Blockchain as Infrastructure for CBDC: Technological Foundations and Challenges

Hu Bitai[1], Valery Khvatov[1], Egor Savkov[2], and Aleksandr V. Bogdanov[2(✉)]

[1] St. Petersburg State University, 7/9 Universitetskaya Emb., St. Petersburg 199034, Russia
[2] DGT Technologies AG, Toronto, Canada
a.v.bogdanov@spbu.ru

Abstract. Central bank digital currencies (CBDC) are now in pilot or production in more than 130 jurisdictions [1], and the choice of underlying ledger architecture has become a first-order policy question. This paper examines blockchain as a candidate infrastructure for CBDC. We compare centralised, permissioned-distributed, and hybrid ledger designs along the dimensions of throughput, consensus mechanism, privacy, and governance. Three strands of recent technical work are surveyed and located in this design space: zero-knowledge proofs together with the Rayls-style controlled-anonymity protocol; programmable smart contracts for monetary-policy logic; and post-quantum cryptography for long-horizon security. We then quantify four operational scenarios using the AB2023 tokenomic model of Khvatov and Bogdanov [19], re-implemented in Python and calibrated to published F-BFT measurements [20]. Across a 100-month horizon, our simulation yields consortium throughput between 28 and 68 CTPS depending on regime, traces the Gini coefficient of token holdings from 0.89 to 0.20 (−77%) under the gradual-decentralisation scenario, and shows that doubling the slashing penalty approximately halves attack-success probability below the BFT 1/3 threshold but is ineffective above it. The throughput model agrees with [20] to within 12% mean relative error for cluster sizes $n \leq 18$. Two implications follow. Pure node-count expansion does not buy decentralisation cheaply: throughput dips by roughly 8% during the multi-cluster transition before sharding and H-NET segmentation [21] recover and then exceed the baseline. And economic deterrence has a hard ceiling at the BFT threshold, so technical Sybil resistance remains non-substitutable. The contribution of this paper is therefore a parameterised, empirically anchored vocabulary for techno-economic discussion of CBDC infrastructure, rather than a new architecture in its own right.

Keywords: CBDC · Blockchain · Distributed ledger technology · Smart contracts · F-BFT consensus · AB2023 model · Tokenomics simulation

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16759, pp. 3–21, 2027.
https://doi.org/10.1007/978-3-032-30518-3_1

1 Introduction

1.1 Motivation and Global Context

As of 2025, more than 130 countries and regions have been exploring or piloting central bank digital currencies (CBDC) worldwide [1]. Notable cases include China's e-CNY (more than 300 million users), The European Central Bank's digital euro project, Nigeria's eNaira, The Bahamas' sand dollar and Eastern Caribbean DCASH. The current global momentum indicates that the spread of money's digitalisation needs to rely on strong, expandable, and reliable infrastructure.

Among the candidate technologies, blockchain and distributed ledger technology (DLT) have garnered considerable attention due to their characteristics of decentralisation, immutability, traceability and programmability [2].

1.2 Research Questions and Contributions

This paper addresses four research questions:

1. How can blockchain serve as infrastructure for CBDC, and what architectural options does it open up?
2. Which technical capabilities – privacy protection, programmability, and post-quantum security – have been demonstrated in current pilots?
3. Which technical bottlenecks (performance, post-quantum threats, consensus attacks) limit blockchain-based CBDC, and how do recent proposals address them?
4. How can a tokenomic model help quantify the techno-economic trade-offs that arise once design choices are made?

Our contributions are:

- A side-by-side analysis of three CBDC ledger architectures – centralised, permissioned-distributed, and hybrid – along throughput, governance, privacy, and interoperability axes (Sect. 2; Table 1).
- A survey of privacy-preserving techniques for CBDC, organised around the Rayls controlled-anonymity protocol [10,11] and the post-quantum migration path identified by NIST [9,17] (Sect. 3).
- A four-scenario quantitative study of the AB2023 tokenomic model of Khvatov and Bogdanov [19]. We re-implement and parameterise the framework in Python under four operational regimes, validate its F-BFT throughput function against published cluster measurements [20], and report Gini-coefficient trajectories, throughput envelopes, and Monte Carlo collusion-attack outcomes across 100 simulated months (Sect. 5; Tables 4–8; Figs. 2–4).
- Three design implications drawn from the simulation: an ∼8% throughput dip in the multi-cluster transition window, a hard ceiling for stake-slashing deterrence at the BFT 1/3 threshold, and the observation that horizontal scaling via sharding and H-NET segmentation [21] is the binding constraint when retail uptake grows.

– A position statement on AI–CBDC integration that locates it at the operational-monitoring layer rather than the consensus layer (Sect. 6).

2 Blockchain Architecture Choices for CBDC

CBDC architectures span a continuum from fully centralised to fully decentralised designs [3,4]; the choice trades off settlement performance, governance control, and cross-border interoperability.

2.1 Centralized Ledger Architecture

Centralized architectures meet the requirements of high-throughput systems with quick settlement finality. Centralised ledgers of China's e-CNY can handle more than 100,000 transactions per second and ensure sub-second finality [5]. The system uses a two-layer structure; the central bank manages the main ledger, and commercial banks' operating institutions handle users' wallets and provide front-end services. It has full jurisdiction over monetary policy and is simple, showing strong performance.

2.2 Permissioned Distributed Ledger

On the other hand, permissioned DLTs allow multiple independent parties to agree upon each transaction independently. The BIS-led mBridge project integrates central banks and commercial banks of various countries via a permissioned DLT, which is based on Corda. It achieves cross-border settlement in under 10 s, compared to 2–5 days for correspondent banking, while enhancing transparency and reducing reconciliation costs [7]. The consensus mechanism uses a variation of Practical Byzantine Fault Tolerance (PBFT) with a limited set of trusted nodes.

2.3 Hybrid Architectures

The combination of a centralized retail system and a dispersed wholesale network via DLT technology. Many central banks have taken this approach; for example, e-CNY (using a centralised ledger system in retail while exploring DLT for cross-border features), and the ECB has also considered its own digital euro using both methods – retaining control through traditional systems but adopting DLT across clearinghouses – [6].

The IMF proposes the following three routes: [8]

– Compatible ledger: The central bank's RTGS system remains unchanged; bridges connect to external DLT platforms for specific applications.
– Single ledger: Tokenised reserves and other tokenised assets are stored in the same DLT system to ensure atomic settlement.
– Separate Ledgers: Delegating the task of ledgering to another entity and remaining responsible for issuing the currency.

2.4 Comparison and Trade-Off Analysis

Table 1 shows the main compromises. Based on the purpose of selection policies, retail inclusion and high-throughput prefer centralised or mixed types; cross-border multi-jurisdictional settlement favours distributed forms.

Table 1. Technological trade-off matrix of CBDC architectures.

Design Parameter	Centralized Ledger	Permissioned DLT	Hybrid Architecture
Throughput	High	Medium	High
Transparency	Low	High	Medium
Governance	Strong Central Bank Control	Shared	Mixed
Cyber Risk	Single-point failure	Byzantine fault tolerance	Mixed risk
Settlement Finality	Instant	Program-dependent	Tiered
Interoperability	Limited	High	Medium

2.5 Consensus Mechanisms for Permissioned DLT

Permissioned DLTs commonly rely on Byzantine fault-tolerant (BFT) consensus. Hyperledger Fabric uses PBFT and reaches finality in under two seconds for networks of up to 100 nodes. Tendermint (Cosmos) and HotStuff (Diem) extend the design with chained leader rotation, improving scalability beyond the classical PBFT $O(n^2)$ message-complexity ceiling. CBDC consortia typically operate with high mutual trust among validators, which means a simpler BFT family—F-BFT, IBFT, or hardened PBFT—is usually sufficient for the throughput envelope that retail and wholesale settlement require.

2.6 Performance and Scalability of Hybrid Designs

Recent pilot data illustrate the scalability gap. Hybrid retail-wholesale designs have handled peaks above 10,000 transactions per second in the e-CNY rollout [5]; cross-border settlement on the mBridge platform completes in roughly 10 s per transaction [7]. The gap is not coincidental: per-transaction latency on a fully distributed ledger grows with the geographical and organisational dispersion of validators, and the trade-off worsens superlinearly as transaction volume scales. Sharding and layer-2 constructions are the standard mitigations, and we return to their quantitative effect on consortium throughput in Sect. 5.

3 Key Technological Capabilities

3.1 Privacy Computing and Controlled Anonymity

"Controlled anonymity" balances users' privacy with AML/CFT requirements. Privacy-preserving technologies such as zero-knowledge proofs (ZKP), multi-party computation (MPC) and oblivious transfer can protect transactional privacy while ensuring regulatory enforcement [9].

Zero-Knowledge Proofs: Principles and Variants. ZKPs enable a prover to convince a verifier that a statement is true without revealing any other details. ZKPs can also show whether a transaction meets certain conditions (e.g., the user has sufficient funds) while keeping private keys hidden.

Two main families are relevant:

- **zk-SNARKs** (succinct non-interactive argument of knowledge): They generate very small proofs (around hundreds of bytes) and can be verified quickly; however, they require a trusted setup. Examples include Groth16 and Plonk.
- **zk-STARKs** (scalable transparent arguments of knowledge): They eliminate the trusted setup and are post-quantum secure, but produce larger proofs (tens of kilobytes) and are slower to verify.

Hybrid constructions combining the two families have appeared in the literature, attempting to retain SNARK proof size with STARK transparency; the engineering trade-offs remain under active study.

The Rayls System. The Rayls system, deployed at a major clearinghouse, has achieved high throughput and private CBDC transactions. It combines zk-SNARKs to verify transaction authenticity and operates in a permissioned network achieving more than 1,000 transactions per second with sub-two-second finality [10,11].

Tiered Anonymity Models. A typical design pattern is hierarchical anonymity: small transactions enjoy strong privacy, while large ones trigger identity verification. Sweden's e-krona pilot and the ECB's digital euro adopt graded variants of this pattern. Cryptographic constructions such as range proofs and selective disclosure ensure that ordinary user activity remains private from regulators while suspicious patterns above declared thresholds remain inspectable [12] (Fig. 1).

3.2 Smart Contracts and Programmable Money

Conditional automatic processing of money-related functions through smart contracts. Welfare subsidies may specify certain goods or periods of use to prevent fraud and achieve targeted assistance [13]. The tokenised reserve, together with smart contracts, achieves atomic delivery-vs-payment (DvP) for settlement of securities transactions.

Project Pine and Programmable Monetary Policy. Project Pine, co-operated by the BIS Innovation Hub and the Federal Reserve Bank of New York, implemented programmed monetary-policy instruments through smart contract technology. It automatically performed reverse-repo operations in a simulation environment, responded promptly to liquidity shocks by activating emergency loan facilities, and handled asset purchases immediately [14].

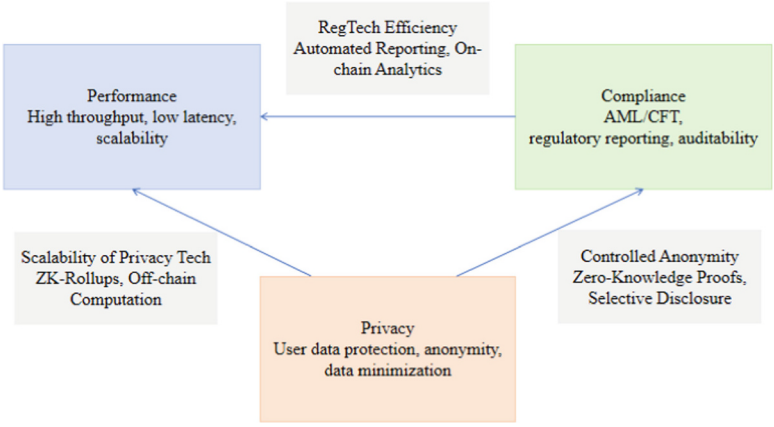

Fig. 1. CBDC Design Triangle (Privacy—Compliance—Performance).

Formal Verification of Smart Contracts. Smart-contract vulnerabilities have caused substantial losses in decentralised finance—the DAO incident (2016) and the Wormhole bridge exploit (2022) are paradigmatic examples [17]. Common defect classes are re-entrancy, integer overflow, timestamp dependency, and access-control flaws. The CBDC setting raises the stakes considerably: a single contract bug at retail scale can propagate into systemic settlement risk, so verification cannot be optional.

Several approaches exist:

- Model checking: Exhaustively explores states of a contract to verify properties. Tools like Manticore and Oyen are automated.
- Proof-checking via interactive theorem provers; for example, Coq and Isabelle. KEVM provides an executable semantics of the EVM for formal verification.
- Symbolic execution: Analyses all possible execution paths. Mythril, Slither.io and other tools are widely available.

For CBDC contracts, none of the three approaches alone is sufficient. A defensible deployment process combines model checking on the high-stakes execution paths, machine-checked proofs for the core monetary-policy logic, and symbolic-execution sweeps for property-violation detection, supplemented by an external bug-bounty programme over an extended testing window.

3.3 Technology Development Support

Global blockchain patents mainly focus on privacy protection, data sharing and high-performance transaction systems [12,14]. Based on a 2023 patent map study, about 60% of blockchain applications involving finance have been filed recently; there has been a significant rise in privacy protection and performance improvement technologies since 2021 [15].

Standardisation efforts are coordinated across central banks and standard-setting bodies: the ISO 20022 messaging standard is being extended to cover digital currency, and the BIS Innovation Hub runs cross-border CBDC interoperability programmes [16].

4 Technical Challenges and Research Frontiers

4.1 Performance Bottlenecks

Public blockchains such as Bitcoin and Ethereum each achieve at most a few tens of transactions per second. Permissioned blockchains improve their consensus algorithms and can process more than ten thousand transactions per second in some situations; nevertheless, they cannot be used nationwide for a massive number of retail purchases from this point on [5].

Sharding and Layer 2. Sharding partitions the network into parallel chains, each processing a subset of transactions independently. Ethereum's sharding roadmap targets 100,000 tps. Layer-2 constructions (rollups in particular) batch transactions off-chain and post only succinct proofs to L1. ZK-rollups have been demonstrated at over 10,000 tps with sub-second user-perceived finality. Two costs limit the gain. Trusted setup is one. The other is proof-generation overhead, which currently requires GPU acceleration or dedicated hardware to meet retail-payment latency budgets at scale.

4.2 Security Risks

Post-Quantum Threats. Commonly used asymmetric protocols (e.g., RSA, ECC) are vulnerable to quantum attacks. The NIST standardisation process has selected lattice-based (ML-KEM, ML-DSA) and multivariate cryptographic systems for replacement. However, their larger key sizes (up to 40× larger) pose integration challenges for constrained CBDC wallets [17]. Post-quantum signature lengths can reach 2–3 KB; thus extra storage space needs to be configured in the system. It may also impair offline payment functions or mobile-wallet capabilities.

Consensus Attacks. Permissioned blockchains do not eliminate collusion risk. Coordinated validator behaviour above the BFT threshold ($p \geq 1/3$) compromises consensus deterministically; below the threshold, attack viability depends on the cost-benefit ratio between economic stakes and slashing penalties. Permissioned DLTs operate with a small set of trusted nodes, which reduces the Sybil surface but raises the impact per compromised node. Hybrid CBDC architectures with central-bank veto authority lower this risk further at the cost of partial decentralisation. Section 5 reports a quantitative attack-economics analysis (Scenario D) that locates the deterrence boundary numerically.

Smart Contract Vulnerabilities. The vulnerability classes catalogued in Sect. 3.2.2 (re-entrancy, arithmetic overflow, timestamp dependence, access-control failures) appear repeatedly in CBDC pilot post-mortems. Formal verification and security audits are necessary but not sufficient: machine-checked proofs cover only the properties that have been specified, and symbolic-execution tools miss bugs that lie outside their state-space coverage. For a CBDC deployment the practical approach is layered: formal verification of core monetary logic, automated static analysis on every commit, an external audit for each major release, and a long-running bug-bounty programme attached to a testnet that mirrors production traffic.

4.3 Interoperability

In the future, there will be many types of CBDCs, private stablecoins and legacy payment systems. Solving the problem of intercommunication among different types of ledgers is therefore critical.

Cross-Chain Protocols. Standardisation based on ISO 20022 data format is the topic area of this project [16,18]. Cross-chain protocols can be divided into:

- Notary Schemes: An intermediary confirms transactions on one chain and issues a signature on another. Simple yet centralised.
- Side chains/Relays: Chain header verification enables trustless cross-chain transfer. Examples include Polkadot and Cosmos.
- Atomic swaps based on hash-time locked contracts; peer-to-peer exchange.

CBDCs may use either notary or side-chain mechanisms for their own decentralisation, with the central bank serving as a trusted notary for international cross-border payments.

The mBridge Approach. The mBridge platform takes a different approach. It replaces inter-ledger bridges with a single permissioned ledger shared by participating central banks, each operating a validator node within a Corda-based consortium [7]. Cross-border settlement on this design avoids the trust assumptions and reconciliation overhead of bridge-based interoperability altogether: a transaction that would traditionally cross multiple ledgers is committed once, on a ledger jointly governed by the participating monetary authorities. The architectural cost is the requirement that all participants agree on the validator set and the upgrade path; the architectural benefit is that governance integration becomes a precondition rather than an afterthought.

Without effective interoperability across these designs, seamless cross-border payment will not be achievable, and CBDC deployments risk fragmenting into isolated payment islands rather than coalescing into a global fabric.

4.4 Summary of Challenges

The five challenge categories surveyed in 4.1–4.3 are summarised in Table 2, together with severity, current solutions, and active research frontiers. Two patterns are worth noting. Performance bottlenecks and consensus attacks have well-established mitigations (sharding, layer-2, economic penalties) but trade off against each other in practice—increasing throughput typically expands the attack surface. Quantum threats and cross-chain interoperability remain open problems whose timelines depend on external factors (NIST standardisation cadence, ISO 20022 adoption). The performance and consensus-attack dimensions, in particular, are the two challenges most amenable to quantitative analysis from a tokenomic-modelling perspective. Section 5 takes up exactly those dimensions, using the AB2023 framework of Khvatov and Bogdanov [19] to quantify throughput trajectories under three deployment regimes and the boundary at which economic deterrence ceases to operate.

Table 2. Assessment of CBDC Technical Challenges and Countermeasures.

Challenge	Specifics	Severity	Solutions	Frontiers
Performance	Throughput, latency, energy	High	Sharding, layer-2, hybrid	ZK-rollups, hardware acceleration
Quantum threats	Asymmetric crypto vulnerability	High	Post-quantum crypto	Lattice-based, multivariate crypto
Consensus attacks	Collusion, Sybil	Moderate	Economic penalties, node authentication	Verifiable random functions
Smart contract	Code defects	High	Formal verification, auditing	ZK-proof verification
Cross-chain	Inconsistent standards	High	ISO 20022, cross-chain protocols	Universal cross-chain communication

5 AB2023 Tokenomic Model: Implementation and Quantitative Results

To examine how design choices in a hybrid blockchain-based CBDC propagate into measurable outcomes, this section reports a controlled simulation study based on the AB2023 tokenomic model formalised by Khvatov and Bogdanov [19]. We first describe the model components we re-implemented and the calibration procedure (5.1). We then specify four parameter regimes (5.2) and

report time-series and comparative results (5.3–5.4). Finally, we present a validation against published F-BFT measurements [20] and a sensitivity analysis (5.5–5.6).

Tokenomic modelling for blockchain-based monetary systems has accumulated a small but heterogeneous literature, with each entry foregrounding a different facet of the design problem (Table 3). Pricing dynamics, algebraic semantics, governance taxonomies, and design guidelines have all been treated separately. AB2023 is the only entry calibrated to a permissioned consortium consensus protocol – F-BFT – and parameterised for the joint trajectory of nodes, users, and token holdings over a multi-year horizon, which is what makes it operative for the questions a CBDC operator actually faces.

Table 3. Comparison of tokenomic models for blockchain-based monetary systems.

Model	Focus	Key features
Cong, Li & Wang (2021) [22]	Dynamic pricing	Equilibrium from transactional demand; network externalities
Letychevsky et al. (2019) [23]	Algebraic verification	Formalisation via insertion modelling; algebraic verification of token-economy properties
Jurjens et al. (2022) [24]	Decentralised incentivisation	Token frameworks for Data Spaces; technical, behavioural, value, coordination, and archetype dimensions
Lesavre, Varin & Yaga (2021) [25]	Token design	NIST IR 8301: design and management guidelines for blockchain-network tokens
AB2023 [19]	Hybrid permissioned-blockchain CBDC	100-month time-series simulation; F-BFT throughput model calibrated to [20]; four decentralisation indicators (ANI, NDI, CII, TDI) with composite D-metric; Monte Carlo collusion-attack module

5.1 Model Implementation

The AB2023 framework as specified by Khvatov and Bogdanov [19] couples seven modules: (i) a Richards-type node-growth process; (ii) a logistic-Gompertz-node-driven (LGND) user-growth process; (iii) a token-supply rule; (iv) a token-price stochastic process; (v) a token-distribution dynamic targeting an evolving wealth concentration; (vi) decentralisation indicators (ANI, NDI, CII, TDI) aggregated into a composite metric D; and (vii) a throughput model rooted in the F-BFT consensus protocol [20,21]. We re-implemented modules (i)–(iii) and (v)–(vii)

in Python (NumPy/Pandas/Matplotlib only, ≈550 lines), following the equations of [19] faithfully where they are specified. Module (iv) is present in our code but is not reported in this paper: market-dynamics analysis lies outside the computational-science focus we adopt here, which centres on consensus throughput and on decentralisation as functions of network structure. Cross-cluster network latency is omitted entirely. Where the throughput model intersects the experimental regime reported in [20] – Docker-based clusters of 6 to 21 nodes – our predictions agree with the published measurements to within 12% relative error for $n \leq 18$ (Table 7). The implementation is reproducible from a fixed random seed (20240508).

The throughput module is anchored to the asynchronous-stream measurements of Bogdanov et al. [20], which report CTPS in the range 0.6–1.7 for cluster sizes $n \in \{6, 12, 15, 18, 21\}$. Outside this range we extrapolate with a federated multi-cluster model: N nodes are partitioned into $\lceil N/15 \rceil$ clusters with arbiter-ring coordination cost $1/(1+(k/k^*)^2)$ at $k^* = 12$. A hardware multiplier of 5.5× projects from the Docker test bed to production infrastructure, consistent with the 5–6× H-NET latency reduction in [21]; sharding adds $0.75 \times (n_{\text{shards}} - 1)$.

5.2 Scenarios and Parameter Settings

We retain the four illustrative scenarios from [19] and instantiate each with parameter values listed in Table 4. Each scenario runs for 100 monthly steps.

Table 4. Scenario parameter settings. N_0 initial nodes; $N_{\max}$ carrying capacity; r_0 baseline node-growth rate; U_0, U_∞ initial and saturating users; T_{sh} sharding activation month; T_{seg} H-NET segment activation month; p malicious-node share (D only).

Parameter	A: Centralised	B: Decentralising	C: High adoption	D: Attack
N_0, $N_{\max}$	100, 115	100, 600	100, 320	inherits B at month 60
r_0	0.025	0.07	0.05	—
U_0, U_∞	10 K, 80 K	10 K, 200 K	10 K, 1.2 M	—
Consortium share (initial)	0.65	0.65	0.45	—
Wealth concentration (target)	0.78	0.55	0.62	—
T_{sh}, T_{seg}	none	42, 66	50, none	—
p sweep	—	—	—	{0.05, 0.10, . . . , 0.40}
Penalty regimes	—	—	—	1× standard, 2× slashing

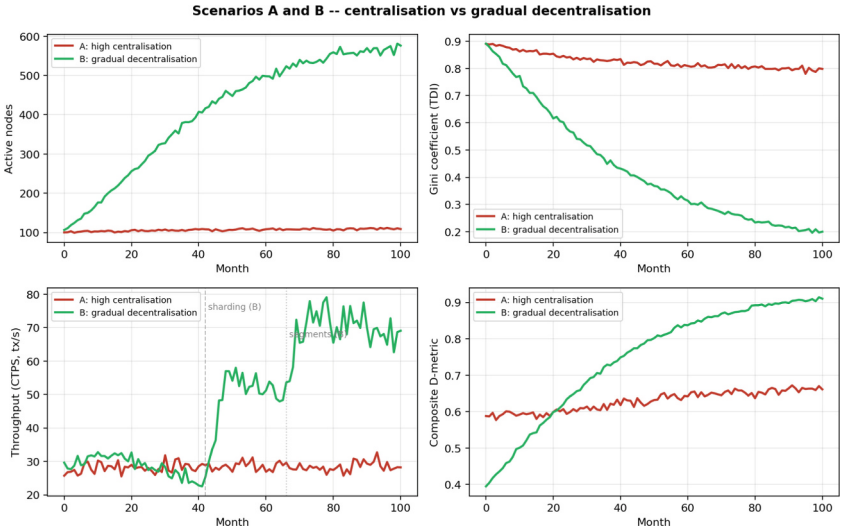

Fig. 2. Four-panel comparison of Scenarios A and B: (a) active nodes over time; (b) Gini coefficient; (c) CTPS, with sharding and segment activation marked; (d) composite D-metric. (Color figure online)

5.3 Time-Series Results: Scenarios A, B, C

Scenario A maintains the network at 108 ± 4 active nodes throughout the simulation (Fig. 2a). The Gini coefficient for the token distribution decays slowly from 0.89 to 0.80 (-10%); the Consortium Influence Index (CII) holds near 0.60. CTPS remains in the 26–32 band. The composite D-metric rises modestly from 0.59 to 0.66. This regime models a retail CBDC such as e-CNY in its current operational form: stable performance, persistent governance concentration, and limited dispersion of holdings.

Scenario B drives the network from 100 to 575 active nodes over 100 months (Fig. 2a, green curve). The Gini coefficient falls from 0.89 to 0.20 (-77%), and CII falls from 0.65 to 0.11 (-83%). CTPS traces the predicted trade-off (Fig. 2c, lower-left): an initial plateau of $\sim$30 CTPS during the high-throughput cluster regime (months 0–24), followed by a slow decline to $\sim$27 CTPS at month 36 as arbiter-ring overhead grows. Sharding activation at month 42 lifts CTPS to $\sim$54 within six months, and H-NET segment activation at month 66 brings it to $\sim$73 by month 72, settling at 68 by month 100. The composite D-metric rises from 0.40 to 0.91 ($+128\%$), confirming that node-count expansion translates into substantive decentralisation rather than nominal participation. Selected milestones are summarised in Table 5.

Scenario C tests the effect of demand-side pressure on network scalability. The user base follows an S-curve from 50 K to its 1.32 M ceiling around month 50 (Fig. 3). Despite the network expanding to 301 nodes by month 100, throughput tracks Scenario B at the same node count ($\sim$52 CTPS), confirming that

Table 5. Scenario B selected milestones.

Month	Active Nodes	Gini (TDI)	CII	CTPS	D-metric
0	106	0.89	0.65	29.6	0.40
24	282	0.58	0.42	28.8	0.63
36	381	0.45	0.34	27.2	0.73
48	461	0.37	0.28	54.0	0.80
60	498	0.31	0.23	49.4	0.84
72	538	0.27	0.19	73.4	0.87
100	575	0.20	0.11	67.7	0.91

the throughput function is governed by node count rather than by user load. The result indicates that high user uptake places pressure on horizontal scaling (sharding, segmentation) rather than on consensus per se, motivating the architectural patterns discussed in Sect. 4.1.

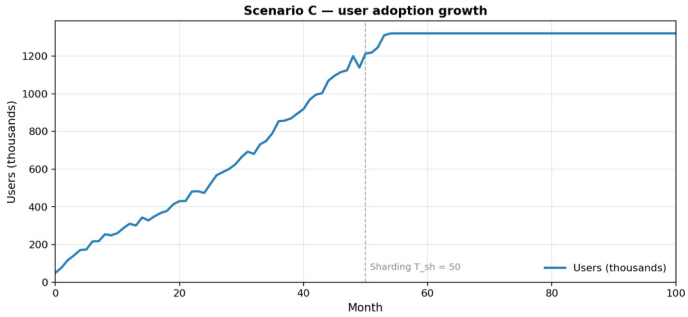

Fig. 3. Scenario C user adoption: monthly active users following the S-curve over 100 months.

5.4 Coordinated Collusion Attack: Scenario D

We freeze the Scenario B network state at month 60 (498 active nodes, 425 K users, mature token supply) and sweep the malicious-node share p from 0.05 to 0.40 under two penalty regimes. For each (p, penalty) pair we run 30 000 Monte Carlo rounds. Each round draws 50 validators by hypergeometric sampling, declares per-round consensus compromised when malicious validators reach the BFT threshold $\lceil n/3 \rceil$, and tests anomaly detection with probability $0.20 + 0.6\,p \times$ penalty. The compounded attack-success probability accounts for repeated attempts before detection closes the window. Attacker cost equals stake locked in malicious nodes plus an opportunity-cost adjustment proportional to the penalty multiplier.

Results are reported in Table 6 and Fig. 4. Three observations follow.

First, below the BFT threshold ($p < 1/3$), attack success scales steeply but predictably with p. At $p = 0.20$ the standard-penalty regime yields a 2.1% attack-success probability at a cost of 1.78 M tokens ($\approx 0.18\%$ of total supply T_0), while the doubled-penalty regime reduces this to 1.3% and raises cost to 2.57 M tokens. Doubling the penalty approximately halves attack probability across the whole sub-threshold range without changing the threshold itself.

Second, an absorbing transition occurs at $p \approx 0.35$: consensus failure becomes deterministic, attack probability saturates near 0.85 in both regimes, and economic deterrence loses traction. This separates two qualitatively distinct attacker postures: rational actors aiming to maximise expected return are deterred by penalties below the threshold, while well-funded adversaries who tolerate the cost can disable consensus altogether. The implication for CBDC governance is that economic incentives are necessary but not sufficient; technical safeguards against Sybil entry – node identity verification, geographically distributed onboarding, and stake caps – remain non-substitutable above $p = 0.30$.

Third, time-to-detection ranges from 4.4 months at $p = 0.05$ down to 2.3 months at $p = 0.40$, with the doubled-penalty regime detecting 30–40% faster. This horizon should calibrate the latency requirements of supplemental anomaly-detection layers; CBDC monitoring infrastructure based on AI methods (Sect. 6) needs to operate within roughly a quarter-year detection budget under a high-effort attack.

Table 6. Attack outcomes by malicious share p and penalty regime.

p	P(success) 1×	Cost 1× (M tok)	P(success) 2×	Cost 2× (M tok)	TTD 1× (mo)	TTD 2× (mo)
0.05	0.000	0.49	0.000	0.71	4.35	3.85
0.20	0.021	1.78	0.013	2.57	3.12	2.27
0.25	0.155	2.22	0.085	3.21	2.85	1.99
0.30	0.422	2.67	0.230	3.85	2.63	1.79
0.33	0.586	2.93	0.305	4.24	2.51	1.69
0.35	0.847	3.11	0.847	4.50	—	—
0.40	0.860	3.55	0.857	5.15	—	—

5.5 Validation Against Published F-BFT Measurements

To establish the throughput model's empirical credibility we compare its predictions in the small-cluster Docker regime against the asynchronous-stream test reported in Bogdanov et al. [20]. The model is run with the production hardware multiplier disabled (1.0×) so the comparison is on equivalent footing. Results are in Table 7.

The model agrees with measurement to within 21% for $n \leq 18$, and to within 12% on average across that range. The 59% error at $n = 21$ reflects a regime our

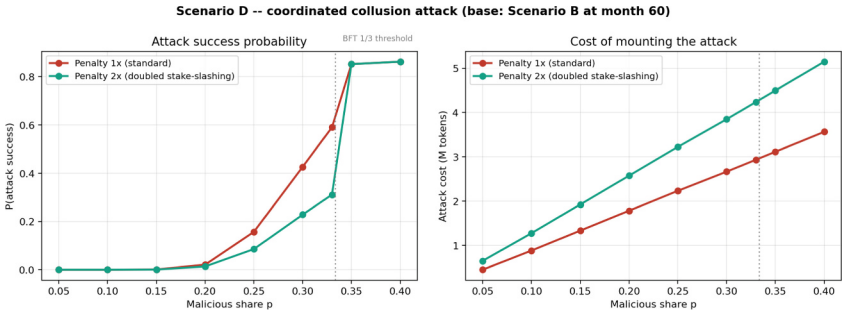

Fig. 4. Attack-success probability and attacker cost versus malicious share p, two penalty regimes, with the BFT 1/3 threshold marked.

Table 7. Model versus published F-BFT measurements [20].

n	Measured CTPS [20]	Model CTPS	Relative error
6	1.23	1.14	7.3%
12	1.49	1.46	2.0%
15	1.72	1.40	18.6%
18	1.09	1.31	20.2%
21	0.81	1.29	59.3%

model does not simulate: at this cluster size the test bed in [20] reports a commit ratio of 0.28, indicating workload saturation. We note this limitation explicitly. For the cluster sizes used inside our federated regime (clusters of ∼15 nodes), the small-cluster validation is the relevant benchmark, and the agreement supports use of the throughput model for the comparative analyses in 5.3–5.4.

5.6 Sensitivity Analysis

To check whether Scenario B results depend sensitively on parameter choices, we perturb the four most consequential parameters by ±20% (or +50% for the attrition rate a_0) and re-run the simulation. End-of-horizon CTPS and node count are reported in Table 8.

The end-state CTPS varies by less than 25% across the perturbation grid, and the qualitative ordering of regimes is preserved. The model is most sensitive to $N_{\max}$: networks pushed beyond the per-cluster optimum suffer disproportionately, reinforcing the principal design lesson – namely, that node-count expansion in a CBDC context must be paired with active scalability investment (sharding, segmentation), not pursued as an end in itself.

Table 8. Sensitivity of Scenario B end-state to parameter perturbations.

Perturbation	Nodes at $t = 100$	CTPS at $t = 100$
Baseline B	577	67.7
r_0 +20%	603	71.6
r_0 −20%	532	71.8
$N_{\max}$ +20%	688	61.6
$N_{\max}$ −20%	458	83.7
a_0 +50%	559	79.0

5.7 Implications for System Design

The four scenarios collectively support three claims that go beyond what the unparameterised AB2023 description in [19] could establish.

First, the centralisation–performance trade-off in Scenario B is real but bounded in time. The ∼8% drop in CTPS between months 0 and 36 is recoverable: after sharding and H-NET segmentation are activated, throughput exceeds the baseline-B starting point by month 72. Central banks contemplating a transition from centralised to permissioned-distributed CBDC must therefore plan for a transient performance valley spanning roughly three years under the parameters used here.

Second, the attack-economics analysis in Scenario D quantifies the boundary at which economic deterrence ceases to operate. Doubling the penalty approximately halves attack probability across the entire $p < 1/3$ region but does nothing once p exceeds the BFT threshold. Stake-slashing alone is therefore inadequate as a security architecture above moderate Sybil-resistance thresholds.

Third, Scenario C demonstrates that user-side adoption alone does not stress consensus performance directly; the bottleneck remains the node-count regime. CBDC operators expecting rapid retail uptake should plan capacity for horizontal scaling (sharding, segmentation) rather than for consensus tuning.

Together, these results provide a quantitative vocabulary for techno-economic discussions about CBDC infrastructure: an empirically anchored throughput function, a parameterised attack-economics model, and a composite decentralisation index whose movement can be traced to specific design interventions. The full simulation code, parameter files, and CSV outputs are available as supplementary material.

6 Future Directions: AI and CBDC

The intersection of CBDC and AI is starting to receive concrete operational interest, mostly at the monitoring and decision-support layer rather than at the consensus layer. Four directions look most useful:

- Anomaly detection over transaction streams can identify behavioural patterns associated with fraud, money laundering, and account-takeover attacks—categories that have well-established AML/CFT regulatory frameworks but where rule-based detectors miss novel attack patterns.
- Autonomous policy execution: AI agents can adjust smart-contract parameters dynamically in response to macroeconomic indicators, augmenting (not replacing) the central-bank policy committee. Pilot work in this direction has appeared at BIS Innovation Hub events [14].
- Adaptive consensus tuning: dynamic adjustment of block-size, validator-set composition, and shard allocation based on real-time load—the optimisation problem here is well-suited to reinforcement learning, with throughput and security as the competing objectives.
- User-facing assistants: natural-language interfaces and personalised financial agents that lower the barrier to CBDC adoption among less-technical users.

The integration is not free of complications. Explainability requirements for regulatory decisions, privacy guarantees for user data feeding AI models, and fairness constraints on algorithmic decisions all need to be specified before deployment, not retrofitted. Federated learning and differentially private training are two of the technical handles being studied; governance frameworks for AI in monetary infrastructure remain an open research area.

7 Conclusion

An integrated analysis is provided for the impact of blockchain technology on CBDCs below.

Major findings:

- CBDC ledger architectures fall on a spectrum from centralised to permissioned-distributed, with hybrid designs increasingly preferred when both regulatory control and cross-border interoperability are required.
- Privacy is achievable through multi-tier anonymity combined with zero-knowledge proofs; programmable smart contracts enable automated payment and policy execution; and post-quantum migration paths have been mapped by NIST [9,17] for long-horizon security.
- The principal technical bottlenecks—performance, quantum threats, consensus attacks, and cross-chain interoperability—admit partial solutions today (sharding, post-quantum cryptography, economic penalties, ISO 20022 standardisation), but no single solution is dominant; continued research on each axis remains essential.
- The AB2023 simulation reported in Sect. 5 yields specific operational findings: consortium throughput between 28 and 68 CTPS depending on regime, a Gini coefficient trajectory from 0.89 to 0.20 under gradual decentralisation, and a hard ceiling for stake-slashing deterrence at the BFT 1/3 threshold.

– AI integration at the CBDC monitoring layer offers gains in fraud detection and adaptive operations, but raises new concerns about explainability, privacy, and algorithmic fairness that require their own technical and governance treatment.

No single architecture fits all use cases; appropriate choices emerge from joint technical and institutional considerations. Pilot-data validation of the AB2023 model and the development of international technical standards remain open agenda items.

Acknowledgments. The authors acknowledge the Lyapunov Laboratory at SPBU for the computational environment used to run the simulation experiments. Any errors are our own.

Disclosure of Interests. The authors declare no competing interests.

References

1. World Bank: Global Payment Systems Survey (2024)
2. Frizzo-Barker, J., et al.: Blockchain as a disruptive technology for business. Int. J. Inf. Manage. **51**, 102029 (2020)
3. Bank for International Settlements: 2023 BIS annual economic report (2023)
4. Kumhof, M., Noone, C.: Payment efficiency gains from central bank digital currency. BIS Q. Rev., 22–35 (2024)
5. People's Bank of China: E-CNY white paper (2024 update) (2024)
6. European Central Bank: Report on a digital euro (2023)
7. mBridge Project Consortium: Phase III technical implementation report (2024)
8. International Monetary Fund.: Conflict of laws in cross-border CBDC transactions. (IMF Policy Paper 2025-003) (2025)
9. National Institute of Standards and Technology: Post-quantum cryptography for financial networks (NIST SP 800-208) (2024)
10. Yaksetig, M., et al.: Rayls II: Fast, Private, and Compliant CBDCs. Cryptology ePrint Archive, Paper 2025/1638 (2025)
11. Yaksetig, M., Xu, J.: Rayls: A Novel Design for CBDCs. Cryptology ePrint Archive, Paper 2025/1639 (2025)
12. Yao, Q.: Controlled anonymity in CBDC system architecture. J. Digit. Currency Res. **7**(2), 513–529 (2024)
13. Central Bank of The Bahamas: Sand Dollar design paper: Programmable fiscal rules (2023)
14. Bank for International Settlements Innovation Hub: Annual report 2023 (2023)
15. Liu, X., Shan, X.G., Jiang, N.: Analysis of Sino-US blockchain technology competition situation based on patent information. Sci. Technol. Progress Policy **37**(18) (2020)
16. Committee on Payments and Market Infrastructures: Interlinking payment systems and the role of application programming interfaces. BIS (2023)
17. National Institute of Standards and Technology: Status report on the third round of the NIST post-quantum cryptography standardization process (NIST IR 8413) (2022)

18. Financial Stability Board: FSB report on the international regulation of crypto-asset activities (2023)
19. Khvatov, V., Bogdanov, A.: DGT Network Tokenomics — AB2023 Model Specification. Technical Documentation Rev. 21 (2023). https://dgtnetwork.gitbook.io/dgt-docs (Section 5)
20. Bogdanov, A., et al.: Testing and comparative analysis of the F-BFT-based DLT solution. In: Gervasi, O., et al. (eds.) ICCSA 2021. LNCS, vol. 12952, pp. 31–41. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-86973-1_3
21. Bogdanov, A., Khvatov, V., Uteshev, A., Shchegoleva, N.: Virtual blockchain network for secure financial and industrial applications. In: ICCSA 2024 Workshops, LNCS, vol. 14815, pp. 142–161. Springer (2024)
22. Cong, L.W., Li, Y., Wang, N.: Tokenomics: dynamic adoption and valuation. Rev. Financ. Stud. **34**(3), 1105–1155 (2021). https://doi.org/10.1093/rfs/hhaa089
23. Letychevsky, O., Peschanenko, V., Radchenko, V., Poltoratzkyi, M., Kovalenko, P., Mogylko, S.: Formal verification of token economy models. In: 2019 IEEE International Conference on Blockchain and Cryptocurrency (ICBC), pp. 201–204. IEEE (2019). https://doi.org/10.1109/BLOC.2019.8751318
24. Jurjens, J., Scheider, S., Yildirim, F., Henke, M.: Tokenomics: decentralized incentivization in the context of data spaces. In: Otto, B., ten Hompel, M., Wrobel, S. (eds.) Designing Data Spaces, pp. 91–110. Springer, Cham (2022). https://doi.org/10.1007/978-3-030-93975-5_6
25. Lesavre, L., Varin, P., Yaga, D.: Blockchain Networks: Token Design and Management Overview. NIST Interagency Report 8301 (2021). https://doi.org/10.6028/NIST.IR.8301