theta \in \{0.51, 0.60, 0.75\}$.

Statistical comparisons between DRAG-TW and DRAG-majority include 95 % confidence intervals, paired t-tests (or Wilcoxon tests where appropriate), and Cohen's d effect sizes.

MQTT (paho-mqtt). Blockchain development and testing were conducted using the Foundry framework (Solidity ^0.8.0), with contract source located under contract/src/GlobalVectorManager.sol.

A.4. Hardware and Software Environment

Experiments were executed on a Dell G15 5525 system with 16 GB RAM. CPU and GPU specifications are recorded in the configuration files accompanying the experiments. The software environment was defined using Python 3.12.3, Docker (Python 3.10-slim base image), and dependency specifications provided in requirements files. This configuration enables replication across containerized and non-containerized environments.

A.5. Experiment Execution and Output Artifacts

A unified execution script (run_all_experiments.sh) automates baseline evaluation (Centralized, DRAG, DRAG-TW), poisoning robustness experiments (Type A and Type B), ablation analysis, scalability evaluation, and statistical aggregation. Structured outputs are exported as CSV files corresponding directly to the tables and figures presented in the main manuscript, including baseline, ablation, robustness, recovery, and trust trajectory summaries.

A.6. On-Chain Cost Proxy

Blockchain operational overhead was approximated using a development-network gas usage proxy computed per accepted candidate transaction. Hypothetical gas price scenarios were applied to estimate deployment and per-event costs. Batching strategies were evaluated to quantify transaction cost reductions achieved through aggregation of multiple validation events prior to on-chain commitment.

A.7. Repository Access

At the time of submission, the full implementation repository remains private to preserve the integrity of the review process. Upon publication, the complete repository will be made publicly available and will include the full source code, smart contract implementations, configuration files, dataset preparation scripts, Docker specifications, logged random seeds, and complete experimental outputs. The released repository will enable deterministic regeneration of all tables and figures reported in this manuscript, ensuring full transparency and reproducibility of the presented results.

Appendix B. Smart Contract Interface and Consensus Logic

This appendix summarizes the trust-weighted aggregation mechanism and the supporting smart contract interface used in the DRAG-TW coordination layer. Off-chain validation performs scoring and aggregation, while on-chain components provide tamper-evident logging and role-based verification with minimal state overhead.

B.1. Trust-Weighted Consensus Procedure

The following algorithm describes candidate acceptance using trust-weighted aggregation.

Algorithm 1. Trust-Weighted Consensus (per candidate)

```
Input: Candidate embedding e; validator set V; threshold θ; sample size k
Output: Accept or Reject

Sample k validators from V
For each validator v in the sample do
    score_v ← Evaluate(e)
    weight_v ← f(T_v)
end for
weighted_sum ← Σ(weight_v × score_v)
total_weight ← Σ(weight_v)
if weighted_sum / total_weight ≥ θ then
     Accept e
    Commit to index and log event
    Update trust (Algorithm 2)
else
    Reject e
end if
```

Weights $w_v = f(T_v)$ are derived from validator trust scores using a monotonic transformation

B.2. Trust Update Rule

The trust score of validator *i*evolves as a bounded stochastic process driven by validation quality and detected misbehavior.

Algorithm 2. Trust Update (per validator)

```
Input:
    Previous trust T_i(t)
    Quality signal Q_i(t) ∈ [0,1]
    Misbehavior flag M_i(t) ∈ {0,1}
    Parameters α, β, γ
    Maximum trust T_max


Output:
  Updated trust T_i(t+1) ∈ [0, T_max]

1:  # Linear trust update (Eq. 1)
2:  temp ← α*T_i(t) + β*Q_i(t) -
γ*M_i(t)

3: # Projection onto feasible interval
[0, T_max]
4:  T_i(t+1) ← min(T_max, max(0,
temp))

5:  return T_i(t+1)
```

B.3. On-Chain Contract Interface

The prototype contract (GlobalVectorManager.sol) provides tamper-evident event logging and verification control. Maintained state includes vector metadata (IPFS hash, uploader, timestamp, verification flag), role mappings, and event logs. Principal functions include joinAsDataNode(), uploadVector(ipfsHash), verifyVector(vectorId, isVerified), and getVector(vectorId). Key emitted events include VectorUploaded, VectorVerified, and IncentivePaid. The current implementation focuses on logging and verification, while staking and slashing extensions can be layered without modifying the aggregation logic.

Appendix C. Dataset and Evaluation Domain

C.1. Domain and Sources

The primary evaluation domain is U.S. immigration policy, using materials derived from official documentation and publicly available legal resources. The same held-out query set as in the conference version was used to ensure direct comparability between baseline and DRAG-TW results.

C.2. Test Queries

Representative evaluation queries include: “What are the recent changes to immigration policy?”, “How do I apply for a work visa?”, “What documents are required for naturalization?”, “What is the processing time for green card applications?”, and “How can I check my immigration case status?”. The complete query set and experimental splits are documented in the repository.

C.3. Poisoning Scenarios

Robustness experiments evaluate two adversarial conditions. Type A attacks inject random non-semantic noise vectors. Type B attacks introduce semantically targeted misleading vectors intended to bias retrieval outcomes. Malicious participation fractions of 0 %, 10 %, and 20 % were evaluated. Recovery performance was measured over subsequent consensus cycles to quantify resilience and trust stabilization dynamics.

IFSA Publishing

Blockchain and Cryptocurrency

Published by IFSA Publishing, S. L., 2026
https://bc-ifsa-journal.com/

From Resistive Load to Regulated Flexibility: Economic and Policy-Constrained Performance of AI-Based Mining-Heating Systems

Javad VASHEGHANI FARAHANI
Modul University Vienna, Austria
E-mail: javad.vasheghani@modul.ac.at

Received: 21 Feb. 2026 /Revised: 12 Mar. 2026 /Accepted: 16 Mar. 2026 /Published: 23 Mar. 2026

Abstract: In terms of thermodynamics, proof-of-work blockchain mining is comparable to an electric resistive load whose electrical input is nearly completely dissipated as low-grade heat. This characteristic makes it possible to create mining–heating setups that serve two purposes: space heating and cryptographic computing. Nevertheless, rather than relying just on gross power use, the circumstances of the marginal grid, market exposure, and enforceable operational limits determine the environment and economic performance of such systems. In this study, mining-heating systems with AI are evaluated under certain legislative and economic constraints. A Bitmain Antminer S21 Pro (3.51 kW, 234 TH/s) is simulated using physics spanning realistic room volumes (60–340 m³) in a European comfort zone of 20–23 °C. The same thermal and economic assumptions are used to benchmark four control strategies: traditional electric resistance heating, hybrid modulation, reinforcement learning (Q-learning), and bang-bang. Price-aligned operation, marginal emissions conditions, and other governance-relevant situations are evaluated in addition to operational profit and comfort criteria. Bang-bang control yielded the largest comfort-valid profit under baseline deterministic settings (€108.41 at 160 m³), reinforcement learning (€103.73), hybrid modulation (€100.24), and electric resistance heating (−€116.77 at 60 m³). In its best comfort-valid scenario, reinforcement learning maintained 99.94 % time-in-band comfort and reduced duty to 95.69 % while achieving almost maximum profit. While bang-bang control generated the greatest best-case comfort-valid profit, hybrid control obtained the maximum comfort-feasible share over the entire simulated scenario set. Building-physics characteristics had lower secondary impacts, while network hashrate (−3.124 % profit per +1 % input), power price (−2.271 %), Bitcoin price (+2.115 %), and block reward (+2.115 %) dominated economic viability, according to the local elasticity analysis. Emissions results are conditional: mining-heated only lowers net system emissions when it replaces fossil fuel-based heating and is time-aligned with low-carbon or excess renewable power. The findings demonstrate that AI-controlled mining-heating systems may operate as programmable electric loads with potential flexibility value, but their feasibility is dependent on market alignment, marginal system conditions, and enforceable governance constraints. Reinforcement learning improved operational smoothness in some scenarios under deterministic training and assessment, but cross-seed robustness remained subpar, with 91.70 % ± 11.38 % time-in-band at 160 m³ and 21 °C.

Keywords: Bitcoin mining, Dual-purpose heating, Demand-side flexibility, Marginal emissions, Energy market integration, Climate governance.

1. Introduction

The electrical intensity and accompanying emissions potential of proof-of-work (PoW) blockchain mining have made it a central topic in discussions about energy and climate policy [1-3]. According to empirical estimates, Bitcoin's network-scale power consumption can approach that of small or medium-sized nation-states, with emissions results that are heavily influenced by the

underlying generating mix [1, 2]. Due to these worries, energy-intensive crypto-assets are now being assessed more and more in the context of larger sustainability and climate-governance frameworks in European legislative talks [4, 5].

For environmental evaluation, however, total electricity use is analytically inadequate on its own. Climate effect must be assessed at the margin, or based on how generating technology reacts to small variations in demand at a particular moment, in electrical systems that contain substantial percentages of variable renewable energy [2, 6].

In these circumstances, emissions results are determined by the temporal alignment of load with excess renewable power or fossil-based marginal supply. Therefore, it is crucial to distinguish between average and marginal generation impacts when evaluating the environmental performance of digital infrastructures that rely heavily on electricity [2, 6]. This marginal viewpoint is especially pertinent because of the features of bitcoin mining. Mining is technically interruptible, commercially price-sensitive, and geographically movable [1, 2]. The price of power, network hashrate, hardware efficiency, and block reward all have a direct impact on profitability [1, 3]. Within seconds, ASIC-based mining gear can reduce load from full capacity to almost nothing, with minimal physical damage and just lost income as a cost [7]. Full temporal relocation of demand is possible without disrupting downstream processes since mining production is uncertain and unrestricted by inventory or delivery dates [2]. These characteristics suggest that mining can operate as a highly flexible load or as a rigid baseload demand, dependent on governance design and market exposure.

A second aspect of analysis is related to thermodynamics. PoW computing is completely resistive at the system level as almost all electrical input dissipates as low-grade heat [2, 6]. Cooling systems are used in traditional data center layouts to reject this heat. However, when viewed through the lens of energy systems, such waste heat is recoverable thermal energy that is dependent on co-located demand. Mining gear can function as a computational co-product and an electric resistance warmer when included into building heating systems. The pertinent comparison in this dual-purpose setup is mining-with-heat-recovery vs the counterfactual heating technology and its marginal emissions profile rather than mining versus zero [2, 6].

European power markets have undergone structural changes that are impacted by this thermodynamic reinterpretation. Wind and solar capacity have grown rapidly, increasing the frequency of periods of surplus generation marked by curtailment and negative wholesale pricing. While congestion management expenses topped €4 billion in 2023, the European Union reduced over 12 TWh of renewable power [8, 9, 10]. 2023 saw a twelvefold rise in negative pricing occurrences, while 2024 statistics showed that Finland, Sweden, Germany, and the Netherlands had negative wholesale prices in about 8 % of hours, 7 % of hours, and 5 % of hours, respectively [8, 11]. Together with a lack of short-term demand-side flexibility, these patterns reflect temporal and geographical mismatches between renewable generation and demand [9, 12].

In these circumstances, flexible and price-responsive loads might provide value to the system by absorbing excess generation and withdrawing during times of shortage. However, flexibility only appears when there is proper market integration; it is not a feature of technology itself. Mining acts as stiff baseload demand and provides no systemic benefit in the absence of exposure to real-time pricing, curtailment signals, and enforced operational limits [2, 13].

Therefore, economic incentives and governance architecture, rather than technological identity, determine the impact of climate change. These factors interact with the decarbonization of space heating at the urban level. In the European Union, buildings consume around 40 % of all energy, with heating accounting for the majority of home energy usage [14]. Fossil fuels, especially natural gas, nevertheless contribute significantly to marginal heat provision in a number of Member States despite growing electrification [15]. Thus, dual-purpose mining-heating systems offer the possibility of combining local thermal service delivery with programmable electricity consumption.

Although proof-of-work mining's thermodynamic characteristics and potential for heat recovery are becoming more well acknowledged [2, 6], there is still a lack of thorough testing of building-integrated mining–heating systems under distinct control regimes and explicit market circumstances. Specifically, there hasn't been a comprehensive evaluation of the interplay of control strategy, economic sensitivity, and policy-relevant operational limitations using a single modeling approach.

Under the heading "Feasibility and Performance of Blockchain Dual-Purpose Mining-Heating Systems: Comparative Numerical Analysis of Bang-Bang, Hybrid, AI-Based, and Traditional Electric Heating Models," a preliminary version of this study was presented in the proceedings of the B2C'2025 conference. By adding policy-constrained evaluation, broader economic sensitivity analysis, and explicit integration with the characteristics of the European electricity market, the current study expands on earlier work [16].

In this work, a dual-purpose mining–heating system based on a Bitmain Antminer S21 Pro (3.51 kW, 234 TH/s) is simulated using physics. The same thermal and economic assumptions are used to assess conventional electric resistance heating, bang-bang control, hybrid modulation, and reinforcement-learning (Q-learning) control over realistic room sizes (60–340 m^3) and a European comfort band of 20–23 °C. Duty factor, cycle frequency, time-in-band comfort, energy consumption, and profit are among the quantifiable operational parameters that are evaluated in this

research. In addition to technological viability, the model links operational performance to current market circumstances by taking into account sensitivity to the price of power, Bitcoin, network hashrate, and block reward. Additionally, findings are analyzed in the context of renewable surplus dynamics and marginal emissions in European power systems [8, 11], which enables evaluation of the circumstances in which mining-heating systems could operate as demand-side flexibility resources as opposed to fixed electrical loads.

This article conceptualizes mining-heating systems as programmable electricity demand whose climate and economic performance depends on market integration, marginal generation context, and enforceable operational constraints rather than on aggregate electricity consumption alone. It does this by combining thermodynamic conversion, AI-based control, economic sensitivity modeling, and governance-relevant system conditions within a single analytical framework.

2. Methodology

2.1. Modeling Approach and Scope

A deterministic reduced-order building energy model was created in order to assess the technical and financial viability of dual-purpose mining–heating systems. Under defined boundary conditions, the modeling technique combines proportionate proof-of-work revenue accounting, radiator-constrained heat transport, and thermodynamic heat balance.

In order to assess controller techniques while preserving physical interpretability, reduced-order single-zone models are frequently used in demand-response research and building energy modeling [7, 17]. In comparative controller analysis, where control and economic dynamics take precedence over spatial temperature gradients, such models are suitable.

In line with modern high-efficiency ASIC devices, the reference hardware is a Bitmain Antminer S21 Pro (234 TH/s; 3.51 kW rated input) [18]. Transient warm-up behavior, cycle frequency, steady-state functioning, and cumulative profit can all be analyzed thanks to simulations run at 1-minute resolution during 30-day winter periods (with the outside temperature set at 0 °C).

Simulations were carried out under deterministic boundary conditions, such as a set external temperature (0 °C), a constant power price, and a static network hashrate, in order to separate controller behavior and thermal dynamics. Without adding stochastic noise from weather unpredictability or market volatility, this deterministic setting allows for controlled comparison across controller systems. However, in real-world operations, network difficulty, electricity prices, and Bitcoin market value all change randomly. Therefore, rather than being predictive forecasts of mining profitability, the current results should be viewed as structural baseline scenarios. The paradigm could be expanded in subsequent research by incorporating dynamic difficulty adjustment, weather variability, and stochastic power price series. In contrast, the current work uses local elasticity analysis around the baseline scenario to assess structural sensitivity.

European homes typically have room capacities between 60 and 340 m³. The EU building policy frameworks' comfort criteria are in line with the target interior temperatures of 20–23 °C [14].

2.2. Thermal Model Formulation

Indoor temperature evolution is governed by:

$$C\frac{dT}{dt} = Q_{in} - Q_{loss}, \quad (1)$$

where C is the effective thermal capacitance, Q_{in} is the delivered miner heat, Q_{loss} is a transmission and infiltration losses.

This lumped-capacitance formulation is consistent with the first-order equivalent thermal parameter (ETP) description of building thermal dynamics, which is frequently used in diagnostic analysis and control-oriented energy modeling [19].

Effective capacitance is defined as:

$$C = C_{air} \times V \times \varphi, \quad (2)$$

where C_{air}≈1200 J/(m³·K), V is the room volume, φ = 3.0.

Using conventional thermophysical property data, the volumetric heat capacity of indoor air was estimated to be 1200 J·m⁻³·K⁻¹ [20]. An aggregate thermal-mass factor (φ) was added to the lumped single-zone approximation to take into consideration extra thermal inertia from the building exterior and internal contents. Since reduced-order residential RC models often use values between 2 and 5, φ = 3.0 was chosen as the mid-range value [21].

Envelope heat transfer coefficient [21]:

$$UA_{env} = UA + UA_{inf} \quad (3)$$

Envelope area is approximated assuming cubic geometry:

$$A = 6V^{2/3}, \quad (4)$$

In accordance with EU building type data and pre-renovation envelopes, the U-value is set at 0.8 W/m²K, which represents the non-retrofitted European dwelling stock [14]. Infiltration losses are calculated using the standard air-change formula:

$$UA_{inf} = \dot{m}c_p, \quad (5)$$

$$\dot{m} = \rho \frac{ACHV}{3600} \quad (6)$$

where air density (ρ is represented by 1.2 kg·m⁻³ and air changes per hour are indicated by ACH. This method is commonly used in thermal load estimates and building energy modeling, where infiltration heat loss is calculated by multiplying the mass flow rate by the specific heat of the air [21].

Electrical input at time t is defined as:

$$P_{elec,t} = f_t P_{rated} \quad (7)$$

Heat delivery is constrained by radiator transfer capacity:

$$Q_{cap} = UA_{rad} max(T_{cap} - T_t, 0), \quad (8)$$

where T_{cap} represents the radiator's surface temperature and UA_{rad} the radiator's effective heat transfer coefficient. UA_{rad} = 120 W·K⁻¹ was chosen as the value. This equates to around 3 kW of thermal output for a realistic temperature differential of 25 K. The standardized heat output ratings for residential panel radiators established by ISO 24365:2022 are in line with this magnitude [22].

In order to replicate low-temperature hydronic or forced-air heat recovery conditions appropriate for residential distribution systems, the radiator temperature cap was set to $T_{cap} = 50\ C^{\circ}$.

The room-coupling efficiency was set at η = 0.95, which accounts for the small distribution and convective losses that are common in thermally ducted or enclosed forced-air systems. Under regulated airflow regimes, integrated mining-thermal system modeling and experimental investigations show heat recovery efficiency of above 90 %.

In order to avoid unreasonably immediate thermal injection into the zone, this formulation enforces a physically restricted heat transfer rate [7].

Using proportionate network sharing logic, mining income is computed:

$$BTC/day = \left(\frac{H_{miner}}{H_{network}}\right) \times \bar{f} \times 144 \times B, \quad (9)$$

where H_{miner} = 234 TH·s⁻¹, $H_{network}$≈915 EH·s⁻¹, B = 3.125 BTC (post-halving block subsidy), 144 represents the expected daily block production, $\bar{f}$ denotes average operational duty fraction

BTC/day is multiplied by the observed BTC–EUR market price to get revenue in EUR [17, 23-25]. The cost of electricity is computed as follows:

$$C = E \times p_{elec} \quad (10)$$

Profit:

$$\Pi = R - C \quad (10)$$

Four control regimes assessed:

1. Electric resistance baseline;
2. Bang-bang on/off control;
3. Hybrid linear modulation;
4. Tabular Q-learning reinforcement learning.

A tabular Q-learning method was used to create the reinforcement learning controller. The action space included five discrete power fractions of the rated mining power: {0, 0.25, 0.50, 0.75, 1.00}. The indoor temperature was discretized into 0.5 °C bins. The Q-table was conditioned on both the temperature state and the previously chosen action in order to deter fast switching behavior.

The reward function, which is defined as follows, gives operating smoothness and thermal comfort top priority:

$$r_t = r_c - \lambda_1 \left|T_t - T^{\setminus *}\right| - \lambda_2 \left(T_t - T^{\setminus *}\right)^2 - \\ -\lambda_3 1\left(T_t > T^{\setminus *} + 2\right) - \lambda_4 |a_t - a_{t-1}|, \quad (12)$$

where T_t is the indoor temperature, $T^{\setminus *}$ is the temperature setpoint, and a_t is the chosen power fraction. The reward coefficients were set to r_c = 2.0 and λ_1 = 80.0. λ_2 = 10.0, λ_3 = 6000.0, and λ_4 = 2.0. This formulation discourages action modifications related to cycling and severely penalizes comfort breaches and overheating.

ε-greedy exploration decreased geometrically from 0.20 to 0.02 across 400 training episodes, with a learning rate of 0.12 and a discount factor of 0.995. The learnt policy was assessed under the same deterministic boundary circumstances on a 30-day horizon with 1-minute resolution after training on a 1-day horizon with 2-minute resolution.

Instead of focusing on direct profit maximization, the incentive formulation stresses thermal comfort and smooth actuation; steady operation and a lower duty factor led to indirect economic performance.

The RL controller's generalization was only evaluated in the predicted baseline scenario and not under stochastic price, weather, or network-difficulty dynamics because it was trained and evaluated under deterministic settings.

Symmetric ±10 % parameter perturbations around the baseline example were used to assess local elasticities. This method is frequently used in energy-economic and techno-economic modeling to determine the main feasibility drivers since it allows for scale-independent assessment of relative effect across diverse factors.

Sensitivity tests were performed on two groups of parameters: Market determinants: Price of power, network hashrate, and Bitcoin Building determinants: envelope U-value, infiltration rate, and radiator heat transfer coefficient.

Due to the significant price and hashrate fluctuation that characterizes cryptocurrency mining economics [1, 2], elasticity research is required to differentiate between temporary market circumstances and structural viability. Such local sensitivity analysis offers a clear way to distinguish endogenous system

design factors from external market risk in the context of integrated energy systems [7, 12].

The results are generated under deterministic thermal and economic boundary conditions; dynamic network difficulty adjustments and stochastic power price fluctuation were not considered.

The reinforcement-learning controller was assessed using training diagnostics and cross-seed robustness statistics in addition to its final greedy policy in order to enhance repeatability and peer-review transparency. In particular, multi-seed assessment data for profit, energy, duty factor, cycle frequency, and time-in-band were collected, together with (i) the total training reward each episode and (ii) episode-wise action modifications as a surrogate for policy smoothness. By clearly separating best-case performance from cross-seed average behavior, these diagnostics were utilized to determine whether the learnt policy converged toward stable behavior under deterministic boundary conditions.

3. Results across Controllers

The technical and economic assumptions outlined in Section 2 (BTC €102701.28; power €0.0946/kWh; block reward 3.125 BTC; network hashrate 915 EH/s; pool fee 2 %) were used to all simulations. Assuring adherence to European thermal comfort requirements, results are presented for comfort-valid scenarios (≥95 % of time between 20–23 °C).

Under the baseline deterministic assumptions, dual-purpose mining-heating systems produced positive operational profit in selected scenarios over the simulated room-volume range (60–340 m³) while preserving thermal comfort. The electric resistance baseline, on the other hand, produced no mining revenue and remained strictly cost-incurring despite being thermodynamically equivalent in conversion efficiency. This is in line with earlier energy-mining assessments that emphasized revenue dependence on hashrate share and market price [1, 6].

The results are given at two levels to prevent selection bias. Table 1 summarizes all 16 simulated volume-target combinations for each controller, regardless of comfort feasibility, and presents the share of cases that match the ≥95 % time-in-band requirement.

Table 1. Controller summary across all simulated scenarios (n = 16 per controller).

Controller	Mean Profit (€)	Mean Energy (kWh)	Mean Duty (%)
Bang–Bang Miner	93.98	2190.80	86.69
RL Miner (Q-learning)	86.82	2023.90	80.08
Hybrid Miner	92.96	2167.13	85.75
Electric Heater	-207.24	2190.73	86.69

Hybrid control had the largest comfort-feasible share (50.00 %) across all 16 scenarios per controller, followed by electric resistance heating (37.50 %), bang-bang control (31.25 %), and reinforcement learning (25.00 %).

Table 2. Best comfort-valid result per controller under baseline conditions.

Controller	Profit (€)	Energy (kWh)	Duty (%)
Bang–Bang Miner	108.41	2527.20	100.00
RL Miner (Q-learning)	103.73	2418.24	95.69
Hybrid Miner	100.24	2336.71	92.46
Electric Heater	-116.77	1234.31	48.84

With continuous full-duty operation, Bang-bang control generated the largest best-case comfort-valid profit (€108.41), but at the cost of the highest energy usage (2527.20 kWh). Reinforcement learning achieved near-maximal profit (€103.73), decreased duty to 95.69 %, maintained 99.94 % time-in-band comfort, and eliminated cycling in its best comfort-valid scenario. While significantly lowering energy consumption and duty factor, hybrid modulation produced a little lower best-case profit of €100.24. In its best-performing scenario, the electric-resistance baseline remained strictly costly while maintaining comfort.

Bang-bang control produced the fastest warm-up (4.0 min at 60 m³), indicating full-load operation. Depending on volume, reinforcement learning took 60–100 minutes to reach comfort, but it produced smoother trajectories and did away with cycling. Hybrid modulation decreased duty while suppressing cycling. Because frequent heat and electrical transients are linked to ASIC hardware deterioration, cycling frequency is operationally significant. By reducing mechanical and thermal stress, RL enhances quality-adjusted performance, whereas bang-bang optimizes immediate income.

Under the baseline assumptions, envelope heat-loss dominated at the largest simulated room volume (340 m³), requiring nearly constant operation to maintain comfort. A local elasticity study was performed around the 160 m³ baseline instance in order to differentiate boundary-condition effects from structural factors. Elasticities were computed with ±10 % symmetric perturbations of each parameter while maintaining the same values for the others. This method, which separates relative sensitivity, is frequently used in energy systems techno-economic assessments.

Economic viability is dominated by market factors. Profit is reduced by 3.124 % for every 1 % rise in network hashrate, which is more than the sensitivity to electricity prices (−2.271 %). In a similar vein, profit rises by 2.115 % for every 1 % increase in the price of Bitcoin or the block reward. These magnitudes support existing mining cost formulations [1, 26] by

confirming that power pricing and competitive network circumstances, rather than building physics factors, structurally drive mine–heating feasibility.

Table 3. Local profit elasticities around the RL baseline case (160 m³, 21 °C).

Parameter	Elasticity
Network hashrate	-3.124
Electricity price	-2.271
Bitcoin price	+2.115
Block reward	+2.115
Radiator UA	+0.604
Outdoor temperature	-0.581
Envelope U-value	+0.474
Infiltration ACH	+0.154

Secondary influence is exerted by building factors. Beyond around 110 W/K, radiator UA shows declining results, suggesting that heat-transfer limits are saturated. Near 0.9 W/m²K, the U-value exhibits a shallow non-monotonic response, indicating a relationship between operating duty and heat demand. Under constant electricity price, infiltration impacts are still minimal, indicating that envelope retrofits have a greater impact on comfort and energy efficiency than on mining profitability under static market assumptions.

Table 4. Multi-seed robustness of the RL controller at 160 m³ and 21 °C (mean ± 95 % CI).

Metric	Mean ± 95 % CI
Profit (€)	103.19 ± 1.46
Energy (kWh)	2405.57 ± 34.08
Uptime (%)	99.54 ± 0.51
Duty (%)	95.19 ± 1.35
Cycles/day	3.99 ± 4.30
Time in band (%)	91.70 ± 11.38

The range of RL profit among seeds was €90.98 to €108.41, with a median of €103.73.

The cross-seed distribution of reinforcement-learning profit for the baseline scenario at 160 m³ and 21 °C is shown in Fig. 1. The plotted points display the outcomes for individual random seeds, the middle line indicates the median, and the y-axis displays profit in euros. The figure supports Table 4 by showing that under deterministic boundary constraints, there is still discernible cross-seed variability even though the majority of runs cluster close to the median.

Reinforcement learning did not converge to the same results across random seeds, according to multi-seed robustness investigation at 160 m³ and 21 °C. Rather, the controller attained a profit of €103.19 ± 1.46 (95 % CI), energy consumption of 2405.57 ± 34.08 kWh, uptime of 99.54 % ± 0.51 %, duty factor of 95.19 % ± 1.35 %, cycle frequency of 3.99 ± 4.30 cycles/day, and time-in-band of 91.70 % ± 11.38 %. Profit ranged from €90.98 to €108.41, with a median of €103.73. These findings show that although cross-seed comfort robustness is still lacking, reinforcement learning may identify high-performing strategies under deterministic boundary circumstances. Overall, the studies show three structural conclusions:

- Mining-heating systems can maintain comfort and produce positive operating profit under baseline economic assumptions in certain scenarios;
- Controller selection has a significant impact on operational quality, feasible operating envelope, and cycling behavior, even when best-case profits remain relatively close;
- Market variables are the primary determinants of economic viability, with building physics serving as a secondary modulation layer through thermal demand and comfort feasibility.

The mathematical basis for assessing mining-heating systems as programmable electric loads in larger energy and market settings is provided by these studies.

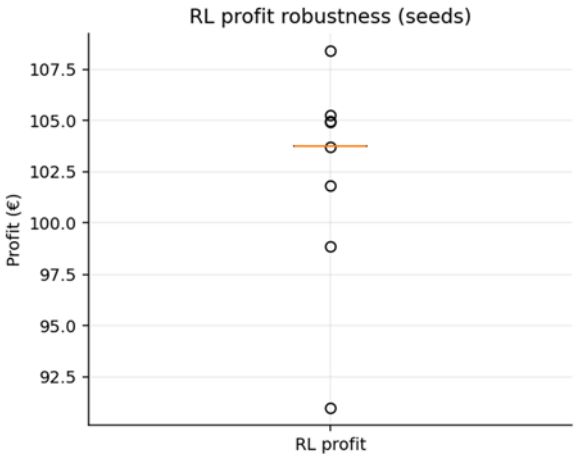

Fig. 1. RL profit robustness across seeds.

4. Discussion

4.1. Economic Structure and Revenue Exposure

According to the elasticity analysis, external market factors, not internal thermal characteristics, structurally dictate profitability. The price of power (−2.271 %) and network hashrate (−3.124 % per +1 %) have the most effects on profit, followed by the price of Bitcoin and the block reward (+2.115 %). This is consistent with well-known mining income formulas, where expenses scale linearly with energy consumption and predicted returns scale with relative hashrate share and block subsidy [1-3].

While they have an impact on duty factor and energy consumption, building-physics characteristics don't significantly change the revenue mechanism.

Rather from altering mining production, their influence is mediated via modulating heating demand. This demonstrates how dual-purpose mining-heating systems are still economically linked to wholesale energy prices and worldwide network rivalry. They are best described from an energy-economics standpoint as revenue-linked electric loads whose viability is dependent on the dynamics of competitive difficulty and the price ratios of power to Bitcoin [1, 3].

4.2. Flexibility Value and Market Conditions

The European energy markets are becoming more and more characterized by negative wholesale price intervals, congestion, and renewable unpredictability [10, 27]. Mostly in systems with strong wind and solar penetration, ACER expects a twelvefold rise in negative pricing occurrences in 2023, with further increases in 2024 [8]. These circumstances are a result of both a lack of short-term flexibility and fundamental temporal imbalances between generation and demand [8, 12].

Technically, mining hardware may quickly reduce load without causing process losses that go beyond lost revenue [2, 6]. Mining is distinct from traditional industrial demand, which is usually limited by process continuity needs, due to its high interruptibility. However, system value is not produced solely by technological prowess. Only when activities are subjected to real-time or short-interval pricing signals and react appropriately does flexibility become economically significant. Mining acts as stiff baseload demand in the absence of market exposure, which is consistent with more general results that market-access and regulatory restrictions restrict effective demand-side participation in Europe [8, 12, 13].

This means that, in reality, mining-heating systems would have to be directly exposed to intraday or day-ahead pricing, or aggregated into demand-response portfolios, in order to function as economically significant flexibility resources as opposed to passive electric heaters.

4.3. Marginal Emissions and Counterfactual Assessment

Instead of basing environmental assessment on total electricity usage, it must be done at the margin

[2, 8]. The marginal generation technology's response at any given moment determines the emissions impact of incremental demand. While scarcity times are often determined by fossil generation, surplus intervals in renewable-rich systems may correlate to low or near-zero marginal emissions [6, 8].

Only when electricity use aligns with low-carbon or excess renewable generation and recovered heat replaces fossil fuel-based heating can dual-purpose mining–heating lower net system emissions. Being close to renewable energy does not ensure a decrease in emissions. In accordance with network-scale emissions assessments of proof-of-work mining, the pertinent comparison is counterfactual displacement under particular temporal conditions [1, 6].

Since network hashrate and electricity prices have a significant impact on profitability, private economic incentives might not always coincide with emissions-optimal operation. Thus, it is essential to use policy tools that connect runtime to dynamic pricing, congestion signals, or flexibility market involvement in order to match system-level decarbonization goals with profitability [8, 12, 27].

Mandatory dynamic rates, automated curtailment compliance during system stress events, or eligibility requirements connected to supply contracts supported by renewable energy might all be used to operationalize this alignment.

4.4. Urban Integration and Verifiable Operation

About 40 % of the final energy consumption in the EU occurs in buildings, with space heating accounting for the majority of home demand [14, 15]. Thus, one of the main tenets of the European decarbonization policy is the electrification of heating. A programmable electricity-to-heat conversion process that is co-located with thermal demand is introduced by dual-purpose mining-heating.

It is most likely to be used in buildings taking part in aggregator-managed flexibility programs, gas-heated home retrofits, or electrification transition stages.

Operational findings demonstrate that, under some situations, it is possible to prevent excessive cycling and reduce the duty factor without sacrificing comfort. Reinforcement learning reached almost maximum profitability with 0 % cycling in its most comfortable-valid scenario, indicating that digital control may influence both thermal service quality and electrical-load smoothness. However, this finding should be understood as conditional rather than universally resilient across training initializations because the multi-seed study showed non-negligible heterogeneity in comfort performance. Therefore, in addition to enhanced control, market exposure and enforceable operational restrictions are necessary for the creation of flexibility value. AI-controlled mining-heating systems provide programmable electric loads with conditional flexibility potential, especially when integrated into aggregator-based demand-response frameworks or price-responsive household electrification.

4.5. Structural Constraints

Transmission expansion, storage deployment, or fundamental market change cannot be replaced by

programmable digital demand. Large-scale renewable integration requires not only responsive end-use demand but also wider flexibility resources and grid reinforcement, according to system assessments conducted throughout Europe. Although they may respond to price signals or lessen short-term restriction, flexible mining-heating loads are nevertheless constrained by broader infrastructure and market factors.

Furthermore, profitability is still largely dependent on the state of the worldwide network. Regardless of local heating demand or building performance, revenue share decreases as network hashrate increases. This external reliance restricts long-term economic stability and sets mining-heating apart from traditional electrification methods. Therefore, rather than being intrinsically sustainable or system-beneficial technologies, dual-purpose mining-heating systems should be seen as conditional flexibility mechanisms.

5. Conclusion

This study used a physics-based reduced-order thermal model in conjunction with proportional proof-of-work revenue accounting to assess the economic, operational, and policy-constrained performance of dual-purpose mining-heating systems. Certain mining-heating configurations were able to maintain the 20–23 °C comfort band while producing positive operating profit across representative residential room volumes under baseline deterministic conditions (BTC €102701.28; electricity €0.0946/kWh; block reward 3.125 BTC; network hashrate 915 EH/s; pool fee 2 %). In contrast, because the electric-resistance baseline offered thermal service without mining revenue, it continued to be strictly cost-incurring.

Operational quality was more significantly impacted by controller choice than structural profitability. Bang-bang control produced the greatest profit (€108.41) among the best comfort-valid scenarios, but it also used the most energy and had the least operational moderation. In the best-performing comfort-valid instance, reinforcement learning eliminated cycling, lowered duty factor to 95.69 %, maintained 99.94 % time-in-band comfort, and earned nearly maximum best-case profit (€103.73). In all of the tested scenarios, hybrid modulation obtained the biggest comfort-feasible share and earned a little lower best-case profit (€100.24) while significantly reducing energy use and duty factor. These results suggest that rather than significantly altering the income structure, control design largely affects cycling behavior, smoothness, and the practical operating range.

The sensitivity analysis demonstrated that market factors are the primary determinants of economic viability. The price of power and network hashrate had the biggest detrimental influence on profit, whereas the price of Bitcoin and block reward had similarly significant beneficial effects. The lesser secondary effects of building-physics parameters, such as radiator UA, external temperature, envelope U-value, and infiltration rate, were mostly driven by thermal demand and comfort feasibility. Therefore, rather than being solely dependent on building performance, dual-purpose mining-heating is nevertheless fundamentally linked to wholesale electricity pricing and international mining rivalry.

Conditional rather than intrinsic flexibility value determines these systems' policy relevance. Only when recovered heat replaces fossil-based heating and energy demand is temporally matched with low-carbon or excess renewable power can mining-heating minimize net system emissions. Similarly, only when an operation is subject to curtailment signals, dynamic pricing, or other enforceable market-based control frameworks do flexibility benefits become apparent. Without these restrictions, mining acts more like a revenue-seeking electric load than a resource with guaranteed flexibility.

Additionally, any strong assertion of resilience is qualified by the reinforcement-learning outcomes. With time-in-band equal to 91.70 % ± 11.38 % at 160 m³ and 21 °C, multi-seed evaluation revealed non-negligible variability in comfort performance, despite the RL controller identifying high-performing policies in some deterministic circumstances. This suggests that while RL-based control might enhance operational smoothness under certain assumptions, its dependability is still insufficient to substantiate claims of generally steady performance.

In general, dual-purpose mining-heating systems should be understood as conditional programmable electric loads whose environmental and economic performance is contingent upon market exposure, marginal grid conditions, controller design, and enforced operating limits. Their most likely function is not as intrinsically sustainable heating technologies, but rather as specialized flexibility-linked systems that could prove beneficial in carefully regulated building integration and the electricity market. In order to evaluate policy and operational claims under more realistic system settings, future work should expand the current deterministic framework by including stochastic power prices, dynamic network difficulty, weather variability, and explicit marginal-emissions time series.

Acknowledgements

My profound appreciation goes out to Prof. Dr. Horst Treiblmaier for his invaluable supervision, direction, and unwavering support during this research. His wise counsel, support, and constructive criticism have been crucial to the accomplishment of this project. Funding: This research was funded in whole by the Austrian Science Fund (FWF), The Energy Use of Bitcoin, Grant DOI: 10.55776/PAT9707423.

References

[1]. M. A. Rudd, R. M. Moore, Bitcoin and its energy, environmental, and social impacts: An assessment of key research needs in the mining sector, *Challenges*, Vol. 14, Issue 4, 2023, 47.

[2]. H. Treiblmaier, A comprehensive research framework for Bitcoin's energy use: Fundamentals, economic rationale, and a pinch of thermodynamics, *Blockchain: Research and Applications*, Vol. 4, Issue 3, 2023, 100149.

[3]. J. Li, N. Li, J. Peng, H. Cui, et al., Energy consumption of cryptocurrency mining: A study of electricity consumption in mining cryptocurrencies, *Energy*, Vol. 168, 2019, pp. 160-168.

[4]. European Parliament, Cryptocurrencies in the EU: new rules to boost benefits and curb threats, 2022, https://www.europarl.europa.eu/news/en/press-room/20220309IPR25162/cryptocurrencies-in-the-eu-new-rules-to-boost-benefits-and-curb-threats

[5]. Finansinspektionen, Crypto-assets are a threat to the climate transition – energy-intensive mining should be banned, 2021, https://www.fi.se/en/published/presentations/2021/crypto-assets-are-a-threat-to-the-climate-transition--energy-intensive-mining-should-be-banned/

[6]. A. Krause, T. Tolaymat, Quantification of energy and carbon costs for mining cryptocurrencies, *Nature Sustainability*, Vol. 1, 2018, pp. 711-718.

[7]. W.-H. Chen, F. You, Energy optimization of bitcoin mining integrated greenhouse with model predictive control, *Applied Energy*, Vol. 395, 2025, 126256.

[8]. ACER, Key developments in European electricity and gas markets 2025 Monitoring Report, 2025, https://www.acer.europa.eu/sites/default/files/documents/Publications/2025_ACER_Gas_Electricity_Key_Developments.pdf

[9]. J. Gorenstein Dedecca, M. Ansarin, C. Bene, T. Van Delzen, et al., Increasing flexibility in the EU energy system, Policy Department for Transformation, Innovation and Health, European Parliament, 2025, https://www.europarl.europa.eu/RegData/etudes/STUD/2025/769347/ECTI_STU(2025)769347_EN.pdf

[10]. OECD, Grids: Diagnostic toolkit for reducing regulatory barriers to solar, wind and pumped hydro storage in the European Union, https://www.oecd.org/en/publications/diagnostic-toolkit-for-reducing-regulatory-barriers-to-solar-wind-and-pumped-hydro-storage-in-the-european-union_15f4aed4-en/full-report/grids_e2c2091b.html

[11]. International Energy Agency, Electricity 2025 – Prices, 2025, https://www.iea.org/reports/electricity-2025/prices

[12]. A. Mellot, M. Pollitt, I. Staffell, Electrification, flexibility or both? Emerging trends in European energy policy, *Energy Policy*, Vol. 206, 2025, 114725.

[13]. Market monitor for demand side flexibility, *smartEn, LCP Delta*, 2025.

[14]. European Commission, Energy Performance of Buildings Directive, https://energy.ec.europa.eu/topics/energy-efficiency/energy-performance-buildings/energy-performance-buildings-directive_en

[15]. Eurostat, Renewable energy statistics, https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Renewable_energy_statistics

[16]. J. Vasheghani Farahani, Feasibility and performance of blockchain dual-purpose mining–heating systems, in *Proceedings of the 4th Blockchain and Cryptocurrency Conference (B2C'25)*, 2025, pp. 13-16.

[17]. Austrian Power Grid, Day-Ahead prices, https://markt.apg.at/en/transparency/cross-border-exchange/day-ahead-prices/

[18]. BITMAIN Technologies, S21 Pro specification, https://support.bitmain.com/hc/en-us/articles/31321354157593-S21-Pro-Specification

[19]. R. Sonderegger, Diagnostic tests determining the thermal response of a house, Technical Report LBL-7815, *Lawrence Berkeley National Laboratory*, 1978.

[20]. P. Michalak, Impact of air density variation on a simulated earth-to-air heat exchanger's performance, *Energies*, Vol. 15, Issue 9, 2022, 3215.

[21]. G. Fraisse, C. Viardot, O. Lafabrie, G. Achard, Development of a simplified and accurate building model based on electrical analogy, *Energy and Buildings*, Vol. 34, Issue 10, 2002, pp. 1017-1031.

[22]. ISO 24365:2022 – Radiators and convectors – Methods and rating for determining the heat output, *International Organization for Standardization*, 2022.

[23]. Blockchain.com, Bitcoin total hash rate (TH/s), https://www.blockchain.com/explorer/charts/hash-rate

[24]. Hashrate.no, WhitePool mining pool for cryptocurrencies, https://www.hashrate.no/pools/WhitePool

[25]. Yahoo Finance, BTC–EUR historical price data, https://de.finance.yahoo.com/quote/BTC-EUR/

[26]. I. Staffell, S. Pfenninger, N. Johnson, A global model of hourly space heating and cooling demand at multiple spatial scales, *Nature Energy*, Vol. 8, Issue 12, 2023, pp. 1328-1344.

[27]. European Commission, Reform of the EU electricity market design, 2023, https://ec.europa.eu/commission/presscorner/detail/en/ip_23_1591

IFSA Publishing

Blockchain and Cryptocurrency

Published by IFSA Publishing, S. L., 2026
https://bc-ifsa-journal.com/

Determining Intent in Smart Contracts: Identification Paths and the Calibration of Interpretive Mechanisms

Rui LU

Southwest Minzu University, School of Law, 16 Section 4, South Yihuan Road, Jiangxi Street, Wuhou District, Chengdu, Sichuan 610041, P.R. China
Tel.: + 86 13738188957
E-mail: Delovely33@qq.com

Received: 25 Feb. 2926 /Revised: 14 Mar. 2026 /Accepted: 18 Mar. 2026 /Published: 23 Mar. 2026

Abstract: As blockchain-based smart contracts gain wider adoption, their nature as code-driven mechanisms for automated contract formation and performance underscores the need for legal regulation. Determining their legal character is essential: the 'code theory' overlooks the core element of party intent, whereas the 'contract theory' better captures their essence, making it more normatively sound. Accordingly, a contractual lens provides the logical foundation for regulating smart contracts. Drawing on civil law's focus on subjective intent and common law's objective approach, this article distinguishes between scenarios of sufficient and constrained manifestation of intent, arguing that parties' deliberate activation of triggering mechanisms can constitute valid contractual assent – even without traditional negotiation or natural language. However, execution risks – semantic distortion, linguistic limitations of code, and over-rigid automation – necessitate a balanced framework integrating technical refinement, legal safeguards, and interpretive flexibility to reconcile efficiency with fairness in digital commerce. This article substantially extends an earlier version presented at B2C'2025 by providing a concrete identification path for determining intent in smart contracts and proposing specific risk mitigation strategies for different contract types.

Keywords: Blockchain technology, Smart contracts, Contract theory, Contractual intent, Semantic interpretation, Risk mitigation, Legal regulation.

1. Introduction

The term ‘smart contract’ was first coined in 1994 by American computer scientist, legal scholar, and cryptographer Nick Szabo, who defined it as ‘a set of digitally specified commitments, including the agreements by which the parties will fulfil these commitments’ [1]. As a computerised transaction agreement [2], smart contracts currently rely on blockchain technology as their core technical foundation, aiming to facilitate, verify, or execute contract negotiations or fulfilment without third-party interference. [3] They rely on decentralised blockchain technology composed of encrypted hash links, automatically recording, judging, and executing pre-set conditions through a node network in a distributed ledger. Based on the characteristics of the blockchain distributed ledger and peer-to-peer network, blockchain smart contracts can automatically complete the establishment, transmission, and fulfilment of instructions without human intervention.

Due to the IF-THEN trigger mechanism, smart contracts have a highly simplified establishment model. By pre-setting and inputting the conditions agreed upon by the parties regarding contract terms and performance methods in code form into the system, once the trigger conditions are met, the program automatically ‘outputs’ the performance results, achieving the integration and automation of contract establishment and performance. Smart contracts are currently widely applied in finance (e.g., Bitcoin, Ethereum, and other token-based methods for securities clearing and settlement, collateral management, derivative contracts, etc. [4]), healthcare (pre-authorisation processes for specific medical procedures), and the Internet of Things (e.g., automatic

purchasing when the quantity of food in a refrigerator falls below a pre-set value). They are increasingly being used as transaction tools [5]. In this coded transaction mechanism, contract content is expressed in programming language rather than natural (human) language, and its enforcement mechanism is highly automated. Therefore, there has been ongoing academic debate that smart contracts weaken or even deprive parties of the ability to fully express their true intentions during contract formation. Moreover, due to the ambiguity of natural language and the precision of code language, issues such as semantic conversion distortion, language span limitations, and excessive semantic migration may arise during contract execution. Therefore, the key to addressing the current debate over the expression of intent in smart contracts lies in seeking a recognition and calibration path that respects the characteristics of blockchain technology while incorporating appropriate technical improvements and legal regulations, thereby achieving a seamless integration between smart contracts and traditional contract law rules to better promote their healthy development.

Building upon the theoretical distinction established in Section 3, this article makes four distinctive contributions to smart contract regulation. First, it develops a novel theoretical framework differentiating between scenarios of sufficient and constrained intent manifestation in smart contracts, demonstrating that deliberate activation of triggering mechanisms constitutes valid contractual assent even in the absence of traditional negotiation or natural language. Second, it introduces a structured analytical framework for reconciling technical and legal language, offering a systematic methodology to determine whether on-chain actions satisfy the legal threshold for intent expression. Third, it proposes a classification schema for smart contracts based on revocability and modification costs, providing actionable guidance for selecting context-appropriate contract types across transactional landscapes. Finally, it delineates targeted risk mitigation strategies that synthesize technical refinement, legal safeguards, and interpretive flexibility – thereby advancing a balanced regulatory paradigm for digital commerce.

2. The Mechanism for Establishing Smart Contracts and Their Contractual Nature

2.1. Differences Between Smart Contracts and Traditional Contract Formation Mechanisms

The widespread concern within the academic community that smart contracts may restrict or even deprive parties of their true intentions primarily stems from the high degree of automation and automatic execution inherent in smart contracts, as well as the significant differences between their technical characteristics and traditional contract formation mechanisms.

Smart contracts supported by blockchain technology possess three key features: tamper-proofing, distributed transactions, and automatic execution. These three features complement one another, forming a delicate balance that collectively establishes a smart contract mechanism for automatic formation and execution without human intervention. Among these, tamper-resistance primarily manifests in blockchain's reliance on cryptography and consensus-based mechanisms [6], and once deployed, it cannot be altered [7]; distributed transactions primarily involve blockchain systems that typically do not rely on a central repository [8] to achieve decentralised information management and execution mechanisms. However, it is their high degree of automation and automatic execution that are considered the core elements hindering parties from reaching a genuine expression of intent in smart contracts.

The automated nature of blockchain technology is manifested in the fact that, once all information is collected, if appropriate parameters meet the conditions (typically expressed as "IF" code), the ledger can act as an electronic agent to replace human involvement in approval and management [8] automatically outputs the results of code execution (typically in the form of "THEN"). Smart contracts operate within the blockchain environment, strictly following the code without requiring approval at each step [9] to produce results. For example, the simplest "IF-THEN" conditional statement: if A sends some cryptocurrency to B, then B will send digital assets to A; and a more complex statement: "if A transfers some cryptocurrency to B every month, then A's advertisement will be placed on website C and remain visible for the next month; if A does not transfer cryptocurrency, then the advertisement will be automatically removed" [10]. The above two "IF-THEN" conditional statements serve as typical examples of smart contract automatic input and output scenarios. Once the "IF" condition is met, the blockchain technology code supporting the smart contract will automatically output the "THEN" conclusion based on the established condition – this is the process by which smart contracts automatically execute based on input code instructions under the guidance of blockchain technology.

It is precisely this highly automated execution feature and capability that allows smart contracts to bypass the traditional "bargaining" process where contract parties must repeatedly negotiate and communicate to establish a contract text, and begin execution after both parties' sign and the contract takes effect. This means that the parties signing a smart contract do not exhibit the outward appearance of "negotiation" during the negotiation phase, thereby exposing the risk of smart contracts restricting the parties' expressions of intent. It is therefore not surprising that such concerns have arisen in academic circles.

2.2. The Essence of Smart Contracts

In the application and development of smart contracts, the blockchain-supported automated transaction and enforcement mechanisms have given rise to "Code Supremacy" [11] "Code is Law" [12] "Code is Contract" [4] and similar perspectives, which contend that smart contracts themselves no longer require the support of law or traditional contract theory, as the power of code alone is sufficient to complete transactions and achieve autonomy. Conversely, scholars and practitioners with legal backgrounds contend that contracts governed solely by code will inevitably become unmanageable. Such contracts cannot fully represent or support the entirety of a transactional "contract." Instead, smart legal contracts – incorporating legal elements and support, and prioritising contractual attributes and essence – [4] are key to steering smart contracts towards sound development. The debate over whether the essence of smart contracts lies in code or contract has persisted unabated within both academic and practical circles.

Adherents of the "Code is Law" perspective contend that smart contracts are essentially programmable code executed on the blockchain [13], which can be understood as event-driven transactional agreements but are not regarded as legally binding contracts [10]. For instance, Ethereum is a prominent smart contract blockchain platform employing the Ethereum Virtual Machine (EVM) to support a Turing-complete scripting language, enabling developers to create and deploy diverse decentralised applications. Each participating node within the network runs the EVM [13] to complete the block validation workflow. Leveraging blockchain's unmanned, automated characteristics, the entire process executes autonomously without involving natural language interactions in the traditional legal sense. Thus, as an autonomous execution mechanism built upon blockchain platforms – integrating cryptographic safeguards and consensus mechanisms – smart contracts fundamentally provide the technical foundation for constructing trustless yet verifiable execution environments. Consequently, some scholars define them as "automatically executed electronic instructions drafted in computer code" [14] "computer protocols" [15] "A system for automatically moving digital assets" [16] "Self-executing computer code" [17] and so forth. Consequently, the "code is law" perspective contends that smart contracts should be understood primarily as a set of code systems or software programs operating on the blockchain, whose core function lies in the automatic execution of predefined operations rather than the traditional construction of legally binding contract texts using natural language.

Adherents of the "contract theory" contend that smart contracts should be interpreted as a "novel form of contract." The white paper Smart Contracts and Distributed Ledger – A Legal Perspective notes that for legal practitioners, contracts typically imply a specific legal obligation relationship, Smart contracts embody this particular legal relationship through code [4] constituting an expression of contractual form and representing a new development in contracts based on automation and standardisation [18]. Furthermore, scholar Chai Zhenguo holds a similar perspective, arguing that smart contracts, as predetermined rules (code), require deployment on blockchain programs. These rules take the form of contracts defined by software, regulating behavioural content within digital spaces. In a certain sense, they represent the expression of traditional contracts within the context of digital entities' behavioural content. In other words, smart contracts resemble the codification of electronic textual contracts, possessing all characteristics of standardised contracts while reducing corresponding regulatory and enforcement costs [19]. Consequently, the "contract theory" perspective indicates that smart contracts may be understood as self-executing agreements possessing legal efficacy as electronic information, with their technical characteristics aiding in overcoming legal impediments to their use as contract formation tools [20]. Smart contracts retain the attributes of traditional contracts, albeit expressed in the form of code.

Having presented the two aforementioned differing perspectives on the nature of smart contracts, it becomes apparent that the core divergence lies in how one approaches and perceives the technical medium of the smart contract itself, as well as how one views the inherent meaning of contracts and law. Adherents of the "Code Doctrine" maintain that smart contracts are entirely synonymous with code. This perspective places significant emphasis on the technical characteristics of smart contracts and blockchain technology, equating the core function of automatic execution with the entirety of a contract's purpose. It interprets the parties' intentions as a specialised form of code, implicitly negating and seeking to replace the existing contractual legal framework. However, this view distorts smart contracts into tools that alienate the parties' will, contravening the fundamental principle of freedom of contract. Its one-sided emphasis on the role of computer technology also overlooks essential elements in contract formation, effectiveness, and interpretation during transactions. Its rejection of the existing contract law system fundamentally misinterprets social fairness and natural justice. In terms of consequences, if the "code theory" prevails, smart contracts may become tools for parties to circumvent legal norms through technological means by leveraging code autonomy. Certain scenarios requiring the maintenance of fairness and justice – such as standard form contracts and grossly unfair terms – would lack legal recourse [21]. Parties' rights may also be compromised [22].

By comparison, the 'contract theory' perspective addresses the shortcomings inherent in the aforementioned 'code theory'. With technological advancement and the evolution of transactional forms, the manifestation of contracts has long transcended the traditional paradigms of 'paper documents' and 'natural language'. If the determination of whether an

agreement constitutes a contract remains confined to formal linguistic expression and offline negotiation processes, it fails to account for the diversity of contractual practices in the digital age. In Prime Sight Ltd v Lavarello, Lord Justice Tolsom noted that parties are generally free to enter into contracts on any terms they choose [23]. When individuals can reach agreements on their own terms, this aligns with the spirit and essence of "freedom of contract" and falls within the scope of traditional contract law [24]. This process of contract formation is not constrained by any prescribed form under traditional law; it suffices that the parties' intentions constitute an extension of their freely expressed will to be encompassed within the scope of a contract recognised by contract law [25]. As scholars of Raz have noted, contracts may be realised through diverse forms [26]. These encompass English language texts, nods or handshakes [27] and electronic data messages [28]. Consequently, in the digital age, smart contracts generated, transmitted, received, and stored on blockchains in the form of data code should analogously possess equivalent legal effect. The intentions expressed by parties through code on-chain, such as "automatically transferring assets to the counterparty upon fulfilment of condition A," though not written in traditional language, rely upon mutually selected platforms, rules, and parameter settings. Their legal function is no different in essence from that of a contract signed on paper. As the law continually adapts to new technologies, smart contracts should not be excluded from the contractual framework due to their programmed, automated, and coded nature. Rather, they should be regarded as a natural extension of traditional contracts within the digital era.

Therefore, the content of smart contracts should be viewed through the lens of the 'contract theory'. When considering a vehicle for transactions and expressions of intent, the crucial factor lies not in whether the method of expression is traditional, but whether the final form of the 'contract' constitutes the 'union' and 'intersection' of the parties' mutual intentions [29]. Conversely, the solution to bridging the gap between smart contract codification and traditional contractual forms lies in legislatively recognising the contractual nature of smart contracts, establishing legally enforceable standards for code conversion and readability, and enhancing mechanisms for translating encoded information into human-semantically comprehensible textual expressions.

3. A Theoretical Distinction: Sufficient Versus Constrained Manifestation of Intent

Building on the established "contract theory," smart contracts should be recognised as a new type of intelligent contract that requires semantic conversion to bridge the gap between code and traditional contract forms [30]. In semantic conversion, the core issue is how to preserve the core features of blockchain technology – such as the automatic execution of smart contracts – without diminishing or distorting the expression of intent, while also ensuring that corresponding legal norms can be clearly applied. Therefore, drawing on the understanding of contractual intent in both the traditional civil law and common law systems, exploring how to adequately express the parties' intent in smart contracts, as well as how to collect and transform such intent within smart contracts, has become a priority issue that must be addressed under contractual regulations.

In response to this challenge, this article posits a critical theoretical distinction between scenarios of sufficient manifestation of intent and constrained manifestation of intent in the context of smart contracts. This distinction serves as the foundation for determining when parties' actions – particularly the deliberate activation of triggering mechanisms – can be deemed to constitute valid contractual assent, even in the absence of traditional negotiation or natural-language terms.

3.1. The Content of Expressions of Intent in Traditional Contract Law

In traditional contract law, the interpretation of contractual intent forms the core of both the theory and practice of contract law, whether in the civil law system or the common law system. The common and core content lies in seeking the true intentions of the parties to maintain transaction security and fairness.

In the civil law system, represented by Germany and France, contract law is deeply influenced by Roman law and places greater emphasis on exploring the "inner intentions" of the parties. Under this system, the validity of a contract is deemed to stem from the parties' free will, and the expression of intent is regarded as an act of externalising one's inner intentions. In his General Principles of German Civil Law, Larenz emphasises that the interpretation of an expression of intent should first seek to uncover the true intentions of the party making the declaration, rather than merely adhering to the literal meaning of the words. He notes that if the party making the declaration has made an error or reservation, their true intentions should take precedence [31]. Additionally, the principle of reliance is introduced as an important supplement to balance the relationship between the declarant's true intent and the reasonable reliance of the counterparty. That is, the declarant's expression gives the counterparty reason to believe that a certain meaning exists, as external behaviour is a means of inferring the parties' true intentions [32], even if the declarant's true intention is not so, the declarant may still be bound by their declaration. This principle is particularly important in maintaining transaction security and protecting the reasonable expectations of the other party in cases of inadequate contract negotiations or unclear declarations. The introduction of the reliance principle has enabled the civil law

system to incorporate certain elements of objective interpretation, but its underlying principle remains the respect for the inner intent of the parties.

In contrast, the common law system, represented by England and the United States, places greater emphasis on the external appearance and objective reasonableness of transactions, stressing the theory of objective expression of intent. It holds that the validity of a contract does not depend on the parties' "secret inner intent," but rather on how a "reasonable third party" would understand the parties' words and actions in a specific context. Lord Dilock explicitly stated in Gibson v Manchester City Council that the formation of a contract is not based on the convergence of the parties' subjective intentions but on an objective assessment of the offer and acceptance [33]. Similarly, the same objective standard of consideration is reflected in Lord Clark's judgment in the Supreme Court case of RTS Flexible Systems Ltd v Molkerei Alois Müller GmbH & Co, where he noted that the formation of an "agreement" does not depend on their subjective mindset, but rather on the consideration of the information conveyed between them through words or actions [34]. This objective approach aims to promote certainty and predictability in commercial transactions, avoiding disputes arising from the difficulty of ascertaining the parties' subjective intentions. In the case of Investors Compensation Scheme Ltd v West Bromwich Building Society, Lord Hoffman outlined five principles for contract interpretation, emphasising that interpretation should consider "all relevant background information" [35] i.e., the "context," and to determine what understanding a "reasonable person" would have of the parties' words and actions given that background information.

3.2. The Content of the Expression of Intent in a Smart Contract

Whether it be the traditional civil law principle of "respecting the parties' genuine intentions" or the common law approach of "inferring intent from objective conduct," both theories are founded upon conventional contract law frameworks. However, as an emerging contractual model, smart contracts present unresolved questions regarding how parties' intentions may be ascertained through semantic transformation. Beyond the "contract theory," scholars hold divergent views: some contend that the "automatic formation" and "automatic execution" inherent in smart contracts inherently constrain parties' expressions of intent; others argue that regardless of negotiation, the core purpose of contract formation lies in "reaching agreement," meaning parties possess a sufficient expression of intent when consensus is attained.

Adherents of the "inadequate expression of intent in smart contracts" view contend that smart contracts suffer from structural deficiencies in core elements of intent formation, making it difficult to satisfy traditional contract law's requirements for sufficiency of expression. Scholar Florian Gamper explicitly notes that compared to conventional contracts, smart contracts often lack comprehensive coverage of transactional details and the parties' expected discretionary freedom in wording – for instance, the term "reasonable" – whose code logic struggles to fully map the intricate allocation of rights and obligations within complex transactions. He further emphasised that the value of natural language in contract formation lies not only in conveying information but also in preserving parties' discretionary and interpretative space, whereas the precision demanded by code renders it ill-suited to accommodate such flexibility [27]. Moreover, Michèle Finck's research reveals that the automatic execution characteristic of smart contracts may conflict with the "persistence of genuine intent" in expressions of will. Under the EU data protection legal framework, data subjects should retain the right to withdraw consent at any time; however, the irreversibility of smart contracts may lead to a disconnect between "formal consent" and "substantive changes in intent" [17].

Scholars who maintain that "smart contracts satisfy the requirement of full expression of intent" contend that although smart contracts employ code language in place of traditional natural language, they nonetheless fulfil the legal requirement for full expression of intent regarding the core elements necessary for contract formation. This conclusion finds corroboration within the theoretical frameworks of both major legal systems. From the perspective of the civil law tradition, German law requires that contract formation be based on the "mutual declaration of intent" by both parties. This principal manifests in smart contracts through the jointly established confirmation mechanism within the code. Scholar Ferreira's research indicates that the pre-execution negotiation of terms and parameter confirmation in smart contracts constitutes a phase where parties reach identical expressions of intent regarding core rights and obligations, exhibiting homogeneity with traditional contractual negotiation processes [36]. Durovic and Janssen note that the technical characteristics of smart contracts merely alter the medium of expression, without changing the fundamental principle that "the attainment of mutual consent constitutes the core of a contract." The manual confirmation stage prior to code execution ensures the authenticity of the expressed intent [37].

From the perspective of the Anglo-American "objective conduct standard," the code operations of smart contracts fully map the elements required for traditional contract formation. At the offer stage, pre-set contractual terms (such as the use of funds and reward mechanisms in crowdfunding smart contracts) constitute a "clear statement of willingness to enter into a contract on those terms" if their content is definite and accessible to specific parties. This aligns with the core characteristic of an offer defined by scholar McKendrick: the inclusion of an intention to be bound upon acceptance [24]. The acceptance phase may be fulfilled through on-chain actions. As established in Carlill v Carbolic Smoke Ball Co,

acceptance need not be expressed verbally; consent demonstrated through conduct remains valid [38]. In smart contracts, the user's act of transferring assets to a liquidity pool constitutes unconditional acceptance of the offer, made with full knowledge of its terms. This aligns with the traditional contractual logic that "performance constitutes acceptance" [24]. The consideration element manifests as the exchange of value between asset transfer and project returns, fulfilling Furmston's core definition: "one party's act or promise as the exchange condition for the other party's promise" [39]. Consequently, smart contracts fully embody the core elements of manifestation of intent through code execution. Their formation logic remains fundamentally consistent with traditional contracts regarding "authenticity of manifestation" and "element completeness"; differences in technological medium do not affect the determination of sufficiency of manifestation.

In summary, the fundamental dispute in academia regarding the sufficiency of parties' expressions of intent in smart contracts centres on how to interpret the semantic conversion mechanism between code language and natural language, and whether it can effectively embody fundamental elements of traditional contract law such as offer, acceptance, and consideration. Furthermore, differing interpretations of the formation of expressions of intent across legal systems further influence the recognition of such expressions in smart contracts. Building upon the foregoing analysis, this article will further explore the theoretical and practical feasibility of establishing "sufficiency of expression of intent" in smart contracts within the contractual framework. This aims to provide foundational institutional analysis and normative guidance for constructing smart contracts under a reasonable system of technological and legal regulation.

4. Path for Determining the Meaning of Smart Contracts

In the process of entering into a smart contract, although the parties did not follow the traditional path of negotiation involving "offer and acceptance," the expressions of intent they made still possess authenticity and sufficiency, and are not, as some scholars claim, "limited in expression." In fact, whether assessed under the standards of the civil law system or the common law system, parties entering into smart contracts possess sufficient expressions of intent, which are realised through a technical method distinct from traditional paradigms but equally legally binding. However, in terms of contract interpretation, the parties' expressions of intent may face certain interpretative risks due to the semantic transformation from "natural language" to "code language" and back to "natural language." Therefore, the improvement of the expression of intent mechanism in smart contracts should be achieved through the synergistic effects of technological innovation and legal regulations to jointly mitigate potential legal risks, thereby constructing a more comprehensive smart contract ecosystem and ultimately achieving a dynamic balance between transaction security and contractual freedom.

4.1. Exploring the Expression of Intent in the Context of contract formation

As scholars Kevin Werbach and Nicolas Cornell have pointed out, while smart contracts use the formulation "if... occurs, you will receive Bitcoin," which differs from the traditional contract statement "I will pay you Bitcoin," this difference merely reflects the technical specificity of their implementation [40].

Conversely, the opposing view that "smart contracts lack sufficient expression of intent" must be addressed from both the core requirements of expression of intent and the logic of contract practice. The simplification or omission of negotiation procedures does not necessarily negate the sufficiency of expression of intent. From the perspective of contract law principles, the sufficiency of expression of intent depends on "clarity of content" and "authenticity of intent," rather than the formal completeness of the negotiation process. The widespread use of standardised clauses in traditional contracts has established a mature solution: when clauses are pre-set and the other party lacks negotiation space, the law ensures the authenticity of the expression of intent through the "duty of reasonable disclosure" and the "rule of explicit consent." This logic of "negotiation simplification" and "procedural supplementation" shares an essential commonality with the operational mechanism of smart contracts: while the pre-set code of smart contracts omits the traditional contract's clause-by-clause negotiation process, parties achieve "explicit consent" through actions such as parameter confirmation and digital signature authorisation. The core requirements of this process are fundamentally consistent with the efficacy supplementation rules of standard terms.

Moreover, the code characteristics of smart contracts determine their inherent nature of "technologically fixed intent." The pre-set and confirmation stages of the code do not represent a "lack of intent," but rather the precise externalisation of intent in a digital environment. Parties' review of the code logic and setting of key parameters (such as the subject matter of the transaction and performance conditions) prior to deploying the contract are, in essence, substantive confirmations of the contract content, which are fully equivalent to the intent logic of "signing constitutes acceptance of all terms" in traditional contracts. As long as the parties' acceptance of the final terms is genuine and explicit, the simplification or even omission of the negotiation process does not constitute a defect in validity.

From a practical value perspective, insisting that smart contracts must reflect the "negotiation process" of traditional contracts not only contradicts the core efficiency requirements of digital transactions but also

ignores the innovative impact of technological characteristics on the form of expression of intent. The automatically executed terms of smart contracts, realised through code, are essentially a technological optimisation of traditional negotiation procedures rather than a negation of them. The immutability of their pre-set terms and confirmation actions actually provides stronger technological safeguards for the authenticity of the expression of intent. This form of expression of intent adapted to the digital ecosystem should be recognised inclusively by contract law, rather than being mechanically applied to the formal requirements of traditional transactions.

To summarise the comparative analysis above, Table 1 provides a clear framework illustrating how on-chain actions in smart contracts satisfy the core intent requirements of both major legal traditions.

Table 1. A Comparative Framework for Determining Intent in Smart Contracts.

Legal Tradition	Core Principle of Intent	Application to Smart Contracts	Justification/Example
Civil Law	Inner Intent / Reliance Principle	Explicit action as externalization of intent	Parties' joint confirmation of IF-THEN code logic is an external representation of their true, internal contractual intent. The act of deployment demonstrates a conscious choice.
Common Law	Objective Conduct	On-chain conduct as a verifiable sequence of legal acts	The series of transparent, immutable on-chain actions forms an objective record from which a rational observer can infer a clear intent to be bound by the agreement.

In summary, while the automated execution mechanism of smart contracts reduces the space for human intervention, it does not weaken the legal intent in the expression of behaviour. From the perspective of the "internal expression of intent" in the civil law system and the "objective behavioural assessment" standard in the common law system, it actually enhances external observers' ability to identify "contractual intent" and reduces interpretation disputes caused by linguistic ambiguity and semantic ambiguity in traditional contract formation. In some ways, it even enhances the clarity and certainty of the intent to form a contract, and can be argued to constitute sufficient expression of intent at the time of contract formation. From the general principles of contract law, the core of determining genuine intent lies in whether the act reflects a "legally binding intention," rather than the form of expression (written or oral). This intention is essentially the parties' right to make autonomous choices, i.e., the "intended legal effect" of the legal act. Such an act is typically the "expression" of this intent, i.e., the "manifestation of intent" [31]. As noted by scholar Atia, the pursuit of interests and reasonable reliance are sufficient to give rise to legal obligations [24]. This provides the basis for determining the validity of non-traditional form agreements.

From the perspective of the civil law system's principle of "respecting the parties' true intentions," smart contracts operate based on the parties' autonomous choices and explicit authorisation. The setting of "IF-THEN" execution conditions must first be jointly confirmed and agreed upon by both parties. In other words, smart contracts do not operate automatically out of thin air but are premised on the parties' explicit choices made after thorough consideration. When both parties make mutually consistent expressions of intent and jointly choose to automatically execute contract terms through triggering the "IF-THEN" mechanism, they have, in another form, expressed their joint acceptance of the constraints imposed by the content of their expressed "choice to have the smart contract execute the terms to achieve their intent." Without such explicit expressions of intent, the program itself cannot execute and will not produce any legal or factual binding effect. Therefore, the determination of intent should not be influenced by changes in form but should instead focus on the parties' genuine intent. Additionally, driven by technology, the objects of trust are undergoing subtle changes. Transactions are no longer based solely on trust in individuals but are increasingly based on trust in technology and its providers. In blockchain transactions, people choose to trust smart contracts and act according to their prescribed operational mechanisms. This is not only a form of technological trust but also institutional trust (institutional-based trust) [21]. Based on the principle of trust, both parties, aware of the automatic execution feature of smart contracts and with a consistent contractual purpose, adopt smart contracts to trigger an automatic fulfilment mechanism. This constitutes a reasonable trust in the automated execution mechanism and blockchain technology, thereby inferring their sufficient expression of intent. Therefore, under the civil law system, parties entering into smart contracts have a reasonable and sufficient expression of intent.

The "objective conduct standard" of the common law system provides the core framework for determining the intent of a smart contract: as long as the parties' conduct is sufficient for a rational third party with both legal and technical knowledge to reasonably infer their intent to be bound through objective evidence, the intent is deemed established. This standard was established in the core principle of

Smith v Hughes – the law focuses on the objective appearance of conduct rather than subjective, hidden intentions [41], and this logic applies equally to smart contract scenarios. During the formation of a smart contract, the sequence of parties' actions constitutes a recognisable trajectory of intent: invoking a standardised contract template (constituting an invitation to treat), setting key parameters for the subject matter of the transaction and performance conditions (constituting a definite offer), sending an execution instruction containing a digital signature to the blockchain system (constituting acceptance), and executing the blockchain smart contract code (constituting consideration). The temporal continuity and content certainty of these actions collectively form the objective appearance of "willingness to be bound by the contract." At this point, the "reasonable third party" should be defined as a professional with both contract law knowledge and blockchain technical expertise, who can determine through on-chain code whether the parties completed the operation with full knowledge that "code execution constitutes legal performance" – this aligns with the presumptive logic in traditional contracts that "signing constitutes acceptance of all terms of the text," but replaces "the ability to read and understand natural language text" with "the technical ability to interpret code logic" [42]. It is important to note that differences in the wording of smart contracts do not affect the substantive determination of the "objective behaviour" of the expression of intent.

4.2. Legal Risks and Regulatory Responses in Smart Contract Interpretation

4.2.1. Legal Risks

The ultimate purpose of establishing a smart contract is to fulfil the original objectives and intentions of the parties involved. However, due to the automated execution of smart contracts, which are not subject to any human intervention, unforeseen circumstances may arise during the execution process. In such cases, legal risks may arise in the interpretation of the contract, typically in the form of semantic distortion, language span limitations, and excessive semantic migration.

From the perspective of semantic distortion, this risk stems from the discrepancy between the code expression and the parties' true intentions under the automated execution mechanism of smart contracts. During the contract formation stage, if there is no effective mapping between the oral agreement and the coded terms, this may result in the execution outcome deviating from the parties' original intent. Typical scenarios include: smart contracts omitting core terms agreed upon in natural language, or code parameters differing substantially from oral commitments (e.g., in a gambling contract where a commitment to transfer a $200 prize is made, but the code is set to $20) [10]. Additionally, smart contracts may contain implicit terms not explicitly agreed upon by the parties, such as self-destruct conditions embedded in the code that result in the termination of rights without cause, even when the parties have not 45xternaliz an automatic termination clause [10]. Such discrepancies fundamentally stem from the inconsistency between the "expressive act" of intent and the "true inner intent" during contract execution, directly reducing the "precision of intent" in contract execution.

Unlike formal deviations caused by semantic distortion, the risks associated with language span limitations focus on content loss during semantic conversion. This stems from the inherent structural and functional differences between code language and natural language, meaning that 45xternaliz languages interpreted by machines cannot adequately capture the abstract nature of legal principles or the flexibility of institutional rules [43]. The code requirements of smart contracts necessitate the conversion of natural language into conditional statements, meaning that all contract terms are translated into the form "if A, then B." This conversion barrier results in certain expressions of intent being unable to be 45xternalized through code [44]. For example, the "IF-THEN" mechanism struggles to translate terms like "reasonable care" and "best efforts," which depend on value judgements [42], thereby creating an "interpretation vacuum" in the contract content.

From the perspective of semantic over-translation, the rigidity of smart contract code may lead to expansive interpretations of the parties' intentions, where interpreters go beyond the literal meaning of the code and the parties' reliance interests to impose additional rights and obligations on the contract. A typical example is the mechanical expansion of "IF-THEN" logic. For example, a supply chain smart contract stipulates that "if goods are delivered 30 days late, 5 % of the payment will be automatically deducted." When the actual delay is 29 days, the interpreter may apply the deduction rule by analogy, citing the "absolute nature of code logic," while ignoring the flexible interpretation space for "reasonable time limits" in traditional contract law [42]. While this interpretation aligns with the formal logic of the code, it exceeds the parties' original intent of punishing malicious breaches rather than minor defects, constituting an excessive expansion of the parties' intentions.

4.2.2. Smart Contract Debugging Path Based on Risk Assessment

The existence of risk means that a more effective risk adjustment mechanism must be established. Only under more comprehensive technical debugging and legal regulation can smart contracts truly fulfil their functions of automatic execution and realisation of the rights and obligations of the parties, thereby promoting fairness and freedom in commercial transactions. There are adjustment paths that can be optimised and improved from a technical perspective,

a legal regulatory perspective, and the coordination between the two.

First, from a technical perspective, to reduce misinterpretation of parties' intentions and situations where execution does not align with parties' original intentions, blockchain software developers can create "permissioned" or private blockchain prototypes [45] add "permission" commands to the blockchain, allowing human intervention for modifications [46]. Additionally, after smart contracts are negotiated and written as programs, legal provisions can be incorporated into the runtime layer to validate the contracts. The purpose of this step is to verify the feasibility of the program from a technical perspective and detect any program vulnerabilities or code errors. The runtime layer dynamically verifies the static contract data of the contract layer. If any content violating legal provisions or the parties' intentions is detected, the contract program may be rolled back, updated, or self-destructed as necessary to ensure compliance with legal provisions and the parties' intentions [47].

Secondly, from the perspective of legal regulation, taking the civil law system as an example, protecting the parties' reliance interests has become a widely recognised important principle in modern civil law [31], which is more conducive to the freedom and fairness of transactions. Therefore, when smart contracts result in semantic distortion or language barriers due to technical limitations, even if the automatic execution mechanism of smart contracts renders misunderstood expressions of intent irreversible, corresponding legal provisions for smart contracts should be established based on the fundamental principles of modern civil law. This would respect and recognise the parties' original intent at the time of contract formation, thereby ensuring that parties have a reasonable expectation of reliance when entering into smart contracts. Additionally, the "safe harbour" rule established in the U.S. Digital Millennium Copyright Act of 1998 can be referenced. When unforeseen "contract interpretation" issues arise that result in damage to the parties' interests, the platform provider should also assume the role of "risk controller," immediately taking measures to prevent further damage to interests. Otherwise, the platform provider should bear joint liability for the expanded portion of the loss. This rule can also be combined with the "pre-review" rule, which involves conducting a formal review of the code for vulnerabilities at the outset of contract formation to reduce legal and transactional risks [46] The pre-review of the Wyndham Hotels booking smart contract by the U.S. Federal Trade Commission (FTC) is a typical example of this [48].

Finally, from the perspective of the interpretative mechanism for reconciling conflicts between technical and legal language, it is evident that when the semantic interpretation of a smart contract conflicts with the parties' intentions, a general rule must be established to clarify which party's terms take precedence when there is a conflict between natural language expressions and smart contract provisions [10]. From the perspective of the interpretation of the essence of the parties' intentions – drawing on the civil law system's emphasis on subjective will and the common law system's objective conduct standard – the parties' manifested intent should serve as the primary basis for interpreting their contractual assent, even in code-based agreements [49]. Therefore, the terms expressed in natural language in the contract or those implied by it should take precedence over the terms of the smart contract used. Furthermore, if there is a third-party interpretation of the contract, special rules should be established based on the principle of fairness to provide effective remedies for all parties to the contract, including third parties, to protect their rights from being infringed upon due to inconsistencies between the written terms of the contract and oral agreements [40]. Furthermore, beyond the priority of interpretation, at the level of smart contract types, they can be categorised into "strong smart contracts" with prohibitive costs for revocation and modification, and "weak smart contracts" that can be relatively easily modified after contract execution. Depending on the complexity of rights and obligations, a contract type can be selected based on the legal relationship at the time of contract formation. In more complex contracting scenarios, "weak smart contracts" can be chosen. When there are improper interpretations of the parties' intentions after the contract is automatically executed, the contract can be appropriately modified based on respecting the parties' intentions at the time of contracting [50].

By improving technology, refining legal rules, and resolving conflicts between technology and contracts, we can explore risk regulation pathways for the interpretation of smart contracts. This approach ensures automated efficiency while respecting the fairness and freedom of transactions, preventing technological means from eroding contract fairness and party rights, and fostering the harmonious coexistence of technology and legal principles to promote transaction circulation.

4.3. A Jurisdiction-Agnostic Framework for Verifying Contractual Intent in Smart Contracts

This article establishes a jurisdiction-agnostic interpretive pathway for intent verification in smart contracts, resolving the critical gap between technical execution and legal sufficiency. The mechanism operates through three sequentially integrated phases, each addressing a distinct dimension of contractual intent formation that transcends conventional negotiation paradigms.

4.3.1. Technical Intent Mapping as Contractual Foundation

The pathway commences with rigorous verification of the code's fidelity to pre-execution

intent. This requires demonstrable alignment between the smart contract's "IF-THEN" logic and documented prior agreements, alongside precise correspondence between coded parameters and expressly stated contractual terms. Crucially, this phase rejects semantic drift – a pervasive risk where code semantics diverge from human intent due to linguistic constraints – by mandating explicit linkage between on-chain conditions and off-chain negotiation records.

4.3.2. Jurisdictional Calibration of Intent Manifestation

The second phase dynamically adapts to legal system paradigms. In civil law jurisdictions, it evaluates whether parties' on-chain actions (e.g., deployment, parameterization) constitute conscious externalization of subjective intent, treating code activation as equivalent to signature. For common law systems, it applies an objective reasonable technologist test: whether a third party possessing relevant blockchain expertise would reasonably infer binding intent from the same actions. This dual-axis assessment resolves the "code theory vs. contract theory" impasse by grounding intent in actional context rather than linguistic formality.

4.3.3. Risk-Contextualized Contractual Validation

The final phase synthesizes transactional context with risk taxonomy to determine legal sufficiency. It requires: (i) verification that execution aligns with original intent (not merely code logic), (ii) identification of specific risk vectors (semantic distortion, language span limitations, or excessive semantic migration), and (iii) context-aware mitigation selection – such as mandatory pre-deployment audits for B2B contracts with high modification costs, versus dynamic consent protocols for consumer-facing smart contracts. This transforms risk assessment from a passive safeguard into an active determinant of contractual validity.

5. Conclusions

With the rapid development of the digital economy, smart contracts – an innovative fusion of blockchain technology and contractual mechanisms – are gradually expanding the operational logic of traditional contractual models. How to determine the true intent of the parties involved in the highly automated and decentralised operation mechanism of smart contracts has become an urgent issue that the field of contract law must address. This article, from the perspective of the "contract theory," examines the theoretical frameworks and criteria for determining intent in both the civil law system and the common law system. It argues that while smart contracts do not follow traditional negotiation pathways or natural language expression methods, they can still be deemed to constitute sufficient and clear contractual intent through the signing trigger mechanism, thereby meeting the legal requirements for sufficient intent.

However, smart contracts still face legal risks during execution due to semantic distortion caused by differences between natural language and technical language, language span limitations, and excessive semantic migration. These risks are particularly evident when there is a disconnect between the code language and the parties' true intentions, potentially leading to execution deviations and legal consequences. Therefore, while acknowledging the validity of smart contracts, efforts should be made from three aspects: technological improvements, legal regulations, and the coordination of technology and law, to establish a "technology-law-institution" synergistic and interactive adjustment pathway.

In summary, as an extension of the concept of contracts in the digital age, the regulatory framework for smart contracts should not be confined to technical logic or formalism. Instead, it should respect the characteristics of blockchain automation while grounding itself in the fundamental principle of party autonomy and the theoretical foundation of contract law. Through the mutual adaptation of technology and law, it should provide a legal framework to support the construction of an efficient and trustworthy digital transaction order. This approach not only addresses the challenge of contract recognition "beyond form" in the digital age but also provides a theoretical foundation and institutional insights for the universal application of smart contracts in the context of global governance, while also establishing a reasonable legal framework for the recognition and use of smart contracts to ensure the safety and freedom of commercial transactions.

References

[1]. N. Szabo, Smart contracts: Building blocks for digital markets, *Extropy: Journal of Transhumanist Thought*, Vol. 16, Issue 2, 1996.

[2]. N. Szabo, Smart contracts, https://www.fon.hum.uva.nl/rob/Courses/InformationInSpeech/CDROM/Literature/LOTwinterschool2006/szabo.best.vwh.net/smart.contracts.html

[3]. S. Bourque, S. F. L. Tsui, A lawyer's introduction to smart contracts, in Scientia Nobilitat: Reviewed Legal Studies (A. Tyc, Ed.), *Scientia Nobilitat*, 2014, pp. 4-23.

[4]. Smart contracts and distributed ledger – A legal perspective, Whitepaper, *European Union Blockchain Observatory and Forum*, 2019.

[5]. Q. F. Xia, Legal nature analysis of smart contracts, *Eastern Law Review*, Issue 6, 2022, pp. 33-43 (in Chinese).

[6]. M. Chevalier, From smart contract litigation to blockchain arbitration: A new decentralised approach leading towards the blockchain arbitral order, *Journal of International Dispute Settlement*, Vol. 12, Issue 4, 2021, pp. 558-584.

[7]. S. Y. Lin, L. Zhang, J. Li, et al., A survey of application research based on blockchain smart

contract, *Wireless Networks*, Vol. 28, 2022, pp. 635-690.
[8]. Lowell Milken Institute for Business Law and Policy, Smart contracts: Is the law ready?, https://lowellmilkeninstitute.law.ucla.edu/wp-content/uploads/2018/08/Smart-Contracts-Whitepaper.pdf
[9]. M. C. Compagnucci, M. Fenwick, S. Wrbka (Eds.), Smart Contracts: Technological, Business and Legal Perspectives, *Hart Publishing*, Oxford, 2021.
[10]. N. Filatova, Smart contracts from the contract law perspective: Outlining new regulative strategies, *International Journal of Law and Information Technology*, Vol. 28, Issue 3, 2020, pp. 217-242.
[11]. P. De Filippi, A. Wright, Blockchain and the Law: The Rule of Code, *Harvard University Press*, Cambridge, 2018.
[12]. L. Lessig, Code 2.0: Law in the Networked Society (X. Li, W. W. Shen, Trans.), *Tsinghua University Press*, *Beijing*, 2018 (in Chinese).
[13]. H. Wang, J. Liu, J. Zhao, Blockchain smart contracts for decentralized matching of counterparties and automatic settlement of financial derivatives, *Blockchain: Research and Applications*, Vol. 6, Issue 4, 2025, 9.
[14]. Gibson v. Manchester City Council, [1979] UKHL 6, *UK*, 1979.
[15]. G. Governatori, F. Idelberger, Z. Milosevic, et al., On legal contracts, imperative and declarative smart contracts, and blockchain systems, *Artificial Intelligence and Law*, Vol. 26, Issue 4, 2018, pp. 377-407.
[16]. G. Wood, Ethereum: A secure decentralised generalised transaction ledger, Ethereum Project Yellow Paper, *Ethereum Foundation*, 2014.
[17]. M. Finck, Smart contracts as a form of solely automated processing under the GDPR, *International Data Privacy Law*, Vol. 9, Issue 2, 2019, pp. 78-94.
[18]. F. Lang, A new interpretation of contracts through smart contracts under blockchain technology, *Journal of Chongqing University (Social Sciences Edition)*, Vol. 27, Issue 5, 2021, pp. 169-182 (in Chinese).
[19]. Z. Chai, Contract law reflections on smart contracts under blockchain technology, *Guangdong Social Sciences*, Issue 4, 2019, pp. 240-241 (in Chinese).
[20]. E. A. Kirillova, V. V. Bogdan, I. B. Lagutin, et al., Legal status of smart contracts: Features, role, significance, *Jurídicas CUC*, Vol. 15, Issue 1, 2019, pp. 285-300.
[21]. Y. Wu, On the private law structure of smart contracts, *Jurist*, Issue 2, 2020, pp. 1-13 (in Chinese).
[22]. Q. Lin, Legal boundaries of smart contract code governance, *Hebei Jurisprudence*, Vol. 43, Issue 7, 2025, pp. 143-161 (in Chinese).
[23]. Cukurova Finance International Ltd v. Alfa Telecom Turkey Ltd, [2013] UKPC 22, [2014] AC 436, *UK*, 2013.
[24]. E. McKendrick, Contract Law (12th ed.), *Palgrave Macmillan*, London, 2020.
[25]. J. Gordley, Contract, property, and the will – The civil law and common law traditions, in The State and Freedom of Contract (H. Scheiber, Ed.), *Stanford University Press*, Stanford, 1998, pp. 79-83.
[26]. J. E. Penner, Voluntary obligations and the scope of the law of contract, *Legal Theory*, Vol. 2, Issue 4, 1996, pp. 325-357.
[27]. F. Gamper, A non-contractual approach to smart contracts, *International Journal of Law and Information Technology*, Vol. 31, Issue 3, 2023, pp. 231-252.
[28]. UNCITRAL, UNCITRAL Model Law on Electronic Commerce, https://uncitral.un.org/en/texts/ecommerce/modellaw/electronic_commerce
[29]. L. Duguit, Traité de Droit Constitutionnel (3rd ed., Vol. 1), *Éditions de la Société Anonyme du Recueil Sirey*, *Paris*, 1927 (in French).
[30]. C. D. Clack, V. A. Bakshi, L. Braine, Smart contract templates: Foundations, design landscape and research directions, https://www.cs.ucl.ac.uk/fileadmin/UCL-CS/research/Research_Notes/Smart_Contract_Templates.pdf
[31]. K. Larenz, General Principles of German Civil Law, Vol. 1 (X. Wang, et al., Trans.), *China Law Press*, Beijing, 2003 (in Chinese).
[32]. M. Planiol, G. Ripert, Traité Élémentaire de Droit Civil, *Librairie générale de droit et de jurisprudence*, *Paris*, 1949 (in French).
[33]. SAGE, Formation of contracts: Gibson v. Manchester City Council, https://sk.sagepub.com/cases/formation-of-contracts-gibson-v-manchester-city-council
[34]. RTS Flexible Systems Ltd v. Molkerei Alois Müller GmbH & Co KG, [2010] UKSC 14, [2010] 1 WLR 753, 2010.
[35]. Stocznia Gdanska SA v. Latvian Shipping Co, [1998] 1 WLR 896, *UK*, 1998.
[36]. A. Ferreira, Regulating smart contracts: Legal revolution or simply evolution?, *Telecommunications Policy*, Vol. 45, Issue 2, 2021, pp. 1-16.
[37]. M. Durovic, A. Janssen, The formation of smart contracts and beyond: Shaking the fundamentals of contract law, in Smart Contracts and Blockchain Technology: The Role of Contract Law (M. Cannarsa, P. Sirena, et al., Eds.), *Cambridge University Press*, Cambridge, 2019, pp. 1-27.
[38]. Carlill v. Carbolic Smoke Ball Co, [1893] 1 QB 256, *UK*, 1893.
[39]. M. P. Furmston, Cheshire, Fifoot & Furmston's Law of Contract (16th ed.), *Oxford University Press*, Oxford, 2012.
[40]. K. Werbach, N. Cornell, Contracts ex machina, *Duke Law Journal*, Vol. 67, 2017, pp. 313-382.
[41]. Business Bliss Consultants FZE, Smith v. Hughes, https://www.lawteacher.net/cases/smith-v-hughes.php
[42]. H.-W. Micklitz, O. Pollicino, A. Reichman, A. Simoncini, et al. (Eds.), Constitutional Challenges in the Algorithmic Society, *Cambridge University Press*, Cambridge, 2021.
[43]. L. Efimova, O. Sizemova, A. Chirkov, Smart contracts: Between freedom and strict legal regulation, *Information & Communications Technology Law*, Vol. 30, Issue 3, 2021, pp. 331-350.
[44]. Y. Ni, Civil law analysis, application, and implications of smart contracts under blockchain technology, *Journal of Chongqing University (Social Sciences Edition)*, Issue 3, 2019, pp. 178-190 (in Chinese).
[45]. M. N. Temte, Blockchain challenges traditional contract law: Just how smart are smart contracts?, *Wyoming Law Review*, Vol. 19, Issue 1, 2019, pp. 185-210.
[46]. L. Han, L. Cheng, Legal deconstruction and risk mitigation of blockchain smart contracts, *Learning and Practice*, Issue 3, 2022, pp. 54-62 (in Chinese).
[47]. X. D. Li, S. Y. Ma, Research on blockchain smart contracts under the contract section of the civil code,

Journal of Shanghai Normal University, Vol. 49, Issue 5, 2020, pp. 58-69 (in Chinese).
[48]. Federal Trade Commission v. Wyndham Hotels & Resorts, LLC, No. 14-3514 (3d Cir. 2015), *USA*, 2015.
[49]. R. Lu, Legal approaches to determining the expression of intent in smart contracts, in *Proceedings of the 4th Blockchain and Cryptocurrency Conference (B2C'25)*, 2025, pp. 97-99.
[50]. M. Raskin, The law and legality of smart contracts, *Georgetown Law Technology Review*, Vol. 1, Issue 1, 2017, pp. 305-341.

IFSA Publishing

Blockchain and Cryptocurrency
Published by IFSA Publishing, S. L., 2026
https://bc-ifsa-journal.com/

Private Blockchain Anonymization-Deanonymization System Preserving Anonymity for the Net

Eligijus Sakalauskas, Antanas Bendoraitis, Syeda Roushan Arshid, Aušrys Kilčiauskas, Aleksejus Michalkovič, Lina Dindienė, Kęstutis Lukšys
Kaunas University of Technology, Scientific Group of 'Cryptography and blockchain systems',
Studentų str. 50, Kaunas, Lithuania
Tel.: + 370 698 784 77
E-mail: eligijus.sakalauskas@ktu.lt

Received: 22 Dec. 2025 /Revised: 19 Mar. 2026 /Accepted: 20 Mar. 2026 /Published: 23 Mar. 2023

Abstract: In a private blockchain, it is essential to provide not only anonymization of users but also deanonymization for certain parts of the Network, such as audit organizations, investment companies, and others, while maintaining anonymity for the other parts of the Network. Traditionally, the anonymization in private blockchain is performed by implementing ring signatures based on elliptic curve cryptography. The deanonymization problem is not considered in connection with the anonymization problem. In this paper, we present a unified anonymization-deanonymization system. Anonymization is based on an arbitrary set of single-user addresses generated independently and secretly. These addresses correspond to the user's private and public keys used for transaction creation based on the Schnorr signature. Deanonymization is based on the Schnorr multi-signature scheme and Non-Interactive Zero Knowledge proof. The presented solution is an integration of Schnorr signature, Schnorr multi-signature, and Schnorr-based Non-Interactive Zero Knowledge Proof. Security and effectiveness analysis are presented.

Keywords: Private blockchain, Transaction's anonymity, Ring signatures, Anonymization, Deanonymization.

1. Introduction

Anonymization is a relevant topic in private blockchain, alongside other confidentiality requirements. For example, this problem in the Monero blockchain is addressed by Ring Signatures (RS) based on Elliptic Curve Cryptography (ECC) [1-3]. On this background, the Elliptic Curve Digital Signature Algorithm (ECDSA) was created and standardised.

Nevertheless, such anonymity can be interpreted as partial anonymity on the Network, since the actual transaction signer is inevitably among the other ring signers, who are members of the ring created by the actual transaction signer. A shortcoming of Monero RS-based anonymity is that it should be performed during transaction signing, which requires additional computational resources. To increase this kind of anonymity, it is required to increase the number of ring members. At the same time, computational resources are increasing. Therefore, there is a need to develop a more anonymous transaction system that requires fewer computational resources.

However, there are also situations in which the opposite process, such as deanonymization, is required. For example, income from anonymous transactions should be redirected to the Investment Company (IC) by revealing the identity of the cryptocurrency owner.

In the case of Monero RS, deanonymization does not require additional computational resources because all funds are held in the transaction creator's account.

In this paper, we present an alternative approach to achieve actual anonymity and a more effective realization based on the classical Schnorr signature together with a deanonymization method based on the Schnorr multi-signature. Deanonymization requires providing a small amount of additional information to IC, together with the required signature for this additional information.

We consider a blockchain system based on the Unspent Transactions Output (UTxO) paradigm [4].

In Section 2, we present an anonymization-deanonymization method based on RS and used in the Monero blockchain. In Section 3, the proposed method of anonymization-deanonymization based on Schnorr signature and Schnorr multi-signature is presented. Section 4 is supported by example, and Section 5 includes security considerations. In Section 6, the conclusions are presented.

2. Anonymization-deanonymization based on Ring Signatures (RS)

This anonymization method is used in the Monero blockchain, providing the transaction's creator with anonymity. RS is a type of digital signature that can be performed by any member of a set of users who each have a pair of public and private keys [5]. Therefore, a message signed with an RS is endorsed by someone in a particular set of people. One of the security properties of a ring signature is that it should be computationally infeasible to determine which of the set's members' keys was used to produce the signature. In RS, there is no way to revoke the anonymity of an individual signature, and any set of users can be used as a signing set without additional setup.

Ring signatures are signer-ambiguous. In a ring signature scheme, there are no prearranged groups of users, there are no procedures for setting, changing, or deleting groups, there is no way to distribute specialized keys, and there is no way to revoke the anonymity of the actual signer (unless he decides to expose himself).

The only assumption is that each member is already associated with the public key of some standard signature scheme.

To produce a ring signature, the actual signer declares an arbitrary set of possible signers that includes himself and computes the signature entirely by himself using only his private key and the others' public keys.

In particular, other possible signers may have chosen their private keys solely to conduct e-commerce over the internet. They may be completely unaware that a stranger uses their public key to produce a ring signature on a message they have never seen and would not wish to sign.

A set of possible signers is named a ring. The ring member who produces the actual signature is the signer, and each of the other ring members is a non-signer.

A ring signature scheme is set-up free: The signer does not need the knowledge, consent, or assistance of the other ring members to put them in the ring - all he needs is knowledge of their regular public keys. The size of the signature depends on the number of ring members.

Verification must satisfy the usual soundness and completeness conditions. It is required that the signatures be signer-ambiguous in the sense that the verifier should be unable to determine the identity of the actual signer in a ring of size r with probability greater than $\frac{1}{r}$.

The construction based on ECDSA provides unconditional anonymity in the sense that even an infinitely powerful adversary with access to an unbounded number of chosen-message signatures produced by the same ring member cannot guess his identity with any advantage, and cannot link additional signatures to the same signer.

In Monero, RS are implemented using elliptic curve cryptography. These curves are defined over a finite field $\mathbb{F}p$, where p is a prime number. The field formed by the set $\{0, 1, 2, \dots, p-1\}$, with arithmetic operations $(+, \cdot)$ calculated $(mod\ p)$.

Typically, elliptic curves are defined as the set of points (x, y) satisfying a *Weierstraß* equation:

$$y^2 = x^3 + ax + b \quad \text{where } a, b, x, y \in \mathbb{F}\text{p}. \tag{1}$$

However, the cryptocurrency Monero uses a special curve known to offer improved security over other commonly used NIST curves, as well as excellent performance of cryptographic primitives. The curve used belongs to the category of so-called *Twisted Edwards* curves, which are commonly used [1]. In EC, the special points addition is defined, creating an abelian group of these points.

A generator G of EC is a point on the curve such that for every other point P in EC there exists k such that $P = kG$.

Public key cryptography algorithms can be devised in an analogous way to modular arithmetic.

Let k be a randomly selected number satisfying $1 < k < N_{EC}$, where N_{EC} is a number or EC points. Number k represents a *private key*. Then the corresponding *public key* $K = kG$, where K is an EC point.

Typically, a cryptographic signature is performed on a cryptographic hash of a message.

2.1. Signature

Assume that Alice has the private/public key pair (k, K). To sign an arbitrary message M univocally, she could execute the following steps [2]:

1. Calculate a hash of the message using a cryptographically secure hash function, $h = H(M)$

2. Generate a random integer r such that $1 < r < N$ and compute $P = (x, y) = rG$.
   If $x = 0$ generate another random integer.
3. Calculate $s = r^{-1}(h + xk)(mod\ N)$. If $s = 0$, then go to the previous step and repeat
4. The signature is $\sigma = (x, s)$.

2.2. Verification

Any third party can verify the signature by calculating

$$u_1 = s^{-1}h \tag{2}$$

$$u_2 = s^{-1}r \tag{3}$$

$$Q = u_1G + u_2K \tag{4}$$

2.3. Correctness

The correctness of the scheme can be derived from the fact that

$$\begin{aligned} Q &= u_1G + u_2K \\ &= s^{-1}hG + s^{-1}rkG \\ &= s^{-1}(h + xk)G \end{aligned} \tag{5}$$

Since $s = r^{-1}(h + xk)$, it follows that $r = s^{-1}(h + xk)$, whereby it is proved that

$$Q = rG \tag{6}$$

In the Monero blockchain, the anonymization is performed using ring signatures (RS) presented in [3].

2.4. Anonymization

The simplified version of this scheme is presented below, assuming that we have the same number of keys for any value of the first index i. Assume that we have a set of public keys $\{K_{i,j}\}$ for $i \in \{1, 2, \dots, n\}$ and $j \in \{1, 2, \dots, m\}$. Furthermore, we also assume that for each *i,* there is an index π_i such that the signer knows the private key k_{i,π_i} corresponding to K_{i,π_i}.

In what follows, we will use *M* for the hash of the message concatenated with keys $K_{i,j}$.

2.5. Ring Signature

1. For each $i = 1, \dots, n$:
   (a) generate a random value $\alpha_i \in_R \mathbb{Z}_q$
   (b) set $c_{i,\pi_i} = H_n(M, \alpha_i G, i, \pi_i)$
   (c) for $j = \pi_i + 1, \dots, m - 1$ generate random numbers $r_{i,j} \in_R \mathbb{Z}_q$ and compute,

$$c_{i,j+1} = H_n\big(M, r_{i,j}G - c_{i,j}K_{i,j}, i, j\big) \tag{7}$$

2. For $i = 1, \dots, n$ generate random numbers $r_{i,m} \in_R \mathbb{Z}_q$ and compute

$$c_1 = H_n\big(r_{1,m}G - c_{1,m}K_{1,m}, \dots, r_{n,m}G - c_{n,m}K_{i,m}\big) \tag{8}$$

3. For $i = 1, \dots, n$:
(a) for $j = 1, \dots, \pi_i - 1$generate random numbers $r_{i,j} \in_R \mathbb{Z}_q$ and compute

$$c_{i,j+1} = H_n\big(M, r_{i,j}G - c_{i,j}K_{i,j}, i, j\big) \tag{9}$$

Here, we interpret references to $c_{i,1}$ as c_1, see the previous step.
(b) set $r_{i,\pi_i} = \alpha_i + k_{i,\pi_i}c_{i,\pi_i}$
The ring signature is

$$\sigma_{RS} = \big(c_1, r_{1,1}, r_{1,2}, \dots, r_{1,m}, \dots, r_{n,m}\big) \tag{10}$$

2.6. Verification

As in the previous section, let m denote the hash of the message to be signed, together with the corresponding set of signing keys. The verification of a given signature is performed as follows:

1. For each $i = 1, \dots, n$ and $j = 1, \dots, m$ compute:

$$R'_{i,j+1} = r_{i,j}G - c'_{(i,j)}K_{i,j} \tag{11}$$

$$c'_{i,j+1} = H_n\big(M, R'_{i,j+1}, i, j\big) \tag{12}$$

   Interpret any $c'_{i,1}$ as c_1
2. Compute $c'_i = H\big(R'_{1,m}, \dots, R'_{n,m}\big)$

The signature will be valid if $c'_1 = c_1$.

2.7. Correctness

1. For $j \neq \pi_i$ and for all i, we can readily see that $c'_{i,j+1} = c_{i,j+1}$.
2. When $j = \pi_i$, for all i

$$\begin{aligned} R'_{i,j+1} &= r_{i,j}G - c'_{i,j}K_{i,j} \\ &= (\alpha_i + ki, \pi_i c'_i, \pi_i)\ G - c'_i, \pi_i K_i, \pi_i \end{aligned} \tag{13}$$

In other words, $c'_{i,\pi_i} = H_n(M, \alpha_i G, i, \pi_i) = c_{i,\pi+1}$. Therefore, we can conclude that the verification step identifies correctly valid signatures.

2.8. Deanonymization

Deanonymization is performed in a trivial way. Alice simply transfers a required sum from her account to IC and signs her transaction with her private key without using any other public keys required for RS creation.

2.9. Efficiency

Most of the proposed algorithms have an asymptotic output size $O(n)$, i.e., the size of the resulting signature increases linearly with the size of

the input (number of public keys) corresponding to the number of chosen ring members. The computational cost is estimated by identifying the operations that require significant computational resources. Such an operation in ECC is a multiplication of EC point by the number, i.e. rG, where r is a number in the field $\mathbb{F}_p$, and G is an EC generator. This operation requires a significant computational resources and, in some sense, is analogous to the exponentiation operation used in the classical ElGamal cryptosystem The number of EC point multiplications by scalars in $\mathbb{F}_p$ required for RS implementation is denoted by N_{ECO}. According to the construction presented in Section 2.5, this number depends on the number n of ring members and is equal to

$$N_{RS} = 4n^3. \quad (14)$$

As we see, to provide more anonymity, a greater number of ring members is needed, and cubically more of EC point multiplication operations of EC points are required.

The other drawback is that these operations must be carried out during the transaction execution, i.e., online.

3. Anonymization-deanonymization based on the Schnorr Signature Approach

This approach is based on Schnorr signature for anonymization [6] and Schnorr multi-signature for deanonymization [7]. As usual, every user has a private and public key pair, which we denote by $(\mathrm{PrK}, \mathrm{PuK})$ for signature creation and verification, respectively. Moreover, every user of blockchain can generate as many $(\mathrm{PrK}, \mathrm{PuK})$ pairs as required.

Let p be a large prime of order 2^{2048}, and $\boldsymbol{Z_p^*}$ is a multiplicative group with operation $mod\ p$. Let $\boldsymbol{G_q}$ be a subgroup of $\boldsymbol{Z_p^*}$ of order q, where q is prime. Then all elements in $\boldsymbol{G_q}$ except 1 are generators. Let $g \in \boldsymbol{G_q}$ be a generator in $\boldsymbol{G_q}$, then the order of the generator g is q.

In the Schnorr signature scheme, public parameters (PP) consist of the following public parameters $PP = (p, q, g)$. In our approach, user anonymization is achieved by generating as many blockchain accounts as required. This simple procedure therefore conceals the true identity of the account owner.

In the Schnorr signature scheme, a key pair $(\mathrm{PrK}, \mathrm{PuK})$ is generated using PP in the following way:

$$\mathrm{PrK} = x \leftarrow rand(\boldsymbol{Z_q}), \quad (15)$$

$$\mathrm{PuK} = g^x\ mod\ p = a. \quad (16)$$

In (15), private key x is generated at random in the semiring $\boldsymbol{Z_q} = \{0, 1, 2, \dots, q-1\}$, where addition and multiplication operations mod q are defined. Moreover, all arithmetical operations in the exponents can be reduced modulo q. Therefore, all exponents are in the set $\boldsymbol{Z_q}$.

The public key a is calculated using a discrete exponential function (DEF), which is considered as a classical conjectured one-way function (OWF). This function is used to ensure security against cryptanalysis by classical computers when public parameters p and q are chosen sufficiently large, i.e., p has a 2048-bit length, and q is more than 1024-bit length.

Assume that the blockchain account address is computed traditionally, i.e., using a secure (collision-resistant) cryptographic function. H-function. This function is specified in the blockchain architecture. Symbolically, we denote it by

$$Addr = H(\mathrm{Puk}). \quad (17)$$

As we see $\mathrm{PrK} = x$ and $\mathrm{PuK} = a$ are linked with the Alice account, we name it the main account.

3.1. Anonymization

Let Alice wish to be anonymous on the Net, except to her business partners, by generating at random N pairs of public and private keys to create N anonymous accounts. First, she generates a set of $N-1$ private key $\mathrm{PrK}_{n-1} = x_{n-1}, n = 1, 2, \dots, N$. The last private key is computed by the expression

$$x_N = x - x_1 - x_2 - \dots - x_N - 1\ mod\ q. \quad (18)$$

Hence, evidently, we have:

$$\mathrm{PrK} = x = x_1 + x_2 + \dots + x_N\ mod\ q. \quad (19)$$

Then, using (16) and (17), public keys and addresses are computed.

The set of addresses and public keys is known to the nodes of the Net. But the Net does not know that the set of addresses $\{Addr_n\}$ and the corresponding set of public keys $\{a_n\}, n = 1, 2, \dots, N$ belong to Alice. To ensure anonymity, Alice uses her anonymous addresses to conduct transactions between her partners. All transactions are signed using the Schnorr signature scheme.

To be self-contained, we present the Schnorr signature scheme to sign the transaction Tx_i

$$j_i \leftarrow \mathrm{rand}(\boldsymbol{Z_q}). \quad (20)$$

$$r_i = g^{j_i} mod\ p. \quad (21)$$

$$h_i = H(\mathrm{Tx}_i || r_i). \quad (22)$$

$$s_i = j_i + x_i \cdot h_i\ mod\ q. \quad (23)$$

Alice's signature on Tx_i is $\sigma_i = (r_i, s_i)$.

The verification of the signature σ_i is performed with Alice $\text{PuK}_i = a_i$ by verifying the following identity

$$g^{s_i} \bmod p = r_i \cdot (a_i)^{h_i} \bmod p. \quad (24)$$

All nodes in the net can verify the validity of the transaction signature by using (24) and acquiring the public key a_i. This verification is secure under the DL assumption in the random oracle model.

For more clarity, let us assume that Alice created two anonymous account addresses $\{Addr_1, Addr_2\}$ as indicated in Fig.1, corresponding to $(PrK_1 = x_1, PuK_1 = a_1)$ and $(PrK_2 = x_2, PuK_2 = a_2)$ respectively.

Let Alice decide to invest some income in her addresses $\{Addr_1, Addr_2\}$ to the Investment Company (IC). Then she creates two transactions Tx1, Tx2, where investment sums are outputs (expenses) of these transactions. These transactions are signed by Schnorr signature using Alice's private keys $PrK_1 = x_1$ and $PrK_2 = x_2$ generated by her in advance. Using (20)-(23) two signatures are computed:

$$\sigma_1 = (r_1, s_1); \ \sigma_2 = (r_2, s_2). \quad (25)$$

IC, together with the received sums in Tx1 and Tx2, verifies signatures on these transactions using (24). But IC does not know that these transactions are sent from Alice's anonymous accounts. IC can also believe that it receives funds from the different investors.

3.2. Deanonymization

After Alice sends Tx1 and Tx2 to the IC, she must prove that this amount of money belongs to her and that she is an investor. For this purpose, Alice uses Schnorr-based Non-Interactive Zero Knowledge Proof (NIZKP) to demonstrate that she knows the $PrK = x$, as defined in (19), without revealing it. The general idea of an NIZKP is illustrated in Fig. 1.

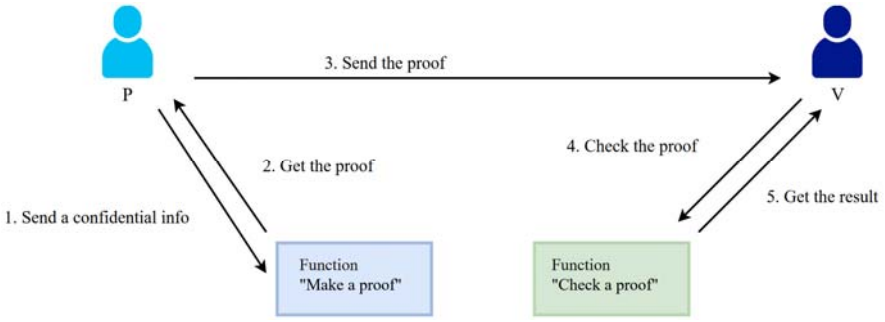

Fig. 1. The general idea of NIZKP.

The idea of a Schnorr-based NIZKP is similar to that of the signature scheme, in that the three-round interactive Schnorr identification protocol is transformed into a single-round non-interactive proof using the Fiat-Shamir heuristic [11].

Following (16) and (19), when $N = 2$, the following equations hold

$$x = x_1 + x_2 \bmod q \quad (26)$$

$$g^x = g^{x_1 + x_2} = a_1 \cdot a_2 = a \bmod p \quad (27)$$

Then, having two signatures (σ_1, σ_2)Alice computes the multiplication of these signatures in a special way, yielding a Schnorr multi-signature. This special multiplication operation, we denote by $*$. As a result, the Schnorr multi-signature is computed in the following way:

$$\begin{aligned} \sigma_{12} &= (\sigma_1 * \sigma_2) = (r_1, s_1) * (r_2, s_2) = \\ &= (r_1 \cdot r_2 \bmod p, s_1 + s_2 \bmod q) = \\ &= (r_{12}, s_{12}). \end{aligned} \quad (28)$$

The Schnorr multi-signature is secure under the DL assumption in a random oracle model. In other words, any polynomial-time adversary cannot produce a forgery with non-negligible advantage based on the challenger's responses to hash and signing queries. This property is known as Existential Unforgeability under Chosen-Message Attack (EUF-CMA).

However, Schnorr multi-signature can be created by anyone on the Net, and, therefore, it itself cannot be a proof of Alice's transactions Tx1 and Tx2. To prove that she is a creator of Tx1 and Tx2, Alice is using Schnorr-based NIZKP to the verifier IC.

To start with, Alice generates at random a number $j \leftarrow rand(\mathbb{Z}_q)$ and computes

$$r = g^j \bmod p \quad (29)$$

Using a collision-resistant H-function, Alice computes the following h-value playing the role of the challenge in Schnorr identification:

$$h = H(Tx1, Tx2, \sigma_1, \sigma_2, \sigma_{12}, a_1, a_2, a || r) \quad (30)$$

She now calculates

$$s = j + x \cdot h \tag{31}$$

Then NIZKP of the prover data we express as follows:

$$\eta = (r, s) \tag{32}$$

The prover Alice sends the following data to the verifier IC:

$$\{Tx1, Tx2, \sigma_1, \sigma_2, \sigma_{12}, a_1, a_2, a, \eta\} \rightarrow IC \tag{33}$$

IC, after receiving this data, performs the following verifications.

1. Are transactions $\{\mathrm{Tx}_1, \mathrm{Tx}_2\}$ on the account of IC.
2. Verifies signatures (σ_1, σ_2) on these transactions using (24).
3. IC Schnorr multi-signature σ_{12} by checking the validity of the following identity

$$g^{s_{12}} mod\ p = r_{12} \cdot (a_1)^{h_1} \cdot (a_2)^{h_2} mod\ p \tag{34}$$

4. IC verifies if

$$a_1 \cdot a_2 = a\ mod\ p \tag{35}$$

5. Then, IC verifies if NIZKP is correct using $\eta = (r, s)$ and checking the identity

$$g^s = r \cdot (a)^h mod\ p \tag{36}$$

The last equation can be rewritten in the form

$$g^{s_1+s_2}\ mod\ p = r_1 \cdot r_2 \cdot (a_1 \cdot a_2)^h \tag{37}$$

If (35) or (37) holds, then Alice proved to IC that she knows her $PrK = x$ satisfying (26) without revealing its value. Consequently, Alice proved that transactions Tx1 and Tx2 belong to her and that she is an investor. This scheme is summarized in Fig. 2.

Schnorr-based NIZKP is standardized by the RFC 8235 document for both finite fields and elliptic curves. It is secure under the DL assumption in the random oracle model.

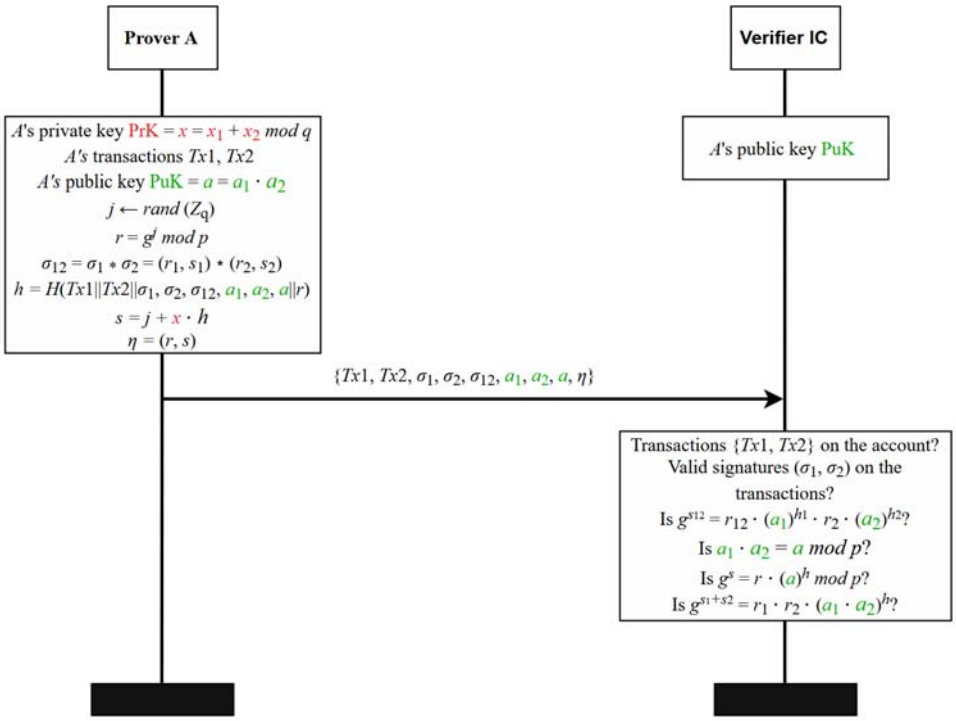

Fig. 2. Deanonymization scheme in our proposal.

4. Example

Let us consider the case when Alice has an account address $Addr$ created according to (17). Then, Alice created two account addresses $\{Addr_1, Addr_2\}$ with corresponding private and public keys to ensure anonymity, as it is depicted in Fig. 3.

According to the Unspent Transactions Output (UTxO) paradigm, Tx1 is realised in $Addr$, with two cryptocurrency incomes $m_1 = 2000$ and $m_2 = 3000$. Hence, the total sum of incomes is $m_{12} = 5000$. The same sum $m_{34} = 5000$ is spent, and is denoted by expenses $m_3 = 1000$ and $m_4 = 4000$. The amount m_3 is transferred to Emily (E), whereas the amount m_4, representing expenses, is transferred by Alice to her anonymous address $Addr_1$.

Let Alice receive two sums $m_6 = 1000$ and $m_7 = 2000$ in the other anonymous address $Addr_2$ in transaction Tx2. She transfers all the received sum $m_8 = 3000$ to IC by this transaction.

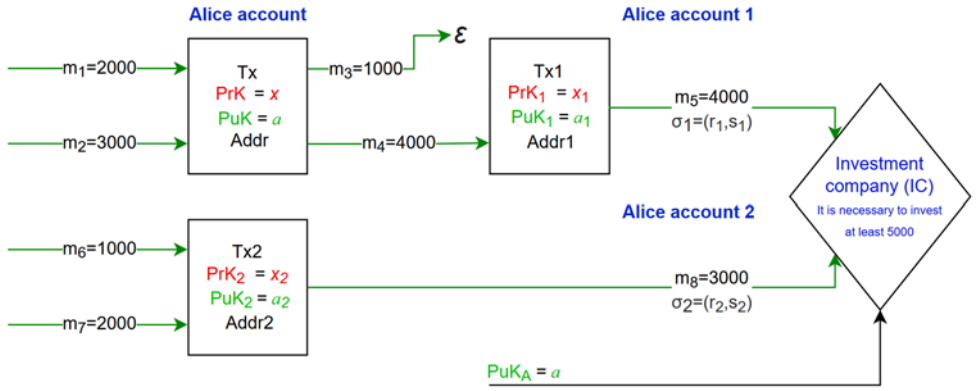

Fig. 3. Anonymization-Deanonymization scheme in a Private Blockchain.

To prove ownership of m_5 and m_8, a combination of Schnorr signatures, Schnorr multi-signatures, and NIZKPs is used to perform the anonymization and deanonymization processes, as presented in Section 3.

5. Security Considerations

The security of anonymization is based on the security of the Schnorr signature scheme.

The security of the Schnorr signature relies on the following two pillars:

- The discrete logarithm (DL) assumption: given a generator g and the public key $a = g^x mod\ p$, no efficient adversary can gain an advantage in obtaining the private key x.
- Fiat-Shamir and random oracle assumptions: the Schnorr signature is obtained from the Fiat-Shamir transform of the Schnorr sigma identification protocol, and uses a collision-resistant H-function, which acts as a random oracle to ensure protection against existential forgeries.

As usual, we understand the adversary's advantage as a probability to improve the random guess of a discrete logarithm value, i.e.

$$AdvDL(g,a,p) = \left|\Pr(\hat{x} = x) - \frac{1}{q}\right|$$

When we say that a certain cryptographic primitive is secure under the DL assumption, we mean that the advantage $AdvDL(g,a,p)$ is negligible, i.e., for any fixed $d > 0$, the advantage $AdvDL(g,a,p)$ tends to zero faster than n^d tends to infinity, where n is the size of q in bits, after a certain n_0, i.e. $\lim_{n\to+\infty}(AdvDL(g,a,p)\cdot n^d) = 0$.

An effective adversary that forges Schnorr signatures can also efficiently compute discrete logarithms in large groups. In other words, any adversary who can efficiently forge a Schnorr signature must have a non-negligible advantage $AdvDL(g,a,p)$. Therefore, this adversary can also be used to solve for x in (16). According to Shor's results [9], such adversaries are possible on quantum computers. Hence, existing cryptographic methods should be replaced by post-quantum cryptographic (PQC) methods in the near future [10].

Therefore, the security of the Schnorr signature scheme against cryptanalytic attacks implemented by classical computers is based on the difficulty of the discrete logarithm problem (DLP) and on the H-function modelled as a random oracle [8]. In real life, random oracles do not exist, and the cryptographic primitives constructed in this model cannot be used directly. Therefore, in practice, the requirement to model H-function as a random oracle is replaced with collision-resistant H-functions.

The security of deanonymization relies on the security of the Schnorr multi-signature scheme and Schnorr-based NIZKP of knowledge of Alice $PrK = x$ satisfying (19), (26), (27).

However, in a multi-user environment, the rogue-key attack can be arranged [12]. In this case, the possible scenario is the following. Anyone on the Net can compute a rogue-key and create a Schnorr multi-signature without using the signatures of other signers. But this attack is prevented according to the results presented in [12]. Referring to [12], the rogue-key attack is naturally prevented in our case, since in the verification equation (34), the public keys a_1 and a_2 are raised by different exponents h_1 and h_2, respectively. Moreover, in the alternative scenario, it is senseless for Alice to act against herself.

The security of Schnorr-based NIZKP based on (32) also relies on the discrete logarithm (DL) assumption, the Fiat-Shamir heuristics, and the random oracle model. This means that an efficient adversary who can impersonate a legitimate user must possess the corresponding secret key and is therefore able to efficiently solve the discrete logarithm problem (DLP) in large groups. As mentioned above, the group $\boldsymbol{Z_p^*}$ and the subgroup $\boldsymbol{G_q}$ are large. Moreover, the same adversary can also efficiently recover x in (19) if they gain access to all but one of the secret values x_i. However, the semiring $\boldsymbol{Z_q}$ is an additive group with the addition operation mod q. Therefore, it preserves the uniform distribution of the sum, given that all the summands are chosen independently and uniformly at

random. This means that the probability of guessing the PrK in (19) is independent of the number of created anonymous accounts N, used for deanonymization and is negligible under the DL assumption. Obviously, the same holds for (32) as well.

Also, the identity of a suspicious user can be restored using the special soundness property of the Schnorr sigma identification protocol in [11]: given two accepting conversations for the statement a_i the IC can compute the value of the witness x_i corresponding to it. Therefore, the IC can identify a dishonest user [11]. This property is preserved in an NIKZP adaptation.

Moreover, due to the soundness property of the Schnorr e-signature, the verifier accepts a false conversation with a negligible probability. Any adversary that can make the verifier accept a false conversation with non-negligible probability can also be used to solve DLP.

Finally, the direct attack is to recover the private key x in $\boldsymbol{Z_q}$ and to perform NIZKP in (32). However, due to the DL assumption, the adversary's advantage $AdvDL(g, a, p)$ is negligible if q is greater than 2^{1024}.

6. Conclusions

The proposed anonymization-deanonymization system is based on a unified cryptosystem, usually named the ElGamal cryptosystem, and therefore uses the general public parameters.

The alternative method is based on ECC ring signatures, requiring more ECC "exponentiation" operations to sign transactions since Alice must sign transactions for all ring members. In this paper, we are using a unified approach based on the integration of Schnorr signature, Schnorr multi-signature and Schnorr-based Non-Interactive Zero Knowledge Proof (NIZKP) to create anonymization, de-anonymization system in a private blockchain under the recognized security assumptions.

The proposed solution provides much more anonymity as compared with the anonymity provided in Ring Signatures (RS) based on the Elliptic Curve Cryptography (ECC) method used in Monero blockchain.

We estimate the effectiveness of realization between RS and our approach by comparing the number N_{ECO} in (14) of EC point multiplication operations to the number of exponentiation operations used in our proposal, based on Schnorr methods, which we denote by N_{EO}.

In our proposal, the effectiveness of anonymization is significantly higher than in the case of RS. In RS schemes, the signer must account for the fact that the number of operations N_{ECO} grows cubically with the number of users n in the ring, as shown in (14).

In contrast, in our construction, the anonymization is linearly dependent on the number of created anonymity addresses. For a single anonymous address, only two exponentiation operations are required. Moreover, anonymization can be performed offline at any time.

For Schnorr multi-signature, no exponentiation operations are required.

For Schnorr-based NIZKP, only two exponentiation operations are required.

The proposed system achieves significantly higher anonymity than RS-based schemes. In the latter case, when the ring consists of n members, the probability of guessing the transaction creator is $\frac{1}{n}$. In our proposal, this probability is equal to $\frac{1}{NoU}$, where NoU stands for the number of users.

This paper is an extension of a B2C' 2025 conference report and proceedings [13].

References

[1]. D. J. Bernstein, T. Lange, Faster addition and doubling on elliptic curves, in Advances in Cryptology - ASIACRYPT 2007, Lecture Notes in Computer Science, Vol. 4833, *Springer*, 2007, pp. 29–50.

[2]. D. Hankerson, A. Menezes, S. Vanstone, Guide to Elliptic Curve Cryptography, *Springer*, 2004.

[3]. G. Maxwell, A. Poelstra, Borromean ring signatures, Technical Report, *Blockstream*, 2015. Available at: https://github.com/Blockstream/borromean_paper/raw/master/borromean_draft_0.01_34241bb.pdf (Accessed: 20 March 2026).

[4]. M. M. T. Chakravarty, J. Chapman, K. MacKenzie, O. Melkonian, et al., The extended UTXO model, in Financial Cryptography and Data Security: FC 2020 International Workshops, Lecture Notes in Computer Science, Vol. 12063, *Springer*, 2020, pp. 525–539.

[5]. R. L. Rivest, A. Shamir, Y. Tauman, How to leak a secret, in Advances in Cryptology — ASIACRYPT 2001, Lecture Notes in Computer Science, Vol. 2248, *Springer*, 2001, pp. 552–565.

[6]. C. P. Schnorr, Efficient signature generation by smart cards, *Journal of Cryptology*, Vol. 4, Issue 3, 1991, pp. 161–174.

[7]. S. Micali, K. Ohta, L. Reyzin, Accountable-subgroup multisignatures, in *Proceedings of the 8th ACM Conference on Computer and Communications Security*, *ACM*, 2001, pp. 245–254.

[8]. J. Katz, Y. Lindell, Introduction to Modern Cryptography, 2nd ed., *Chapman and Hall/CRC*, 2015.

[9]. P. W. Shor, Algorithms for quantum computation: Discrete logarithms and factoring, in *Proceedings of the 35th Annual Symposium on Foundations of Computer Science*, *IEEE*, 1994, pp. 124–134.

[10]. Penta Security Inc., Post-quantum cryptography: A new security paradigm for the post-quantum era, *Penta Security Blog*, 5 June 2025. Available at: https://www.pentasecurity.com/blog/security-issue-post-quantum-cryptography/ (Accessed: 20 March 2026).

[11]. D. Boneh, V. Shoup, A Graduate Course in Applied Cryptography, Draft (Version 0.6), 2023. Available at: https://toc.cryptobook.us/book.pdf (Accessed: 20 March 2026).

[12]. M. Bellare, G. Neven, Multi-signatures in the plain public-key model and a general forking lemma, in *Proceedings of the 13th ACM Conference on Computer and Communications Security*, *ACM*, 2006, pp. 390–399.

[13]. E. Sakalauskas, A. Kilčiauskas, A. Bendoraitis, A. Michalkovič, et al., Anonymization-deanonymization System in Private Blockchain, in *Proceedings of the 4th Blockchain and Cryptocurrency Conference (B2C' 2025)*, Innsbruck, Austria, 25–27 November 2025, pp. 6-10.

ISSN 2938-2602