From Supply Chain Finance to Paradigm Reconstruction: The Synergistic Evolution of CBDC and Blockchain

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. More than 130 jurisdictions now run CBDC pilots [1]; supply chain finance (SCF) is one of the densest test beds for joint deployment with distributed ledger technology [13]. We organise the joint-deployment problem along three axes: blockchain plus CBDC against the three SCF frictions (information asymmetry, settlement delay, SME credit constraint); three institutionally distinct national patterns (Singapore PBM [8], EU privacy-centred digital euro [6], U.S. market-driven framework [9]); and a four-feature paradigm—intelligence, trustworthiness, inclusion, cross-border—tested on a tokenomic simulation. We re-implement the AB2023 model of Khvatov and Bogdanov [14] in Python and run four scenarios over 100 months: e-CNY-style retail (115 nodes, ∼28 CTPS), wholesale permissioned-distributed (575 nodes, Gini 0.89 → 0.20, CTPS 30 → 27 → 68 with sharding at month 42 and H-NET segmentation [16] at month 66), high-adoption SCF (301 nodes, 1.32 M users, ∼52 CTPS), and a 30 000-round Monte Carlo collusion sweep. Doubling the slashing penalty halves attack-success below the BFT 1/3 threshold; above it both regimes saturate near $P = 0.85$. The throughput model agrees with F-BFT measurements [15] to within 12% mean error for $n \leq 18$. Two implications follow. The retail/wholesale CBDC choice is a quantifiable trade-off; the throughput valley between months 24 and 42 in Scenario B costs roughly three years if scalability investment lags behind node growth. Trustworthiness needs three layers: token-economic deterrence below BFT, technical Sybil resistance above it, and AI anomaly detection on a 2–4 month latency budget. The contribution is a parameterised vocabulary linking paradigm goals to operational outcomes, not a new architecture.

Keywords: CBDC · Blockchain · Supply chain finance · Cross-border payments · Programmable money · 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. 22–38, 2027.
https://doi.org/10.1007/978-3-032-30518-3_2

1 Introduction

1.1 Motivation and Global Context

CBDC have crossed from research literature to operational infrastructure in over 130 jurisdictions, and 94% of central banks report some form of CBDC engagement [1,4]. China's e-CNY has surpassed 300 million users [5]; the Eurosystem's digital-euro project sits in its preparation phase [6]; mBridge has reached Phase III multi-jurisdiction settlement [7]. The open question is no longer adoption but design: how the architectural and policy choices behind a CBDC interact with the application domains that absorb it first.

Supply chain finance (SCF) is the densest of those domains. The sector's three legacy frictions—multi-tier information asymmetry, five-to-ten-day letter-of-credit settlement, persistent SME credit constraint—map onto what blockchain and CBDC jointly address [12,13]: blockchain provides shared, immutable record-keeping and programmable execution; CBDC provides final settlement in central-bank money without intermediary chains. We take that joint deployment as our empirical anchor and use a tokenomic simulation to attach quantitative price tags to the policy choices it forces.

1.2 Literature Review

Four overlapping research streams cover the CBDC–blockchain–SCF intersection.

Stream 1—Blockchain in Supply Chain Finance. Frizzo-Barker et al. [2] surveyed 158 articles on blockchain as a disruptive technology and identified trade finance among the densest application clusters. Tang and Chen [19], on 88 listed Chinese logistics firms 2010–2023, report a 0.759 difference-in-differences coefficient of blockchain adoption on supply-chain resilience, mediated by innovation capacity and total factor productivity. Li and Du [18], surveying 1,760 papers, find digital currency, market efficiency, and SCF applications dominant. The stream treats the settlement asset as exogenous.

Stream 2—CBDC Design and Macro-effects. Kumhof and Noone [10] formalise CBDC design principles for financial stability; the BIS [3] provides annual cross-jurisdictional surveys. The stream is theoretically strong on monetary effects but rarely engages the technical substrate.

Stream 3—Cross-Border and Multi-CBDC. mBridge [7] has produced operational evidence that wholesale CBDC bridges compress settlement from days to minutes; the BIS Innovation Hub's Project Agorá [11] extends the architecture to a tokenised correspondent-banking layer. Empirical analysis remains thin because the implementations are recent.

Stream 4—Comparative National Approaches. MAS has published Project Orchid materials including the Purpose Bound Money protocol [8]; the ECB [6] documents digital-euro design with privacy as the governing constraint; the U.S. Federal Reserve [9] has issued discussion papers without committing to retail CBDC.

The streams rarely meet. SCF–blockchain studies fix the settlement asset; CBDC studies black-box the substrate; comparative studies stay qualitative. What is missing is a parameterised framework in which CBDC design choices, blockchain architecture choices, and SCF deployment patterns generate measurable joint outcomes. This paper supplies one.

1.3 Research Design and Contributions

We combine three methodological elements: comparative case analysis of national approaches, a paradigm-reconstruction synthesis built on the comparative material, and a parametric simulation of the AB2023 tokenomic model linking paradigm goals to operational outcomes. The research questions are: (i) which SCF pain points yield to joint blockchain–CBDC solutions and which resist them? (ii) what differentiates the Singapore, EU, and U.S. patterns in terms of deployment viability? (iii) at what threshold do paradigm goals become parameter-attainable rather than aspirational?

Contributions: a structured analysis of CBDC–blockchain synergy in SCF organised by three pain points and three solution mechanisms (Sect. 2; Table 1); a comparative framework over the three national patterns with explicit institutional and SCF implications (Sect. 3; Table 2); a four-dimensional paradigm derived from comparison rather than asserted (Sect. 4); a 100-month parametric simulation under four CBDC archetypes, validated against published F-BFT measurements [15] to within 12% mean error and reporting quantitative trade-offs that link paradigm goals to operational outcomes (Sect. 5; Tables 3–7; Figs. 1–2). Every paradigm claim in Sect. 4 maps to a measurable AB2023 indicator in Sect. 5; every Sect. 5 number reproduces from Table 3 with random seed 20240508.

The paper proceeds: Sect. 2 covers SCF pain points and joint solutions. Section 3 presents the comparative cases. Section 4 develops the four-feature paradigm. Section 5 reports the simulation. Section 6 closes with policy implications and an agenda.

2 Synergistic Solutions in Supply Chain Finance

SCF moves working capital and liquidity along multi-party trade networks [13]. Three structural inefficiencies persist that neither blockchain nor CBDC fully solves alone: information asymmetry across multi-tier supplier networks [18,19], in which financial institutions cannot verify at low cost that an invoice corresponds to a genuine trade, leading to elevated SME financing premia; operational delay, with letter-of-credit transactions taking five to ten business days from application to fund release [12], the bottleneck being manual reconciliation across banks, freight forwarders, customs, and treasuries rather than any single step; and SME-tier liquidity constraints [13,18], where bank programmes favour anchor enterprises over the long tail of suppliers, with the resulting credit shortfall being one of the most-cited frictions in the empirical SCF literature [19].

2.1 Blockchain Mechanisms for Trust and Transparency

Three mechanisms account for most reported benefits. *Immutable record-keeping* [2]: order placement, shipment events, inspection reports, and invoice issuance write to a shared ledger that authorised participants verify independently, which suppresses invoice double-pledging—one of the larger historical loss categories for trade-finance lenders. *Smart-contract automation* [12]: business rules such as payment on inspection or against bill of lading encode as contracts that trigger on condition match; IoT sensor evidence (RFID inventory, GPS shipping data) feeds them as oracles, eliminating most manual reconciliation. *Tokenisation of real-world assets* [18]: accounts receivable, warehouse receipts, and other trade assets become fungible or non-fungible tokens with fractional ownership, which expands the funding pool for suppliers who lacked a liquid receivables market. The empirical case is no longer purely conceptual: Tang and Chen [19] report a 0.759 DiD coefficient on supply-chain resilience over 88 firms (2010–2023), with innovation capacity and total factor productivity as significant mediators. The settlement-asset choice—in particular whether that asset is a CBDC—is the next question.

2.2 The Enabling Role of CBDC: Value-Flow Integration

CBDC, as a direct liability of the central bank, supports final settlement in central-bank money without intermediary chains [10]; integrated with a blockchain trade platform, payment occurs atomically with title transfer (DvP at machine speed), collapsing multi-day letter-of-credit settlement risk. CBDC issuance can carry conditional restrictions on use [14]: in an SCF context, a buyer releases payment spendable only by designated sub-suppliers, preventing leakage to unrelated obligations—the mechanism Singapore's PBM implements directly [8]. For cross-border settlement, the largest cost line in international trade finance, multi-CBDC platforms such as mBridge let participating commercial banks transact directly on a shared ledger, compressing settlement from days to seconds and reducing aggregate cost by roughly 40–60% [7].

3 Global Exploration Paths: A Comparative Analysis

National CBDC–blockchain integration diverges along institutional axes that the technology alone does not determine. Three cases bracket most of the variation observed: Singapore's design-driven, ecosystem-building approach; the EU's privacy-first, compliance-bound trajectory; and the U.S. market-driven, regulation-first stance. The differences condition which deployment patterns are SCF-viable in each jurisdiction.

Table 1. Supply chain finance use cases under joint blockchain–CBDC deployment.

Use case	Blockchain role	CBDC role	Reported gain	Example
Receivables financing	Authenticity, double-pledge prevention	Real-time, programmable release	Cycle days → hours	e-CNY pilot [5]
Cross-border payment	Multi-lateral shared ledger	Synchronised cross-currency settlement	2–5 days → seconds; cost ↓ 40–60%	mBridge Phase III [7]
Trade letters of credit	Full documentation on-chain	Conditional automatic payment	5–10 days → within 24 h	Singapore PBM [8]
Advance-payment financing	Logistics tracking	Targeted fund release	Risk-controlled	Bahamas Sand Dollar (illustr.)

3.1 Singapore: Purpose Bound Money and a Layered Stack

MAS has anchored CBDC innovation in two flagship programmes, Project Ubin (wholesale) and Project Orchid (retail and programmable money) [8]. The architecture has three features. First, a layered ecosystem: wholesale CBDC for inter-bank settlement, tokenised bank liabilities for commercial-bank money on chain, and regulated stablecoins for specific commercial use cases, with PBM as the universal programmable layer across all three [8]. Second, PBM proper: an issuer attaches conditions—payee, purpose, expiry—at the protocol layer rather than the application layer; the Amazon–HSBC pilot used PBM to tokenise accounts payable, letting merchants obtain unsecured working capital against the resulting digital receivables [8]. Third, wholesale CBDC live testing for inter-bank settlement, with planned extension to cross-border securities settlement.

The pattern is design-driven: the regulator participates in protocol design rather than waiting for market consolidation, shaping standards while the ecosystem is still small and tractable.

3.2 European Union: Privacy at the Core

The ECB has pursued the digital euro since 2021 with privacy and data protection as governing constraints [6]. Public consultations identified privacy as the top user concern, and the design embeds GDPR-aligned protections including offline transaction modes for low-value retail payments. Programmability is deliberately limited compared with Singapore's PBM, emphasising payment sovereignty and inclusion over conditional spending [6]. While the retail digital euro receives most public attention, the ECB and several national central banks (notably Banque de France) have run wholesale CBDC experiments for cross-border and securities settlement [11]. The pattern is compliance-first: innovation must fit existing legal frameworks, not reshape them.

3.3 United States: Market-Driven, Regulation-First

The Federal Reserve has moved cautiously on retail CBDC while encouraging private-sector innovation [9]. The U.S. has prioritised establishing a regulatory framework for digital assets, including stablecoins, through executive orders and Treasury policy with consumer protection, financial stability, and illicit-finance prevention as stated priorities [9]. Innovation in blockchain payments and SCF is driven by technology companies and financial institutions; the Canton Network targets wholesale settlement and trade-finance use cases. The Fed has not committed to a digital dollar, citing the need for Congressional authorisation [9]. The market-driven approach exploits private-sector dynamism but carries fragmentation risk: without mandated interoperability standards, parallel platforms can proliferate without converging.

3.4 Comparative Insights

Table 2 summarises how each pattern translates into SCF-relevant outcomes. The three paths represent stable points in the design space rather than transient choices. A fourth pattern—China's e-CNY rollout, with its centralised governance and large user base [5]—supplies the reference for Sect. 5's retail-CBDC scenario.

Table 2. Comparison of global CBDC exploration paths and their SCF implications.

Region	Core characteristic	Focus in SCF	Regulatory orientation
Singapore	Layered ecosystem; PBM	Enterprise programmable financing; wholesale settlement	Proactive design; protocol-level regulator involvement
European Union	Privacy first; limited programmability	Retail SME payment; wholesale pilots (Banque de France)	Compliance-first; strong legal framework
United States	Market-driven; private-sector-led	Blockchain SCF platforms; future regulated stablecoins	Sandbox guidance; risk-control emphasis

4 Paradigm Reconstruction: Towards a New Financial System

CBDC–blockchain integration is not an incremental upgrade to payment rails but a reshaping of how the financial system is governed, audited, and connected

across borders. Four features define the emerging paradigm: intelligence, trustworthiness, inclusion, and cross-border functionality. Each is illustrated below with an SCF hook—since SCF remains our empirical anchor—and each maps to a measurable AB2023 indicator in Sect. 5. The institutional preconditions for the paradigm (monetary classification as M0 [10]; legal-tender status under amended central-bank acts as in eNaira and the e-CNY [5]; data-protection regimes that constrain programmability under GDPR [6] or accept it under PBM [8]; and financial-stability holding caps such as the ECB's ∼3,000 retail cap [6] or the e-CNY's tiered identity-keyed limits [5]) shape Sect. 5 Scenario A's adoption ceiling directly.

4.1 Intelligent Finance

CBDC carrying spending conditions in the protocol layer makes conditional payment routine [14]: a trade-finance loan releases automatically when a bill of lading hashes onto the ledger; a green-bond disbursement flows only to certified low-carbon suppliers [10]. Both patterns hit Table 1's accounts-receivable-financing and milestone-triggered-LC use cases. Anonymised, aggregated CBDC transaction data, processed through statistical and ML pipelines, can give central banks early warning of liquidity stress at the sectoral or regional level [3]—shifting monetary aggregates from a coarse unit-of-account variable to a finer addressable grain, which makes targeted SCF relief programmes feasible in days rather than quarters.

4.2 Trustworthy Finance

Legacy finance trust rests on the reputation and regulatory backing of centralised institutions: banks, clearinghouses, rating agencies. Blockchain layers cryptographic trust on top, reducing what each institution carries single-handedly [2]. Key trade and settlement records write to a shared ledger that no single participant rewrites unilaterally [18]; this addresses the asymmetries embedded in asset-backed securities, where buyers historically had to trust originators on receivable quality. With receivables on chain, the verification cost shifts from the originator's reputation to cryptographic attestation. Through zero-knowledge proofs and verifiable credentials, users prove specific attributes—creditworthiness, regulatory status—without exposing underlying data [14], so SMEs demonstrate credit history without surrendering full transaction logs and corporate buyers verify supplier compliance without aggregating supplier data centrally.

4.3 Inclusive Finance

Financial inclusion remains binding in much of the global South where physical banking is sparse [1]. A CBDC wallet runs on a smartphone or, in low-bandwidth deployments, a feature phone with a SIM-based hard wallet. The Bahamas Sand Dollar was designed precisely to serve residents across more

than 700 islands where conventional banking is uneconomic [1]; Nigeria's eNaira targets the unbanked. CBDC transaction data, combined with other digital footprints—agricultural IoT readings, e-commerce histories, mobile-money records—supports credit assessment for smallholder farmers and informal businesses that legacy bank systems cannot rate [19]. For SCF this opens working-capital credit down to the long tail of small suppliers, conditional on the data-protection rules covered above.

4.4 Cross-Border Finance

Cross-border payment is the single most inefficient global-finance segment: settlement still takes days, retail remittance costs run 6–11% of principal, reconciliation across correspondent-banking chains is the dominant error source [3]. mBridge lets participating banks exchange CBDC directly across the platform's shared ledger, bypassing correspondent intermediaries [7]: settlement compresses from days to seconds and aggregate cost falls by 40–60%. Project Agorá extends the architecture to a tokenised correspondent-banking layer [11]. For SCF this collapses the asynchronous-LC class of transactions to synchronous settlement, removing principal risk from a class that has carried it for centuries. Currency competition follows: a CBDC participating in mBridge or Agorá inherits the platform's reach, while a non-participating one carries a structural disadvantage in trade settlement.

5 Quantitative Calibration via the AB2023 Tokenomic Model

The four paradigm goals of Sect. 4 are aspirational until tied to measurable parameter regimes. Each goal is mapped here to a CBDC policy archetype, each archetype to AB2023 parameter settings [14], each parameter setting to a 100-month trajectory. The four scenarios should be read as calibrations of policy intent, not predictions of any specific deployment.

5.1 Shared Simulation Framework and Scope

The framework, parameter values, and CSV outputs reported below are shared with the companion paper Hu and Bogdanov (2026) [17], which discusses the same AB2023 implementation from a technical-validation standpoint. The present section re-interprets the four scenarios as policy parameterisations of Sect. 4's paradigm goals rather than as evidence about consensus or scalability per se. Raw CSV data and source code are supplementary material common to both papers.

The AB2023 framework 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 with mint and burn; (iv) a stake-weighted token-distribution dynamic; (v) decentralisation indicators—Active Node Index (ANI),

Network Diversity Index (NDI), Consortium Influence Index (CII), Token Distribution Inequality (TDI, Gini)—aggregated into composite metric D; (vi) an F-BFT throughput model anchored to Bogdanov et al. [15]; (vii) a Monte Carlo collusion-attack module. We re-implemented these modules in Python (NumPy and Pandas only, ≈550 lines) following the equations of [14]. Two simplifications are explicit: a geometric Brownian motion replaces the full equilibrium token-price formulation, and cross-cluster network latency is omitted. Reproducibility is anchored at random seed 20240508.

The throughput module deserves a separate note since it carries the empirical weight of the section. We anchor its small-cluster behaviour to the asynchronous-stream measurements of [15] (CTPS in 0.6–1.7 for $n \in \{6, 12, 15, 18, 21\}$). Outside that range, N nodes partition into $\lceil N/15 \rceil$ clusters with arbiter-ring coordination cost $1/(1 + (k/k^*)^2)$ at $k^* = 12$. A 5.5× hardware multiplier projects from the Docker test bed to production infrastructure, consistent with the 5–6× H-NET latency reduction reported in [16]. Sharding adds $0.75 \times (n_{\text{shards}} - 1)$.

5.2 Mapping Policy Archetypes to Scenarios

The four scenarios (Table 3) tie back to Sects. 3 and 4. Scenario A (e-CNY archetype, $N_{\max} = 115$, $U_\infty = 80\,\text{K}$) models a tightly governed retail rollout where the central bank retains direct node-admission control [5]. Scenario B (mBridge / Agorá archetype, $N_{\max} = 600$, sharding at month 42, H-NET segmentation at month 66) models active expansion with explicit interoperability and gradual decentralisation goals [7]. Scenario C (Singapore PBM / tokenised-

Table 3. 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}, T_{seg} sharding and H-NET segmentation activation months; p malicious-node share (D only).

Parameter	A: Retail	B: Wholesale	C: High-adoption SCF	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	—	—	—	$\{0.05, 0.10, \ldots, 0.40\}$
Penalty regimes	—	—	—	1× standard, 2× slashing

receivables, $U_\infty = 1.2\,\mathrm{M}$, sharding at month 50) tests whether consensus remains the binding constraint under aggressive demand-side pressure [8]. Scenario D freezes Scenario B at month 60 (498 nodes, 425 K users) and sweeps malicious-share $p \in [0.05, 0.40]$ under two penalty regimes, quantifying the economic-incentive boundary at which CBDC governance deters collusion without off-chain enforcement.

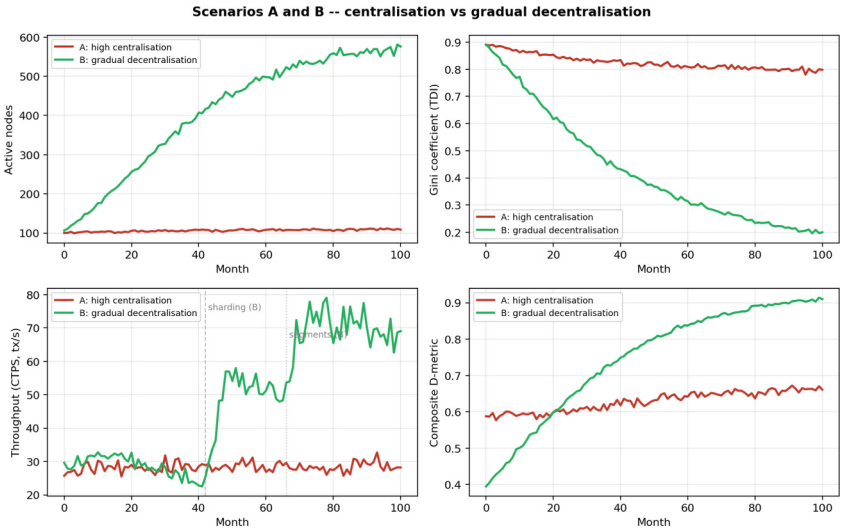

Fig. 1. Four-panel comparison of Scenarios A and B: (a) active nodes; (b) Gini coefficient; (c) CTPS, with sharding (month 42) and H-NET segment (month 66) marked; (d) composite D-metric. (Color figure online)

5.3 Time-Series Results

Scenario A. The network stabilises at 108 ± 4 active nodes (Fig. 1a). Gini drifts $0.89 \rightarrow 0.80$ (-10%); CII holds near 0.60; CTPS settles in 26–32; composite D rises $0.59 \rightarrow 0.66$. The trajectory matches the e-CNY profile in [5], where the two-tier issuance arrangement preserves the central bank as sole issuer and authorised intermediaries handle distribution.

Scenario B. Active nodes rise $100 \rightarrow 575$ (Fig. 1a, green). Gini falls $0.89 \rightarrow 0.20$ (-77%); CII falls $0.65 \rightarrow 0.11$ (-83%). CTPS traces a non-monotone path: an initial $\approx$30 plateau (months 0–24), a slow decline to $\approx$27 at month 36 as arbiter-ring overhead grows, sharding lifting it to $\approx$54 by month 48, H-NET segmentation carrying it to $\approx$73 by month 72, end-of-horizon settling at 68. Composite D rises $0.40 \rightarrow 0.91$ ($+128\%$): node-count expansion translates into substantive decentralisation rather than nominal participation. Milestones are in Table 4. The CTPS minimum at month 36 is not a model artefact—it is the

predictable consequence of arbiter-ring overhead growing faster than per-cluster throughput in the federated regime [16].

Table 4. 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

Scenario C. The user base scales 50 K $\rightarrow$ 1.32 M around month 50 (S-curve); node count expands to 301 by month 100 as financial institutions, logistics providers, and corporate treasurers enter the platform supporting PBM-style conditional payments and accounts-receivable tokenisation. CTPS at month 100 is $\approx$52, tracking Scenario B at the same node count. Throughput here is governed by node count and consensus topology, not user load: even at >1 M active wallets the binding constraint is federated cluster topology, not per-transaction consensus load. SCF deployments scaling toward Singapore-PBM volumes need horizontal scaling (sharding, segmentation) more than BFT optimisation.

5.4 Governance Under Coordinated Attack: Scenario D

The trustworthiness paradigm of Sect. 4.2 rests on the claim that economic incentives deter collusion within a CBDC ledger. Scenario D quantifies the boundary at which that claim breaks down. We freeze the Scenario B network at month 60 (498 nodes, 425 K users, mature token supply) and sweep malicious-share p from 0.05 to 0.40 under two penalty regimes. For each $(p, \text{penalty})$ 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 $0.20{+}0.6\,p{\times}\text{penalty}$. Compounded attack-success accounts for repeated attempts before detection closes the window. Attacker cost equals stake locked in malicious nodes plus an opportunity-cost adjustment scaled by the penalty multiplier.

Three observations follow from Table 5 and Fig. 2. (i) 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% success probability at a cost of 1.78 M tokens ($\approx$0.18% of total supply T_0); doubling slashing drops success to 1.3% and raises cost to 2.57 M tokens. The penalty multiplier roughly halves attack probability across the entire sub-threshold range without shifting the threshold

itself. (ii) An absorbing transition occurs near $p = 0.35$: consensus failure is deterministic, attack probability saturates near 0.85 in both regimes, economic deterrence loses traction. Rational adversaries who maximise expected return are deterred below the threshold; well-funded adversaries who tolerate cost can disable consensus altogether. Above $p = 0.30$, technical Sybil safeguards—node identity verification, geographically distributed onboarding, stake caps—are non-substitutable. (iii) Time-to-detection runs 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. AI monitoring infrastructure must operate within roughly a quarter-year detection budget under high-effort attack, giving the AI–CBDC research agenda a quantitative target rather than a qualitative wish.

5.5 Validation Against Published F-BFT Measurements

To establish empirical credibility we compare model predictions in the small-cluster Docker regime against the asynchronous-stream test of [15], with the production hardware multiplier disabled (1.0×). Results are in Table 6: the model

Table 5. Attack outcomes by malicious share p and penalty regime. TTD: time-to-detection; em-dash denotes consensus failure beyond the BFT threshold.

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
0.10	0.001	0.99	0.000	1.42	3.85	3.12
0.15	0.005	1.48	0.001	2.13	3.45	2.63
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	—	—

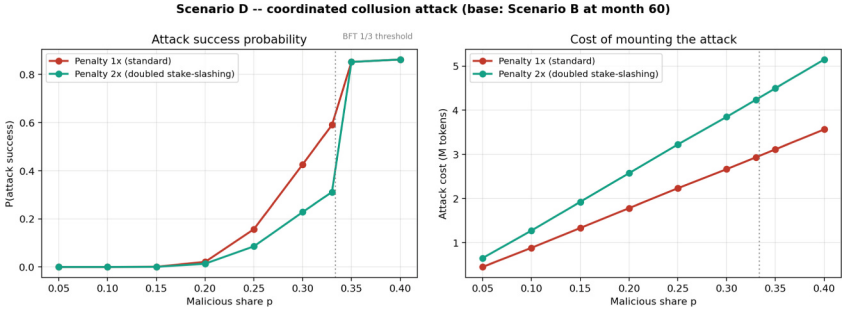

Fig. 2. Attack-success probability and attacker cost versus malicious share p, two penalty regimes, with the BFT 1/3 threshold marked.

agrees with measurement to within 21% for $n \leq 18$ and within 12% on average across that range. The 59% error at $n = 21$ reflects a regime our model does not simulate—at this cluster size the test bed reports a commit ratio of 0.28, indicating workload saturation. We note this limitation explicitly. For the cluster sizes used inside our federated regime (≈15-node clusters), the small-cluster validation is the relevant benchmark.

Table 6. Model versus published F-BFT measurements [15].

n	Measured CTPS [15]	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%

5.6 Sensitivity Analysis

We perturb the four most consequential Scenario B parameters by ±20% (or +50% for the attrition rate a_0) and re-run. End-of-horizon CTPS varies by less than 25% across the perturbation grid (Table 7); the qualitative ordering of regimes is preserved. The model is most sensitive to $N_{\max}$—networks pushed beyond the per-cluster optimum suffer disproportionately—which reinforces the design lesson that node-count expansion in a CBDC context must be paired with active scalability investment, not pursued as an end in itself.

Table 7. 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 Paradigm Reconstruction

The four scenarios connect Sect. 4's paradigm goals to AB2023 indicators on three concrete fronts. The Sect. 4.3 inclusion-versus-stability trade-off surfaces quantitatively as the contrast between Scenarios A and B: the retail-baseline regime preserves throughput stability with a Gini that barely moves over 100 months (0.89 → 0.80), while the wholesale-distributed regime drives Gini to 0.20 at the price of three years of throughput stagnation before sharding rescues performance. Section 4.4's cross-border functionality goal maps onto the post-sharding regime of Scenario B (CTPS 68, Gini 0.20, CII 0.11)—a candidate operating point for multi-CBDC bridges, where the transition path matters more than the endpoint: capacity expansions must be sequenced, not deferred, or the network spends years in the CTPS valley with elevated operational risk and political exposure. Section 4.2's trustworthiness goal cannot be satisfied by economic mechanisms alone—Scenario D shows doubling slashing penalties roughly halves sub-threshold attack probability but does nothing once p exceeds the BFT 1/3 threshold, so CBDC trust architectures need three layers: token-economic deterrence (effective for rational adversaries below BFT), technical Sybil resistance (necessary above the threshold), and AI anomaly detection (latency budget ≈2–4 months). The paradigm-reconstruction agenda becomes parameterised rather than aspirational, with the AB2023 model as the link.

6 Conclusion

This paper has examined the synergistic evolution of CBDC and blockchain through three integrated lenses: pain-point resolution in supply chain finance, comparative national approaches, and quantitative calibration via the AB2023 tokenomic model. Five findings stand out. Blockchain and CBDC address complementary frictions (information versus value flow [10,13]); neither alone resolves the structural credit and verification problems blocking SME participation in trade finance. National approaches diverge along stable institutional axes rather than transient technical choices: Singapore's design-driven, ecosystem-building strategy [8], the EU's privacy-first trajectory [6], the U.S. market-driven, regulation-first stance [9] are the principal stable points in the design space. The four-feature paradigm of Sect. 4 is a parameterisable specification—each feature maps to specific AB2023 indicators in Sect. 5. AB2023 calibration delivers operational findings (consortium throughput 28–68 CTPS by regime, Gini 0.89 → 0.20 under wholesale-distributed settings, hard ceiling for stake-slashing deterrence at the BFT 1/3 threshold), with the throughput model agreeing with [15] to within 12% mean error for $n \leq 18$. Realising the paradigm requires joint advancement of technical standards (sharding, segmentation, post-quantum), legal frameworks (privacy regulation, cross-jurisdictional harmonisation), and institutional infrastructure (multi-CBDC platforms, deposit-cap policy); no single thread is sufficient and Sect. 5 quantifies the cost of each lagging behind.

6.1 Policy Implications

For jurisdictions designing CBDC–SCF integration, four recommendations follow. First, articulate policy objectives before architectural choices: Sect. 5 shows the cost of leaving the inclusion–efficiency–stability–cross-border trade-off implicit, as Scenarios A and B differ by a factor of more than two on Gini reduction (10% vs 77%) over the same horizon precisely because the policy goal differs [10]. Second, adapt rather than copy: PBM [8] suits jurisdictions whose central banks join protocol design; the EU privacy-first model [6] fits regimes where strong data-protection rules already constrain programmability. Third, invest in interoperability before fragmentation locks in—mBridge participation [7], ISO 20022 alignment, and Project Agorá architecture [11] all reduce later integration cost; deferring interoperability tends to make it permanently more expensive. Fourth, stress-test in simulation before deployment: AB2023 calibration surfaces non-obvious dynamics (the three-year throughput valley during cluster expansion, the BFT cliff in attack economics, the saturation of stake-slashing deterrence above $p \approx 0.33$) that descriptive analysis would miss. Live pilot data complement parametric simulation; they do not substitute for it.

6.2 Limitations and Future Research

Three limitations qualify the conclusions. First, the analysis is conceptual and computational, drawing on publicly available data and on a faithful but not exhaustive AB2023 re-implementation [14]; as more CBDC projects move from pilot to production, empirical validation against operational data will strengthen or refute simulation findings. Second, the AB2023 model does not capture the political-economy dimensions of CBDC adoption—distributional consequences of monetary digitalisation, redistribution of seigniorage between central and commercial banks, geopolitical implications of cross-border CBDC platforms all sit outside the parameter space. Third, the comparative analysis covers three high-income jurisdictions plus referenced examples from China, Nigeria, and the Bahamas; application to lower-income, infrastructure-constrained settings—where Sect. 4.3's inclusion paradigm matters most—requires recalibration of the user-growth and node-growth modules.

A natural research agenda follows. The most pressing item is empirical validation against pilot deployments such as e-CNY, eNaira, and Sand Dollar [1,5]. The second is the political economy of CBDC design choices, including who bears the cost of paradigm transitions when the throughput valley of Sect. 5 lasts three years. The third is technical: extension of AB2023 with cross-cluster latency modelling and a richer attack taxonomy, both consistent with the H-NET virtualisation work in [16]. CBDC–blockchain integration in supply chain finance has no single optimal design; appropriate choices emerge from the interplay of policy intent, institutional constraint, and technical capacity, with empirical validation of the AB2023 simulation against pilot deployments and cross-jurisdictional legal harmonisation as the principal 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. World Bank, Washington, DC (2024)
2. Frizzo-Barker, J., Chow-White, P.A., Adams, P.R., Mentanko, J., Ha, D., Green, S.: Blockchain as a disruptive technology for business: a systematic review. Int. J. Inf. Manage. **51**, 102029 (2020). https://doi.org/10.1016/j.ijinfomgt.2019.10.014
3. Bank for International Settlements: Annual Economic Report 2023. BIS, Basel (2023)
4. Kosse, A., Mattei, I.: Making headway — results of the 2022 BIS survey on central bank digital currencies and crypto. BIS Papers No. 136. Bank for International Settlements, Basel (2023)
5. People's Bank of China: E-CNY White Paper (2024 update). People's Bank of China, Beijing (2024)
6. European Central Bank: Report on a Digital Euro – Second Progress Report. European Central Bank, Frankfurt (2023)
7. mBridge Project Consortium: Project mBridge – Phase III Technical Implementation Report. BIS Innovation Hub, Hong Kong SAR (2024)
8. Monetary Authority of Singapore: Project Orchid — Programmable Digital SGD and the Purpose Bound Money Protocol. MAS Whitepaper (2023)
9. Board of Governors of the Federal Reserve System: Money and Payments: The U.S. Dollar in the Age of Digital Transformation. Federal Reserve Discussion Paper. Federal Reserve Board, Washington, DC (2022)
10. Kumhof, M., Noone, C.: Central bank digital currencies – design principles for financial stability. Econ. Anal. Policy **83**, 553–572 (2024). https://doi.org/10.1016/j.eap.2024.07.012
11. BIS Innovation Hub: Project Agorá — Tokenisation of cross-border payments. BIS Innovation Hub Annual Report 2023. Bank for International Settlements, Basel (2023)
12. Harish, A.R., Liu, X.L., Zhong, R.Y., Huang, G.Q.: Blockchain for logistics 4.0: a systematic literature review. Transp. Res. Part E Logist. Transp. Rev. **201**, 104269 (2025). https://doi.org/10.1016/j.tre.2025.104269
13. Song, H., Yang, Y., Tao, Z.: The application of blockchain in enterprise financing. Nankai Bus. Rev. **25**(2), 34–48 (2022)
14. Khvatov, V., Bogdanov, A.: DGT Network Tokenomics — AB2023 Model Specification. Technical Documentation Rev. 21 (2023). https://dgtnetwork.gitbook.io/dgt-docs
15. 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
16. Bogdanov, A., Khvatov, V., Uteshev, A., Shchegoleva, N.: Virtual blockchain network for secure financial and industrial applications. In: Gervasi, O., Murgante, B., Garau, C., et al. (eds.) ICCSA 2024 Workshops. LNCS, vol. 14815, pp. 142–161. Springer, Cham (2024)

17. Hu, B., Bogdanov, A.V.: Blockchain as infrastructure for CBDC: technological foundations and challenges. In: O. Gervasi et al. (eds.) ICCSA 2026. LNCS 16759, pp. xx–yy. Springer, Cham (2026)
18. Li, D., Du, X.: Research status and prospects of blockchain technology in digital currency and supply chain finance. Chin. J. Manag. Sci. (2025). https://doi.org/10.16381/j.cnki.issn1003-207x.2023.1852
19. Tang, Y., Chen, G.: Blockchain technology empowering logistics supply chain resilience: theoretical mechanism, empirical examination, and improvement paths. Bus. Econ. Res. **7**, 82–86 (2026)