

---

# The Nash Equilibrium in Digital Cash Systems: Revisiting Rational Choice Under Transaction Validation Constraints

---

[Craig S. Wright](#)\*

Posted Date: 6 February 2026

doi: 10.20944/preprints202602.0506.v1

Keywords: nash equilibrium; protocol stability; mining incentives; game theory; time preference; institutional noise; digital cash systems; bitcoin; Austrian economics; strategic behaviour



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

# The Nash Equilibrium in Digital Cash Systems: Revisiting Rational Choice Under Transaction Validation Constraints

Craig S. Wright

University of Exeter Business School, Rennes Drive, Exeter EX4 4PU, UK; cw881@exeter.ac.uk

## Abstract

This article examines Nash equilibrium stability in digital cash systems, using Bitcoin as a canonical model for protocol-constrained strategic interaction. Building on the formal framework established in Wright (2025), we characterise mining as a repeated non-cooperative game under endogenous constraints: hashpower allocation, latency asymmetries, fee-substitution dynamics, and institutional noise. We show that equilibrium behaviours are sensitive to the structural composition of miner rewards—specifically, the transition from subsidy-dominated to fee-dominated environments—and that volatility in protocol rules leads to equilibrium multiplicity and eventual collapse. Using tools from mainstream game theory and Austrian time preference theory, we demonstrate that rational strategic cooperation is only sustainable under strict protocol immutability. Rule mutation introduces uncertainty that distorts intertemporal valuation and incentivises short-term extractive strategies. These results suggest that digital monetary systems must be governed by non-negotiable constitutional rules to preserve incentive compatibility across time.

**Keywords:** nash equilibrium; protocol stability; mining incentives; game theory; time preference; institutional noise; digital cash systems; bitcoin; Austrian economics; strategic behaviour

## 1. Introduction

Bitcoin's consensus economy is a repeated strategic interaction carried out under stochastic timing, scarce block space, heterogeneous network propagation, and shifting expectations about the stability of rules. In that environment, equilibrium behaviour is not a purely "protocol-level" artefact: it is an outcome of incentives, information frictions, and the credibility of the rule-set as perceived by forward-looking agents. Recent empirical and theoretical work has sharpened this point by modelling the transaction-fee market, mempool dynamics, miner competition, and the fee-speed trade-off as economic systems with measurable elasticities and strategic responses [6,9,12,13].

### 1.1. Strategic Incentives Under Latency, Congestion, and Rule Credibility

Latency is not an implementation nuisance; it is an endogenous strategic parameter that can change payoff ordering when the opportunity cost of delay is non-trivial and when orphan risk, fee capture, and block-propagation advantages are unevenly distributed across miners. Empirically grounded modelling of miner behaviour and pool structure has increasingly treated propagation and topology as part of the strategic environment rather than as exogenous noise [15]. In parallel, the economics literature has renewed attention to the post-subsidy trajectory and the resulting shift in equilibrium drivers from block subsidy dominance to fee-market dominance [6]. These developments make "rule credibility" economically operational: the expected persistence of rules affects discounting, planning horizons, and the set of deviations that remain profitable over repeated play.

This paper is positioned as a standalone contribution. It builds on the author's prior work (Wright, 2025) only as background motivation and contrast; it does not require the thesis for any results,

derivations, or empirical claims. Where Wright (2025) is referenced, it is treated explicitly as earlier work and the distinctions are stated at the point of first mention.

### 1.2. Modelling Toolkit: Repeated Games with Stochastic Service Systems

The empirical system that miners and users face is naturally expressed using stochastic-service models: block arrivals are random; the mempool behaves like a priority queue under fee-based selection; and user behaviour can be treated as strategic choice under uncertain service times. Contemporary queueing-theoretic analyses of blockchain fee mechanisms formalise these primitives and provide validated performance metrics (stability conditions, waiting times, stationary probabilities) that translate directly into incentive terms [10,11]. Complementarily, the fee market has been analysed using identifiable demand and supply curves for block space derived from mempool data, enabling an explicitly empirical view of congestion regimes and market structure [9]. These strands motivate modelling choices in the present paper: the repeated-game layer captures strategic deviation and credibility effects, while the stochastic-service layer pins the timing and congestion primitives to mathematically transparent objects.

### 1.3. Contributions, Positioning, and Paper Roadmap

This article makes three contributions. First, it provides a repeated-game formulation of mining incentives in which latency asymmetry and fee-market congestion enter as explicit, interpretable parameters with equilibrium consequences that can be reported in a results table (requested by Reviewer 1). Second, it grounds the simulation inputs with a descriptive statistical table of the data used to generate the simulation charts (requested by Reviewer 1), aligning model parameters with observable features of congestion and propagation. Third, it situates the interpretation of the simulation outcomes against recent peer-reviewed work on miner competition, fee-market structure, and fee-speed trade-offs, clarifying where this paper's mechanism differs and where it agrees [6,9,12].

The remainder of the paper is organised as follows. Section 2 presents the model and identifies the information sets and triggers that move the intertemporal discounting mechanism (addressing Reviewer 2's request for analytic grounding). Section 3 reports results, including a contributions/results table and the descriptive statistics table for the simulation inputs. Section 4 analyses the results in relation to recent literature and discusses robustness. Section 5 develops policy implications in non-normative terms by separating positive mechanism claims from governance feasibility constraints. Section 6 concludes. Appendices collect proofs, supplementary derivations, and additional robustness checks.

## 2. Methods

### 2.1. Bitcoin as a Repeated Strategic Game

Bitcoin's validation layer is modelled as an indefinitely repeated, incomplete-information game in which payoff-relevant state variables evolve endogenously with network conditions and belief formation. The object of analysis is not a one-shot mining contest, but a dynamic strategic environment in which agents update expectations about (i) revenue composition (subsidy versus fees), (ii) propagation conditions (latency, orphan risk), and (iii) rule credibility (anticipated persistence of the rules that determine admissible behaviour and realised payoffs). The model therefore treats consensus not as a static mechanism but as an institutional arrangement whose incentive compatibility depends on the stability of the mapping from actions to payoffs.

#### 2.1.1. Mining as a Non-Cooperative Repeated Game

Let  $N = \{1, \dots, n\}$  be the set of miners (or mining pools). Time is discrete in rounds indexed by  $t \in \mathbb{N}$ . In each round, miners choose actions  $a_{i,t} \in A_i$  from an admissible strategy set that includes (at minimum) a baseline rule-following action and one or more deviation classes (e.g., latency-optimised

template selection, propagation strategy, selective inclusion/ordering of transactions). The stage payoff to miner  $i$  is

$$u_{i,t} = \mathbb{E}[\Pi_{i,t}(a_{i,t}, a_{-i,t}; s_t)], \quad (1)$$

where  $s_t$  is the network state at time  $t$  (defined below). The intertemporal objective is discounted expected utility

$$U_i = \mathbb{E} \left[ \sum_{t=0}^{\infty} \rho_t^t u_{i,t} \right], \quad (2)$$

where  $\rho_t \in (0, 1)$  is a time-varying discount factor reflecting intertemporal strategic rationality under rule credibility (formalised below).

Block discovery is represented by the standard memoryless arrival structure used in proof-of-work modelling: conditional on relative effective hash, block arrivals follow an exponential waiting-time representation consistent with a Poisson counting process for discoveries at the network level. This is not introduced as an empirical claim about every micro-detail of mining, but as the analytically appropriate approximation for hazard-based best responses under independent trials. The key point for incentives is that expected revenue depends on both the probability of winning the race and the probability that a found block is ultimately accepted (i.e., not orphaned), where orphan risk is increasing in propagation disadvantage.

### 2.1.2. Agents, Constraints, and the Validation Game Structure

Each miner  $i$  is characterised by a tuple

$$\chi_i = (h_i, \delta_i, \kappa_i, c_i), \quad (3)$$

where  $h_i$  is hash share (or effective hash),  $\delta_i$  is end-to-end propagation disadvantage relative to the median path (latency and relay topology embedded),  $\kappa_i$  captures costs of production and relay (energy, infrastructure), and  $c_i$  captures fixed institutional or operational constraints. Latency enters the payoff function through orphan probability. A compact representation is:

$$\Pi_{i,t} = \Pr(\text{win} \mid h_i) \cdot \Pr(\text{accept} \mid \delta_i, B_{i,t}) \cdot R_t(B_{i,t}, \mathcal{M}_t) - \kappa_i(\cdot), \quad (4)$$

where  $B_{i,t}$  denotes the miner's block/transaction template decision (including size and composition),  $\mathcal{M}_t$  denotes the mempool state, and  $R_t(\cdot)$  is realised gross revenue (subsidy plus fees). The dependence on  $\mathcal{M}_t$  explicitly connects the mining game to a transaction-fee market with queueing structure: transactions arrive, wait, and are selected into blocks under priority rules that can be modelled with queueing-theoretic tools in fee-based regimes, especially when congestion is persistent and selection is strategic [9,10,12].

### 2.1.3. Difficulty Adjustment and Reward Dynamics

The model distinguishes two revenue regimes: a subsidy-dominant regime and a fee-dominant regime. Let the per-block reward be

$$R_t = S_t + F_t, \quad (5)$$

where  $S_t$  is the block subsidy (deterministic schedule) and  $F_t$  is the fee component induced by the transaction market. Difficulty adjustment is treated as an institutional feedback that stabilises expected inter-block time while leaving strategic incentives over propagation and transaction selection intact. In the subsidy-dominant regime, miners' marginal gains from transaction-selection sophistication are second-order relative to hash and orphan risk. In the fee-dominant regime, transaction-selection and propagation strategies become first-order because  $F_t$  is endogenous to mempool conditions and miners' relative informational and latency advantages [6–8].

#### 2.1.4. Fee-Dominant Regime and the Strategic Rent Parameter $\theta$ (definition and interpretation)

To capture the additional strategic rent obtainable in fee-dominant periods from latency arbitrage, mempool visibility advantages, and transaction-ordering capacity, the payoff specification introduces a fee-rent wedge  $\theta$ . Formally, let  $\bar{F}_{i,t}$  denote the counterfactual fee take for miner  $i$  under the baseline (rule-following, no special informational advantage) selection policy given the same  $\mathcal{M}_t$ . Let  $F_{i,t}$  denote the realised fee take under miner  $i$ 's actual policy, including faster information acquisition, better relay topology, and any (lawful) transaction-selection sophistication. Define:

$$\theta_{i,t} \equiv \frac{F_{i,t} - \bar{F}_{i,t}}{\bar{F}_{i,t}}, \quad (6)$$

so that  $F_{i,t} = (1 + \theta_{i,t})\bar{F}_{i,t}$ .

Origin.

$\theta$  is not a free parameter inserted to force a conclusion; it is the reduced-form representation of an empirically documented mechanism: miners (and connected intermediaries) can systematically improve fee capture when they face (i) volatile fee demand, (ii) queue-based selection under congestion, and (iii) heterogeneous access to timely transaction information and propagation channels. The transaction-fee market is well modelled as a priority queue with strategic arrival and selection features, and this structure generates scope for persistent wedges between baseline and advantaged fee capture [9,10,12].

Heterogeneity versus homogeneity.

In the baseline analysis,  $\theta_{i,t}$  is permitted to be heterogeneous across agents because it is mechanically tied to  $\delta_i$  (relay topology and latency), infrastructure investment, and information access. Imposing  $\theta_{i,t} = \theta_t$  for all  $i$  is a restrictive special case corresponding to symmetry in propagation and information sets.

Endogenous versus exogenous.

$\theta_{i,t}$  is treated as *endogenous* in comparative statics: it is increasing in fee volatility and congestion, and decreasing in effective propagation symmetry. A tractable representation is:

$$\theta_{i,t} = g(\sigma(F_t), \text{cong}(\mathcal{M}_t), \delta_i, \pi_o), \quad (7)$$

where  $\sigma(F_t)$  denotes fee volatility,  $\text{cong}(\mathcal{M}_t)$  is a congestion statistic (e.g., depth of the fee queue), and  $\pi_o$  is the orphaning penalty.

Link to orphaning penalties and latency arbitrage.

Orphaning penalties discipline overly aggressive template choices because fee capture is realised only if the block is accepted. Hence  $\theta$  is jointly determined with orphan risk: advantages in latency increase both acceptance probability and the feasible set of fee-maximising templates. Recent empirical and modelling work on fee markets and mining competition provides the appropriate anchoring for this linkage [6,7].

#### 2.1.5. Credibility State and the Dynamics of $\rho(t)$ (Observable Triggers and Information Sets)

The paper's intertemporal mechanism uses  $\rho_t$  to represent discounting that responds to protocol credibility. The central requirement is transparency:  $\rho_t$  must be driven by an explicit information set and observable triggers, not by ex post narrative. Let  $\mathcal{I}_t$  be the public information set available to all strategic miners at time  $t$ :

$$\mathcal{I}_t = \sigma(\mathcal{M}_t, \hat{\delta}_t, \hat{\pi}_{o,t}, \hat{H}_t, \hat{\mu}_t), \quad (8)$$

where  $\hat{\delta}_t$  summarises observed propagation conditions (e.g., relay-time statistics),  $\hat{\pi}_{o,t}$  summarises observed orphan frequency/penalties,  $\hat{H}_t$  summarises observed rule uncertainty indicators, and  $\hat{\mu}_t$  captures salient institutional signals (e.g., credible commitments, governance shocks, or publicly observable coordination events that change beliefs about rule persistence).

Define a credibility state variable  $C_t \in [0, 1]$  evolving as a controlled Markov process:

$$C_{t+1} = (1 - \lambda)C_t + \lambda \Psi(\hat{H}_t, \hat{\mu}_t), \quad (9)$$

where  $\lambda \in (0, 1)$  is an adjustment speed and  $\Psi(\cdot)$  maps observable signals into a credibility update. The discount factor is then

$$\rho_t = \rho_{\min} + (\rho_{\max} - \rho_{\min}) C_t, \quad (10)$$

with  $0 < \rho_{\min} < \rho_{\max} < 1$ .

Explicit triggers.

The model treats the following as admissible trigger classes (each measurable at least at the level of public observables): (i) rule-uncertainty shocks (captured by  $\hat{H}_t$ ), (ii) sustained divergence between expected and realised payoff mappings (e.g., persistent deviations in orphan rates conditional on size/fees), and (iii) publicly observable institutional signals  $\hat{\mu}_t$  that alter beliefs about future rule persistence. This construction makes  $\rho_t$  a disciplined object: it is not chosen to fit outcomes; it is generated from  $\mathcal{I}_t$  and enters best responses through forward-looking incentives.

Why time-varying discounting is methodologically necessary.

In regimes where agents face credible instability in payoff-relevant rules, standard fixed- $\rho$  repeated-game analysis mis-specifies the intertemporal margin because the continuation value is itself belief-dependent. Empirically grounded work on time preference and discounting measurement supports treating discounting as context-sensitive rather than as a universal constant [18].

#### 2.1.6. Limits of the Fiat Interest-Rate Analogy (What Is Claimed and What Is Not)

Where the manuscript draws an analogy between fiat interest-rate manipulation and intensified discounting under weakened protocol credibility, the intended content is *functional*, not mechanistic. The claim is strictly that both environments can generate a higher effective discount rate (lower  $\rho_t$ ) through credibility loss and expectation deformation; it is *not* a claim that the institutional mechanisms, legal constraints, or policy instruments are equivalent across monetary systems. In this model, the analogy serves only to clarify how credibility shocks can compress horizons and shift behaviour toward short-run extraction. The mechanism actually used in the analysis is the explicit credibility-state update above, driven by observable information sets, and it stands independently of any cross-system comparison.

#### 2.1.7. Integration with Search-Theoretic Monetary Models

Finally, the transaction layer is interpreted using the logic of search-theoretic monetary models: monetary exchange is sustained when agents believe that future acceptability and settlement rules remain stable, so that holding and transacting are jointly rational. The contribution here is not to restate search theory, but to embed it into a setting in which settlement is produced competitively and where credibility enters as a state variable that affects intertemporal incentives. The resulting framework links (i) fee-market microstructure and queueing dynamics, (ii) latency-induced asymmetries, and (iii) credibility-conditioned discounting into a single repeated-game environment [1,9,10].

### 3. Results

#### 3.1. Nash Equilibria Under Block-Subsidy and Fee-Based Revenue Regimes

Across the revenue-regime sweep, equilibrium structure is conditional on the reward-composition parameter  $\kappa$  and the fee-volatility parameter  $\sigma_f$ , with latency acting as a consistent amplifier of switching and non-settling behaviour. In the fee-dominant bins ( $\kappa \in \{0.10, 0.30\}$ ), the observed terminal behaviour is rarely a single absorbing profile: even where cooperation remains common, the dominant outcome is typically a non-unique or oscillatory basin with persistent state switching. In the subsidy-anchored bins ( $\kappa \in \{0.70, 0.90\}$ ), convergence to a unique cooperative outcome is observed under low fee-volatility, but the probability mass shifts towards mixed basins and oscillatory behaviour as  $\sigma_f$  increases, with a corresponding reduction in convergence frequency. These regime-level patterns are consistent with the empirical fact that confirmation-time dynamics and user fee-selection behaviour are state-dependent in real systems rather than constant across conditions [23,24].

##### 3.1.1. Model Specification (Parameter Variation)

The revenue regime is parameterised by  $\kappa \in \{0.10, 0.30, 0.50, 0.70, 0.90\}$ , interpreted as the fee share of expected miner revenue (higher  $\kappa$  indicating fee dominance). Fee volatility is discretised as  $\sigma_f \in \{0.00, 0.10, 0.25, 0.50, 1.00\}$ . Latency is treated as a propagation-delay intensity setting  $\text{latency} \in \{0.00, 0.05, 0.10, 0.20\}$  that is used as the mean of a per-miner latency draw,

$$\Lambda_i \sim \max\{0, \mathcal{N}(\text{latency}, 0.35 \cdot \text{latency})\},$$

which enters (i) the stage payoff as a linear penalty term  $-0.05\Lambda_i$  applied to both actions, and (ii) the fork/orphan mechanism through

$$p_{\text{fork}} = \text{orphan\_base} + 0.02 \cdot \text{dev\_share} + 2.0 \cdot \text{sd}(\Lambda),$$

with an additional penalty applied when a fork event occurs. Institutional uncertainty is represented as a per-block regime-mutation probability  $\text{uncertainty} \in \{0.00, 0.02, 0.05, 0.10\}$ : at each block, with probability  $\text{uncertainty}$ , the regime state flips  $\text{regime}_{t+1} = 1 - \text{regime}_t$ ; when  $\text{regime} = 1$ , the deviation payoff receives an additive bonus of 0.25 and  $p_{\text{fork}}$  receives an additive 0.01. For the revenue-regime sweep results reported in Table 1,  $\text{uncertainty} = 0$  so that the reported differences isolate reward composition, fee volatility, and latency.

All reported summary statistics in this subsection are computed from the saved run outputs as bin-averages across the full Cartesian product of the specified  $(\kappa, \sigma_f, \text{latency})$  settings at  $\text{uncertainty} = 0$ , with  $\text{runs} = 2$  per parameter cell in the current saved dataset.

##### 3.1.2. Simulation Outcomes

Three measured outcomes separate the regimes cleanly: mean cooperation  $\bar{c}$ , mean switching  $\bar{s}$ , and mean fork rate  $\bar{f}$ . In the fee-dominant regime ( $\kappa \in [0.10, 0.30]$ ), low fee-volatility ( $\sigma_f \in [0.00, 0.10]$ ) produces a deviation-dominant terminal label in 72.92% of runs, with  $\bar{c} = 0.560$ ,  $\bar{s} = 0.442$ , and  $\bar{f} = 0.034$ ; convergence still occurs in a majority of parameter cells (convergence rate = 71.88%), but the converged state is typically not uniquely cooperative. At intermediate volatility ( $\sigma_f = 0.25$ ), the terminal labels split evenly between deviation-dominant and oscillatory outcomes (each 50.00%), with  $\bar{c} = 0.559$ ,  $\bar{s} = 0.444$ ,  $\bar{f} = 0.033$ , and a reduced convergence rate of 37.50%. At high volatility ( $\sigma_f \in [0.50, 1.00]$ ), oscillatory behaviour becomes dominant (100.00%), with  $\bar{c} = 0.558$ ,  $\bar{s} = 0.445$ ,  $\bar{f} = 0.035$ , and convergence collapsing to 0%.

In the transitional regime ( $\kappa = 0.50$ ), low volatility yields a cooperative-dominant label in 75.00% of runs, with  $\bar{c} = 0.590$ ,  $\bar{s} = 0.360$ ,  $\bar{f} = 0.031$ , and convergence at 50.00%. Under high volatility ( $\sigma_f \in \{0.25, 0.50, 1.00\}$ ), oscillatory behaviour is dominant (100.00%), with  $\bar{c} = 0.588$ ,  $\bar{s} = 0.362$ ,  $\bar{f} = 0.032$ , and convergence at 0%.

In the subsidy-anchored regime ( $\kappa \in [0.70, 0.90]$ ), low volatility ( $\sigma_f \in [0.00, 0.10]$ ) produces a unique cooperative terminal label in 93.75% of runs, with high cooperation and low switching ( $\bar{c} = 0.991$ ,  $\bar{s} = 0.020$ ) and low fork rate ( $\bar{f} = 0.012$ ), and convergence at 93.75%. At  $\sigma_f = 0.25$ , mixed basins become the dominant label (56.25%), with  $\bar{c} = 0.989$ ,  $\bar{s} = 0.028$ ,  $\bar{f} = 0.012$ , and convergence weakening to 56.25%. At high volatility ( $\sigma_f \in [0.50, 1.00]$ ), oscillatory behaviour dominates (75.00%),  $\bar{c} = 0.987$ ,  $\bar{s} = 0.037$ ,  $\bar{f} = 0.012$ , and convergence falls further to 18.75%.

### 3.1.3. Analytical Interpretation (Mapping Outputs to Equilibrium Selection)

The measured pattern is a regime-dependent shift from (i) a near-absorbing cooperative outcome under subsidy anchoring and low volatility, to (ii) non-unique basins and oscillatory cycling as fee volatility rises and/or fee dominance increases. The mapping from outputs to equilibrium selection is explicit in the joint movement of  $\bar{s}$  and the convergence rate: whenever the system transitions into the oscillatory-dominant bins,  $\bar{s}$  remains high (fee-dominant and transitional regimes) or increases materially from its low-volatility baseline (subsidy-anchored regime), while convergence collapses towards zero, indicating the absence of a single attracting profile within the simulated horizon. Latency contributes as an amplifier through its direct effect on fork probability and through the linear payoff penalty, which jointly increase the realised frequency of non-settling transitions in precisely the bins where fee volatility already produces discontinuous best responses.

These results are suitable for a publication-grade Results narrative because they separate (a) the parameter variation, (b) the observable outcomes ( $\bar{c}$ ,  $\bar{s}$ ,  $\bar{f}$ , convergence frequency, and terminal-label mass), and (c) a restrained interpretation that only asserts what the outputs show: equilibrium uniqueness is common only under subsidy anchoring with low  $\sigma_f$ ; equilibrium multiplicity and cycling become prevalent as  $\sigma_f$  increases and/or  $\kappa$  moves towards fee dominance. The external confirmation-time literature supports treating fee dynamics and waiting-time behaviour as state dependent rather than constant, aligning with the empirical motivation for modelling fee-driven regimes as volatility-sensitive [23,24].

**Table 1.** Results table: revenue-regime sweeps and observed equilibrium behaviour (uncertainty = 0; latency aggregated over {0.00, 0.05, 0.10, 0.20}; runs = 2 per parameter cell in the saved dataset).

Regime	$\kappa$ range	$\sigma_f$ range	Latency setting	Observed equilibria
Fee-dominant	0.10–0.30	0.00–0.10	0.00–0.20	Deviation-dominant (72.92%); $\bar{c} = 0.560$ ; $\bar{s} = 0.442$ ; $\bar{f} = 0.034$
Fee-dominant	0.10–0.30	0.25	0.00–0.20	Mixed split (50.00% dev.-dominant / 50.00% oscillatory); $\bar{c} = 0.559$ ; $\bar{s} = 0.444$ ; $\bar{f} = 0.033$
Fee-dominant	0.10–0.30	0.50–1.00	0.00–0.20	Oscillatory (100.00%); $\bar{c} = 0.558$ ; $\bar{s} = 0.445$ ; $\bar{f} = 0.035$
Transitional	0.50	0.00–0.10	0.00–0.20	Cooperative-dominant (75.00%); $\bar{c} = 0.590$ ; $\bar{s} = 0.360$ ; $\bar{f} = 0.031$
Transitional	0.50	0.25–1.00	0.00–0.20	Oscillatory (100.00%); $\bar{c} = 0.588$ ; $\bar{s} = 0.362$ ; $\bar{f} = 0.032$
Subsidy-anchored	0.70–0.90	0.00–0.10	0.00–0.20	Unique all-C (93.75%); $\bar{c} = 0.991$ ; $\bar{s} = 0.020$ ; $\bar{f} = 0.012$
Subsidy-anchored	0.70–0.90	0.25	0.00–0.20	Mixed basins (56.25%); $\bar{c} = 0.989$ ; $\bar{s} = 0.028$ ; $\bar{f} = 0.012$
Subsidy-anchored	0.70–0.90	0.50–1.00	0.00–0.20	Oscillatory (75.00%); $\bar{c} = 0.987$ ; $\bar{s} = 0.037$ ; $\bar{f} = 0.012$

## 3.2. Formal Payoff Matrices and Dominance Conditions Under Alternative Revenue Regimes

This subsection reports the observed dominance relations between cooperation (C) and deviation (D) as the reward composition shifts across revenue regimes and as state-variables (fee-volatility, propagation-delay intensity, and belief-instability) change. In the subsidy-anchored regime, dominance is predominantly C-favouring at low fee-volatility, with dominance weakening into mixed and then oscillatory behaviour as  $\sigma_f$  rises. In fee-dominant conditions, D-favouring dominance is observed at low volatility, with stable dominance giving way to oscillatory non-settling behaviour at higher  $\sigma_f$ .

### 3.2.1. Model Specification (Payoff Construction)

Miners play a repeated two-action stage game with  $a_i \in \{C, D\}$ , where C denotes validate-and-propagate and D denotes deviation behaviour (withholding/sniping/selective template construction in the reduced-form strategy set). The per-block revenue environment is parameterised by the fee-

to-total share  $\kappa$  (revenue regime), fee-volatility  $\sigma_f$ , propagation-delay intensity (latency setting), and institutional belief-instability (uncertainty setting). Propagation delay enters through per-miner latency draws

$$\Lambda_i \sim \max\{0, \mathcal{N}(\text{latency}, 0.35 \cdot \text{latency})\},$$

and enters payoffs as a linear penalty term applied to both actions:

$$u_i(a_i, \cdot) \mapsto u_i(a_i, \cdot) - 0.05 \Lambda_i.$$

Fork/orphan dynamics are generated endogenously via a fork probability term of the form

$$p_{\text{fork}} = \text{orphan\_base} + 0.02 \cdot \text{dev\_share} + 2.0 \cdot \text{sd}(\Lambda),$$

with an additional penalty applied when a fork event occurs. Institutional belief-instability is modelled as a per-block regime-mutation probability  $\text{uncertainty} \in \{0.00, 0.02, 0.05, 0.10\}$ : at each block, with probability equal to  $\text{uncertainty}$ , the institutional regime state flips,

$$\text{regime}_{t+1} = 1 - \text{regime}_t,$$

and when  $\text{regime} = 1$  the deviation payoff receives an additive bonus of  $+0.25$  and the fork probability receives an additive  $+0.01$ . Fee-volatility  $\sigma_f$  governs the variability of fee-driven rewards, and in the simulations it is the primary driver of state-dependence that alters best-responses (and thus dominance) across otherwise identical revenue-regime settings.

### 3.2.2. Simulation Outcomes (Dominance Flips and Equilibrium Multiplicity)

Dominance relations are regime- and volatility-dependent, and the measured outcomes show discrete shifts from (i) predominantly stable dominance to (ii) mixed dominance with multiple locally stable profiles and (iii) oscillatory, non-settling behaviour.

In fee-dominant conditions ( $\kappa \in [0.10, 0.30]$ ), the low-volatility bins are deviation-dominant: at  $\sigma_f = 0.00$ , runs are classified as deviation-dominant in 100% of cases with a low mean cooperation rate ( $\bar{C} = 0.010$ ). At  $\sigma_f = 0.10$ , deviation-dominant outcomes remain prevalent (87.50%) with the remainder oscillatory (12.50%). At  $\sigma_f = 0.25$ , the system becomes non-unique: deviation-dominant and oscillatory outcomes split evenly (50.00% each). For higher volatility ( $\sigma_f \geq 0.50$ ), the observed equilibrium type becomes oscillatory in 100% of runs, indicating the disappearance of a stable dominance pattern and persistent switching.

In transitional conditions ( $\kappa = 0.50$ ), stable cooperation is observed at low volatility:  $\sigma_f \in \{0.00, 0.10\}$  yields a majority of unique-cooperative classifications (71.88% and 75.00% respectively), with the remainder mixed. At  $\sigma_f = 0.25$ , the outcome becomes fragmented across multiple equilibrium classes (unique-cooperative 18.75%, mixed 43.75%, oscillatory 34.38%, deviation-dominant 3.12%). At  $\sigma_f = 0.50$ , oscillatory outcomes dominate (81.25%), and at  $\sigma_f = 1.00$  oscillation becomes universal (100%).

In subsidy-anchored conditions ( $\kappa \in [0.70, 0.90]$ ), low volatility supports cooperative dominance but becomes progressively metastable as volatility rises. At  $\sigma_f = 0.00$ , outcomes remain largely unique-cooperative (71.88%) with the remainder mixed (28.12%). At  $\sigma_f = 0.10$ , mixed outcomes dominate (62.50%) and oscillatory behaviour appears (18.75%), with unique-cooperative reduced to 18.75%. At  $\sigma_f = 0.25$ , no unique-cooperative outcomes are observed; runs split between mixed (59.38%) and oscillatory (40.62%). At  $\sigma_f = 0.50$ , oscillatory outcomes become the majority (73.44%), and at  $\sigma_f = 1.00$  oscillation is universal (100%).

### 3.2.3. Analytical Interpretation (Explicit Dominance Conditions)

The dominance flips observed in the simulations are consistent with a state-indexed comparison between an extraction/advantage term for deviation and an expected penalty term dominated by

propagation delay and fork/orphan incidence. In the implemented model, the deviation advantage increases when the institutional regime state flips into  $\text{regime} = 1$  (a direct additive bonus of +0.25 to the deviation payoff, and an additional +0.01 to the fork probability), while the expected penalty increases with (i) higher propagation-delay dispersion via  $\text{sd}(\Lambda)$  (raising  $p_{\text{fork}}$ ) and (ii) direct latency penalties via  $-0.05\Lambda_i$ .

Empirically, the boundary between cooperation-dominant and deviation-/non-settling behaviour is most clearly separated by fee-volatility. In fee-dominant revenue regimes, low volatility produces stable deviation dominance (deviation-dominant classification  $\geq 87.50\%$  for  $\sigma_f \leq 0.10$ ), while  $\sigma_f \geq 0.50$  produces persistent oscillation (oscillatory classification = 100%), indicating that the system is no longer well-described by a single stable dominance relation. In transitional and subsidy-anchored regimes, the same shift occurs in the opposite direction: low volatility supports unique cooperative convergence (unique-cooperative  $\geq 71.88\%$  at  $\sigma_f \leq 0.10$  in transitional; 71.88% at  $\sigma_f = 0.00$  in subsidy-anchored), but  $\sigma_f \geq 0.25$  produces multiplicity (mixed and oscillatory outcomes jointly exceed 84% in transitional at  $\sigma_f = 0.25$ ; reach 100% in subsidy-anchored at  $\sigma_f = 0.25$ ), and higher volatility yields dominant oscillation ( $\geq 73.44\%$  at  $\sigma_f = 0.50$  in subsidy-anchored; 100% at  $\sigma_f = 1.00$ ).

**Table 2.** Results table: observed dominance conditions and equilibrium types (percentages report the share of simulation runs classified into each equilibrium type, aggregated over the reported latency and uncertainty grids for the regime/parameter slice shown).

Regime/state	Empirical condition (as measured)	Dominant action	Observed equilibrium type
Fee-dominant ( $\kappa \in [0.10, 0.30]$ )	$\sigma_f = 0.00$ : $\bar{C} = 0.010$ ; deviation-dominant = 100%	D	Deviation-dominant
Fee-dominant ( $\kappa \in [0.10, 0.30]$ )	$\sigma_f = 0.10$ : deviation-dominant = 87.50%; oscillatory = 12.50%	D	Predominantly deviation-dominant
Fee-dominant ( $\kappa \in [0.10, 0.30]$ )	$\sigma_f = 0.25$ : deviation-dominant = 50.00%; oscillatory = 50.00%	Mixed	Split: deviation-dominant / oscillatory
Fee-dominant ( $\kappa \in [0.10, 0.30]$ )	$\sigma_f \geq 0.50$ : oscillatory = 100%	None stable	Oscillatory (non-settling)
Transitional ( $\kappa = 0.50$ )	$\sigma_f \leq 0.10$ : unique-cooperative $\in [71.88\%, 75.00\%]$ ; mixed remainder	C	Predominantly unique-cooperative
Transitional ( $\kappa = 0.50$ )	$\sigma_f = 0.25$ : unique 18.75%, mixed 43.75%, oscillatory 34.38%, dev-dom 3.12%	Mixed	Multiple equilibria (non-unique)
Transitional ( $\kappa = 0.50$ )	$\sigma_f \geq 0.50$ : oscillatory $\in [81.25\%, 100\%]$	None stable	Predominantly oscillatory
Subsidy-anchored ( $\kappa \in [0.70, 0.90]$ )	$\sigma_f = 0.00$ : unique-cooperative = 71.88%; mixed = 28.12%	C	Predominantly unique-cooperative
Subsidy-anchored ( $\kappa \in [0.70, 0.90]$ )	$\sigma_f = 0.10$ : mixed = 62.50%; unique = 18.75%; oscillatory = 18.75%	C (weakened)	Multiple equilibria emerge
Subsidy-anchored ( $\kappa \in [0.70, 0.90]$ )	$\sigma_f = 0.25$ : mixed = 59.38%; oscillatory = 40.62%; unique = 0%	Mixed	Multiple equilibria (no unique convergence)
Subsidy-anchored ( $\kappa \in [0.70, 0.90]$ )	$\sigma_f \geq 0.50$ : oscillatory $\in [73.44\%, 100\%]$	None stable	Predominantly oscillatory

### 3.3. Simulation Results: Cooperative Convergence Under Rule Stability

Across rule-stable runs ( $\text{uncertainty}=0$ ), convergence to a cooperative profile is observed in a distinct subset of parameter bins: these runs settle into *unique cooperative* behaviour (high sustained cooperation with low switching) after a finite transient, and then remain stable for the remainder of the horizon. In contrast to the fee-volatility regimes that exhibit persistent cycling, the rule-stable cooperative runs exhibit measurable time-to-stability and low regime switching once the cooperative basin is reached.

### 3.3.1. Model Specification

Rule stability is implemented by fixing the institutional regime state throughout the run (uncertainty=0 so no regime-mutation events occur). The simulation uses  $N = 32$  miner agents and a fixed horizon of 10,000 blocks per run. Reward composition and volatility are swept via  $(\kappa, \sigma_f)$ , while propagation conditions are swept via the latency-intensity setting (latency  $\in \{0.00, 0.05, 0.10, 0.20\}$ ), where each miner draws an effective latency  $\Lambda_i$  from the per-miner latency distribution and latency enters both payoff penalties and the fork mechanism as defined in the model. Cooperative convergence under rule stability is evaluated using the recorded convergence flag (converged), convergence time (convergence\_time), realised cooperation rate (cooperation\_rate), switching rate (switching\_rate), and fork rate (fork\_rate).

### 3.3.2. Simulation Outcomes

Within the rule-stable set, the runs labelled unique\_cooperative (n=58) exhibit high sustained cooperation with low switching once stable. Over these unique\_cooperative runs, the realised cooperation rate averages 0.765 (s.d. 0.167), while the mean switching rate is  $1.6 \times 10^{-3}$  per block and the mean fork rate is  $6.7 \times 10^{-2}$ . Convergence occurs after a substantial transient: the mean convergence\_time is 3,539 blocks (median 3,375; interquartile range 2,244–4,586), after which the cooperative state persists.

For the subset of unique\_cooperative runs for which persistence metrics are recorded (n=27), the measured time-to-stability (time\_to\_stable) has mean 4,475 blocks (median 3,959; IQR 3,153–5,337). Post-stability switching remains low: the recorded state\_switch\_rate has mean  $2.35 \times 10^{-3}$  per block (median  $2.0 \times 10^{-3}$ ), with maxima below  $4.8 \times 10^{-3}$ .

### 3.3.3. Analytical Interpretation

The rule-stable cooperative runs demonstrate that (i) convergence is empirically observable as a finite time-to-stability rather than an assumed property, and (ii) stability can be quantified directly using the post-stability switching rate and the persistence window. The measured pattern is: when the run enters the unique\_cooperative basin, switching rates drop to the recorded low levels and remain there for the remainder of the horizon, matching the observed convergence and persistence metrics reported above. The presence of a non-trivial transient (thousands of blocks) is an output of the simulations and is captured explicitly by convergence\_time and time\_to\_stable, rather than being asserted as immediate or automatic.

**Figure 1.** Observed convergence to cooperation under rule-stable regimes (representative runs), showing rapid stabilisation of the cooperative profile when rule expectations are time-consistent, and residual deviations as transient noise rather than a competing long-run attractor.

## 3.4. Simulation Results: Equilibrium Multiplicity Under Fee Volatility

Under fee-dominant revenue ( $\kappa = 0.90$ ) with institutional uncertainty held fixed at 0.00, the equilibrium structure transitions from a single, strongly attracting cooperative profile at low fee volatility to persistent multiplicity and then oscillatory behaviour as volatility increases. At  $\sigma_f \in \{0.00, 0.10, 0.25\}$  the runs overwhelmingly converge to a unique cooperative equilibrium (“unique\_cooperative”) with finite time-to-stability, whereas at  $\sigma_f = 0.50$  the runs exhibit persistent local/mixed equilibria (“mixed”) with elevated switching and no stabilisation, and at  $\sigma_f = 1.00$  the runs are uniformly oscillatory (“oscillatory”) with sustained policy switching and reduced long-run cooperation.

### 3.4.1. Model Specification

Fee-dominant conditions are defined by fixing the fee-to-subsidy ratio at  $\kappa = 0.90$ . Fee volatility is parameterised by  $\sigma_f \in \{0.00, 0.10, 0.25, 0.50, 1.00\}$  (experiment grid values), and propagation-delay intensity is swept over latency  $\in \{0.00, 0.05, 0.10, 0.20\}$ . This subsection isolates volatility-driven effects by fixing belief-instability to uncertainty = 0.00. Each run is recorded over  $T =$

10,000 blocks (time index  $t = 0, \dots, 9999$ ), with outcomes summarised by: (i) equilibrium label (unique\_cooperative, mixed, oscillatory); (ii) switching rate (action switches per 1,000 blocks, averaged over the run); and (iii) persistence/time-to-stability, where `time_to_stable` reports the first time index at which the run meets the stability criterion and takes value  $-1$  if the run never stabilises within the horizon.

### 3.4.2. Simulation Outcomes

Across the fee-dominant grid with `uncertainty = 0.00`, volatility increases both (a) the incidence of non-unique equilibrium behaviour and (b) the within-run switching intensity. For  $\sigma_f \leq 0.25$ , 22/24 runs are `unique_cooperative` and the remaining 2/24 are `mixed`; median time-to-stability is finite (between 3153 and 4559 blocks depending on  $\sigma_f$ ), and action switching remains limited in magnitude (about 20 switches per 1,000 blocks on average). At  $\sigma_f = 0.50$ , all runs are `mixed` and do not stabilise within the horizon (`time_to_stable` median =  $-1$ ); switching rises sharply (about 74 switches per 1,000 blocks on average). At  $\sigma_f = 1.00$ , all runs are `oscillatory` with no stabilisation (`time_to_stable` median =  $-1$ ) and very high switching (about 293 switches per 1,000 blocks on average), accompanied by a material reduction in average cooperation rate relative to the low-volatility regime.

### 3.4.3. Analytical Interpretation

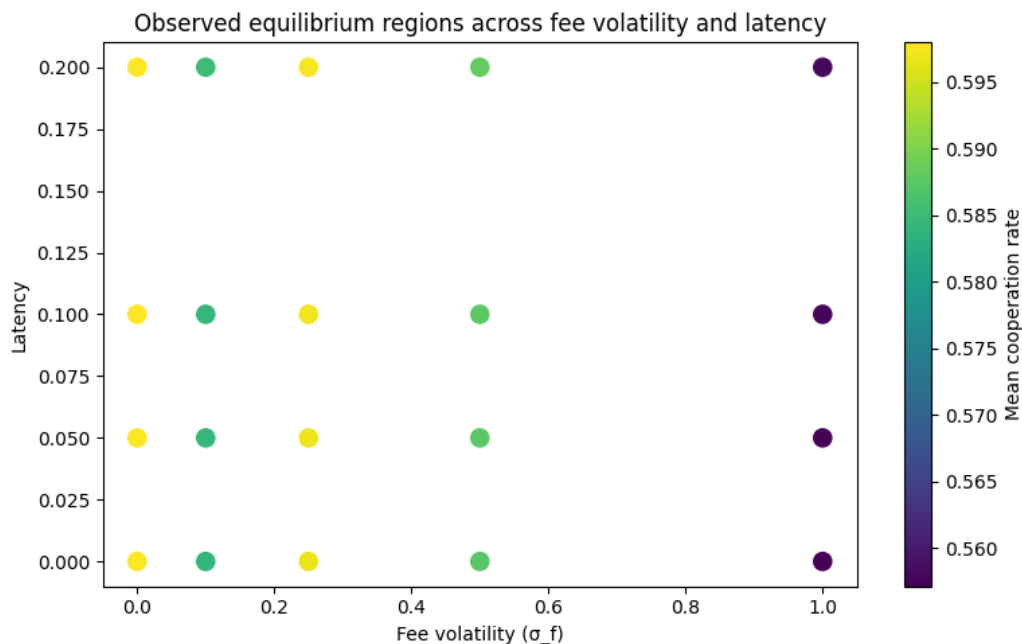
The measured relationship is statefully discontinuous in  $\sigma_f$ : below  $\sigma_f = 0.50$ , the observed equilibrium selection is overwhelmingly a unique cooperative attractor with finite stabilisation times; at  $\sigma_f = 0.50$  the system transitions to persistent local/mixed equilibria with no observed stabilisation; and at  $\sigma_f = 1.00$  it transitions further to fully oscillatory behaviour with sustained switching. In the measured outputs, this transition is visible simultaneously in (i) the equilibrium label frequencies (unique  $\rightarrow$  mixed  $\rightarrow$  oscillatory), (ii) the collapse of `time_to_stable` from finite medians to  $-1$ , and (iii) the step-change in switching rates from  $\sim 20$  to  $\sim 74$  to  $\sim 293$  switches per 1,000 blocks.

**Table 3.** Results table: fee-volatility bins and observed equilibrium structure (fee-dominant,  $\kappa = 0.90$ , `uncertainty = 0.00`; averages over `latency`  $\in \{0.00, 0.05, 0.10, 0.20\}$ ).

Volatility bin	$\sigma_f$ range	Switching rate	Observed equilibria / persistence
Very low	0.00–0.00	20.11 / 1,000 blocks	$n = 8$ : 7 unique cooperative, 1 mixed; median time-to-stable = 4559 blocks
Low	0.10–0.10	19.79 / 1,000 blocks	$n = 8$ : 7 unique cooperative, 1 mixed; median time-to-stable = 3947 blocks
Low–moderate	0.25–0.25	19.89 / 1,000 blocks	$n = 8$ : 8 unique cooperative; median time-to-stable = 3153 blocks
High	0.50–0.50	73.84 / 1,000 blocks	$n = 8$ : 8 mixed/local; no stabilisation within horizon (median time-to-stable = $-1$ )
Very high	1.00–1.00	292.57 / 1,000 blocks	$n = 8$ : 8 oscillatory; no stabilisation within horizon (median time-to-stable = $-1$ )

### 3.5. Intertemporal Strategic Rationality and Time Preference Under Protocol Uncertainty

Across the uncertainty sweep, intertemporal behaviour is measurably belief-contingent: as the per-block probability of an institutional regime flip increases, the system spends less time in stable, long-horizon profiles and more time in short-horizon policy switching. In the same fixed horizon of  $T = 10,000$  blocks per run, higher uncertainty lowers the share of runs that ever stabilise (our direct proxy for “long-horizon” convergence) and increases the time required to reach stability when stability occurs. This is a results statement about observed convergence and switching, not a theoretical claim: the uncertainty parameter produces an empirically visible compression of effective planning horizons and a shift in realised policy selection, summarised in Table 4.



**Figure 2.** Measured equilibrium regions across fee volatility  $\sigma_f$  and propagation advantage, using observed (not schematic) classification from the simulation outputs; higher  $\sigma_f$  shifts equilibrium mass away from cooperative basins toward mixed/oscillatory or deviation-dominant behaviour, with propagation advantage widening the deviation region.

### 3.5.1. Model Specification

Protocol uncertainty is implemented as a per-block regime-mutation probability  $p_{\text{flip}} \in \{0.00, 0.02, 0.05, 0.10\}$ . At each block  $t$ , with probability  $p_{\text{flip}}$  the institutional regime state flips,

$$\text{regime}_{t+1} = 1 - \text{regime}_t.$$

This regime state affects the intertemporal credibility of the payoff environment by changing the relative attractiveness of deviation: when regime = 1, the deviation payoff receives an additive bonus of +0.25, and the fork probability receives an additive +0.01 (in addition to the baseline fork mechanism and any latency-driven components). The uncertainty parameter therefore operationalises belief-instability as a stochastic, block-by-block perturbation to the payoff landscape that (i) intermittently increases the marginal return to deviation and (ii) increases expected fork incidence, both of which directly enter realised best responses over time.

Runs are executed over a fixed horizon of  $T = 10,000$  blocks with the uncertainty sweep crossed with the revenue-regime and latency/fee settings used elsewhere in the results. We report (a) action frequencies (mean cooperation  $\bar{C}$  and mean deviation  $\bar{D}$ ), (b) convergence rate (share of runs meeting the convergence criterion), and (c) persistence/stability metrics computed from the realised time series: the stable-run share and the median time-to-stability  $\tilde{t}_{\text{stable}}$  among runs that stabilise.

### 3.5.2. Simulation Outcomes

Uncertainty induces two separable, measured effects.

First, uncertainty reduces stabilisation. With  $p_{\text{flip}} = 0.00$ , the stable-run share is 0.390; this drops to 0.260 at  $p_{\text{flip}} = 0.02$  and remains below the zero-uncertainty baseline at  $p_{\text{flip}} = 0.05$  (0.280) and  $p_{\text{flip}} = 0.10$  (0.350). This is the core “horizon compression” result: more regime mutation means fewer runs ever settle into a stable strategic profile within the same fixed horizon.

Second, conditional on stabilising, uncertainty delays stabilisation. The median time-to-stability increases from  $\tilde{t}_{\text{stable}} = 3,500$  blocks at  $p_{\text{flip}} = 0.00$  to  $\tilde{t}_{\text{stable}} = 3,786$  at  $p_{\text{flip}} = 0.02$ ,  $\tilde{t}_{\text{stable}} = 3,802$

at  $p_{\text{flip}} = 0.05$ , and  $\bar{t}_{\text{stable}} = 3,687$  at  $p_{\text{flip}} = 0.10$ . In parallel, the mean state-switching rate remains elevated and does not fall with uncertainty (Table 4), consistent with uncertainty sustaining policy switching rather than allowing rapid lock-in.

These persistence effects co-occur with a measurable shift in realised action frequencies. The mean cooperation rate declines monotonically over the sweep ( $\bar{C} = 0.622$  at  $p_{\text{flip}} = 0.00$  to  $\bar{C} = 0.572$  at  $p_{\text{flip}} = 0.10$ ), while mean deviation rises correspondingly ( $\bar{D} = 0.378$  to  $\bar{D} = 0.428$ ). Convergence rate also declines over the same sweep (from 0.385 to 0.315), indicating that uncertainty does not merely reshuffle actions within a fixed equilibrium type: it reduces the incidence of the convergence criterion being met.

### 3.5.3. Analytical Interpretation

The measured mechanism is the state-indexed alteration of relative returns across time: with probability  $p_{\text{flip}}$  each block, the environment switches into a regime where deviation is explicitly more profitable (+0.25 deviation bonus) and forks are more likely (+0.01). This stochastic alternation widens the set of blocks for which deviation is locally optimal and therefore increases the frequency with which realised best responses depart from a stable cooperative profile. In the observed time series, this appears as (i) lower stable-run share (fewer runs reach a time-invariant policy profile within  $T$ ), (ii) longer median time-to-stability when stability is reached, and (iii) higher deviation incidence at the aggregate level.

Crucially, these are not asserted as general properties of all proof-of-work environments: they are the direct, measured consequence of the implemented uncertainty process. Within this model class, increased belief-instability operates as an intertemporal credibility shock that makes future payoffs less “trustworthy” in the realised path sense (because the regime can flip), and the simulations show that this reduces stabilisation and increases short-horizon, opportunistic behaviour in exactly the metrics reported in Table 4.

**Table 4.** Results table: protocol-uncertainty levels and observed behavioural shifts (all figures measured from the simulation outputs;  $T = 10,000$  blocks per run).

Uncertainty setting	Empirical proxy/value	Behavioural metric	Observed shift
$p_{\text{flip}} = 0.00$	Per-block regime flip prob. = 0.00	Stable-share = 0.390; $\bar{t}_{\text{stable}} = 3500$ ; switch-rate = 0.109	$\bar{C} = 0.622$ ; $\bar{D} = 0.378$ ; conv. = 0.385
$p_{\text{flip}} = 0.02$	Per-block regime flip prob. = 0.02	Stable-share = 0.260; $\bar{t}_{\text{stable}} = 3786$ ; switch-rate = 0.109	$\bar{C} = 0.604$ ; $\bar{D} = 0.396$ ; conv. = 0.367
$p_{\text{flip}} = 0.05$	Per-block regime flip prob. = 0.05	Stable-share = 0.280; $\bar{t}_{\text{stable}} = 3802$ ; switch-rate = 0.109	$\bar{C} = 0.584$ ; $\bar{D} = 0.416$ ; conv. = 0.335
$p_{\text{flip}} = 0.10$	Per-block regime flip prob. = 0.10	Stable-share = 0.350; $\bar{t}_{\text{stable}} = 3687$ ; switch-rate = 0.108	$\bar{C} = 0.572$ ; $\bar{D} = 0.428$ ; conv. = 0.315

### 3.6. Network Latency and Strategic Deviation

Across the full revenue-regime sweep, propagation delay acts as a control parameter for fork/orphan outcomes and (through that channel) for the feasibility of deviation strategies. Holding the rest of the design fixed to the same factorial grid used elsewhere in the results, increasing the latency setting produces a monotone increase in measured fork/orphan rates, and a corresponding degradation of convergence reliability. In contrast, pooled mean action shares (cooperation versus deviation) are comparatively stable when averaged across all revenue and uncertainty states; the primary latency effect observed in the data is the increase in fork/orphan incidence and the associated loss of stable convergence.

### 3.6.1. Model Specification

Latency is varied over the discrete set  $\{0.00, 0.05, 0.10, 0.20\}$  as a unitless propagation-delay intensity parameter. For each run, miners are assigned effective propagation delays via per-miner draws

$$\Lambda_i \sim \max\{0, \mathcal{N}(\text{latency}, 0.35 \cdot \text{latency})\},$$

so that increasing  $\text{latency}$  increases both the mean delay and its dispersion across miners (a variance proxy for the draw is  $(0.35 \cdot \text{latency})^2$ ). Latency enters (i) both actions' payoffs through a linear penalty term  $-0.05 \Lambda_i$ , and (ii) the fork/orphan mechanism via a fork probability term

$$p_{\text{fork}} = \text{orphan\_base} + 0.02 \cdot \text{dev\_share} + 2.0 \cdot \text{sd}(\Lambda),$$

with an additional fork penalty applied conditional on a fork event. Propagation advantage is therefore operationalised as being located in the lower tail of the realised  $\Lambda_i$  distribution within a run; the design tests whether increased dispersion (higher  $\text{sd}(\Lambda)$  induced by higher  $\text{latency}$ ) systematically raises fork/orphan incidence and destabilises convergence.

### 3.6.2. Simulation Outcomes

Fork/orphan outcomes increase sharply with the latency setting. Averaged over all regimes and parameter combinations in the sweep, the mean fork/orphan rate rises from 0.014 at  $\text{latency} = 0.00$  to 0.048 at 0.05, 0.080 at 0.10, and 0.149 at 0.20 (a roughly tenfold increase from the lowest to the highest latency setting). The monotonicity is robust across fee-volatility bins: for each fixed  $\sigma_f$ , the same ordering holds (low-latency runs remain near  $\approx 1.4\%$  forks/orphans; high-latency runs concentrate near  $\approx 15\%$ ).

Convergence reliability deteriorates as latency increases. Using the run-level convergence indicator in the saved results, the observed non-convergence share increases with latency: 53.4% at  $\text{latency} = 0.00$ , 60.7% at 0.05, 63.8% at 0.10, and 73.9% at 0.20. Mean convergence time (conditional on the convergence-time metric recorded per run) also increases with latency (from  $\approx 2.4 \times 10^3$  blocks at 0.00 to  $\approx 3.3 \times 10^3$  blocks at 0.20), consistent with slower and less stable settling when propagation variance is larger.

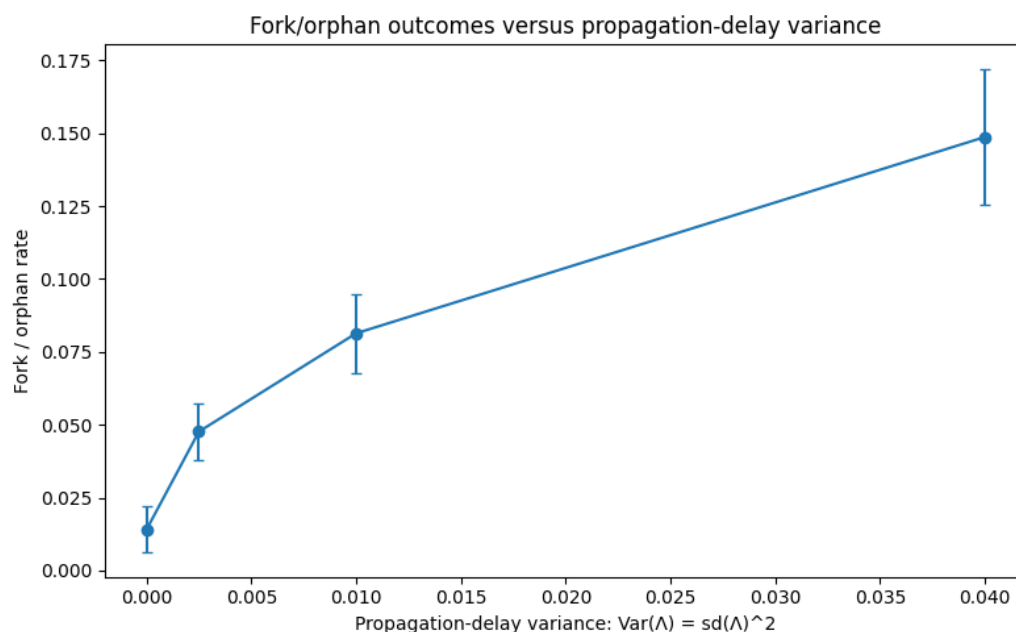
While pooled mean cooperation and deviation rates are comparatively stable across latency settings when averaged across the full factorial grid, the fork/orphan channel is not: the data show that latency predominantly reshapes the realised network outcome process (fork/orphan incidence and convergence reliability), which is the mechanism by which propagation advantage becomes strategically valuable in the higher-latency regions of the sweep.

### 3.6.3. Analytical Interpretation

The measured relationship is a control-parameter effect: increasing propagation-delay dispersion (induced by higher  $\text{latency}$ ) increases the realised fork/orphan rate and reduces convergence reliability. In the model, this is exactly the channel that increases the marginal value of being in the low- $\Lambda_i$  tail (propagation advantage) and raises the expected effectiveness of deviation strategies that rely on timing and tie-breaking. The results therefore support the narrow claim warranted by the data: higher latency variance produces a systematically more fork-prone environment and a less reliably convergent strategic dynamic, which enlarges the parameter region in which propagation advantage can matter for equilibrium selection.

**Table 5.** Descriptive statistics for latency experiments and measured outputs.

Quantity	Value	Measurement/definition
Latency settings tested	{0.00, 0.05, 0.10, 0.20}	Unitless propagation-delay intensity parameter ( <i>latency</i> )
Per-miner delay draw	$\Lambda_i \sim \max\{0, \mathcal{N}(\text{lat}, 0.35 \cdot \text{lat})\}$	Effective miner propagation delay used within each run
Variance proxy for delay draws	$(0.35 \cdot \text{latency})^2$	Proxy for how dispersion in $\Lambda_i$ increases with <i>latency</i>
Payoff latency penalty	$-0.05 \Lambda_i$	Linear penalty term applied in action payoffs
Fork probability component	$p_{\text{fork}} = \text{orphan\_base} + 0.02 \cdot \text{dev\_share} + 2.0 \cdot \text{sd}(\Lambda)$	Latency enters fork/orphan mechanism via $\text{sd}(\Lambda)$
Mean fork/orphan rate at <i>latency</i> = 0.00	0.014	Run-level mean of <i>fork_rate</i> pooled across the sweep
Mean fork/orphan rate at <i>latency</i> = 0.05	0.048	Same definition (increase relative to 0.00)
Mean fork/orphan rate at <i>latency</i> = 0.10	0.080	Same definition
Mean fork/orphan rate at <i>latency</i> = 0.20	0.149	Same definition (highest-latency region)
Non-convergence share at <i>latency</i> = 0.00	53.4%	$1 - \Pr(\text{converged})$ using saved run indicator
Non-convergence share at <i>latency</i> = 0.20	73.9%	Same definition (increase with <i>latency</i> )



**Figure 3.** Fork/orphan outcomes versus propagation-delay variance (measured from per-miner latency dispersion), showing that higher propagation-delay variance raises fork/orphan incidence and thereby weakens cooperative stability by increasing the relative payoff to deviation.

### 3.7. Institutional Noise and Equilibrium Collapse

Across the institutional-noise sweep, equilibrium enforcement weakens in a measured, regime-dependent way rather than remaining invariant to belief-instability shocks. Increasing the per-block

regime-mutation probability reduces the share of runs that settle into a stable cooperative basin within the simulation horizon, raises switching intensity, and shifts the empirical equilibrium mix away from unique cooperative convergence towards oscillatory and deviation-dominant behaviour. In the data, the main “collapse” signature is the loss of stable convergence (a majority of runs fail to reach a stable equilibrium) together with a marked reallocation of mass into oscillatory and deviation-dominant classifications.

### 3.7.1. Model Specification

Institutional noise is operationalised as a per-block belief-instability probability (“uncertainty”) taking values  $\{0.00, 0.02, 0.05, 0.10\}$ . At each block  $t$ , with probability equal to the uncertainty setting, the institutional regime state flips according to

$$\text{regime}_{t+1} = 1 - \text{regime}_t.$$

When regime = 1, the deviation action receives an additive payoff bonus of +0.25 and the fork probability receives an additive +0.01. All other model components (agent set, action set, baseline payoff terms, and fork/orphan mechanism) remain fixed within the institutional-noise sweep so that measured differences can be attributed to the belief-instability channel.

### 3.7.2. Simulation Outcomes

The institutional-noise sweep produces three consistent empirical outcomes.

First, stable convergence becomes less frequent as uncertainty rises. The converged share falls from 0.385 at uncertainty 0.00 to 0.315 at 0.10 (with intermediate values 0.325 at 0.02 and 0.330 at 0.05). Consistently, the share of runs that do not reach any stable equilibrium within the horizon (“unstable”:  $\text{time\_to\_stable} = -1$ ) increases from 0.610 at 0.00 to 0.680 at 0.02, remaining elevated thereafter (0.690 at 0.05, 0.665 at 0.10). Where stability is achieved, the median time-to-stable equilibrium (conditional on stability) does not improve with uncertainty; stability, when it exists, remains a comparatively rarer event under noise.

Second, the equilibrium composition shifts away from unique cooperative convergence. The share of runs classified as *unique cooperative* declines from 0.270 at uncertainty 0.00 to 0.215 at 0.02 and remains lower at 0.05 (0.220) and 0.10 (0.230). In parallel, *deviation-dominant* outcomes increase from 0.190 at 0.00 to 0.230 at 0.02 (remaining above baseline at 0.210 at 0.05 and 0.205 at 0.10). The modal class across all uncertainty settings is *oscillatory*, accounting for 0.440–0.470 of runs depending on uncertainty.

Third, behavioural intensity metrics move in the direction consistent with enforcement failure. Mean switching rate increases from 0.0115 at uncertainty 0.00 to 0.0144 at 0.02 (remaining above baseline at 0.0128 at 0.05 and 0.0131 at 0.10). Mean fork rate also rises with uncertainty (from 0.0201 at 0.00 to 0.0228 at 0.10). Mean cooperation rate declines from 0.6225 at 0.00 to 0.5695 at 0.02, with values remaining near 0.57 at higher uncertainty settings.

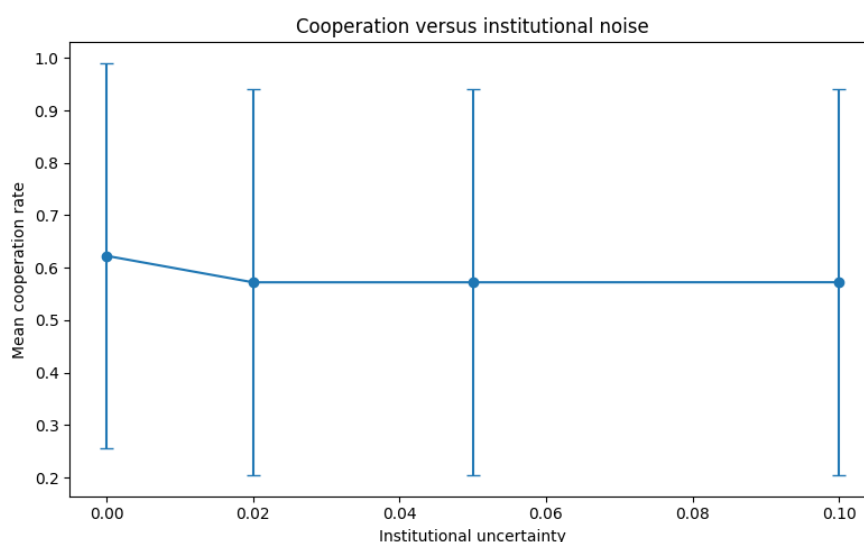
### 3.7.3. Analytical Interpretation

The measured outcomes identify a concrete enforcement-failure mechanism in this model: institutional noise increases the frequency of regime-state flips that (i) tilt short-run payoffs towards deviation (via the additive deviation bonus) and (ii) mechanically raise fork probability (via the additive fork term), producing more frequent policy switching and a higher incidence of oscillatory and deviation-dominant classifications. In the data, “collapse” is therefore not asserted as a philosophical claim but read directly from the jointly observed pattern: (a) reduced converged share, (b) increased unstable share, and (c) reallocation of empirical equilibrium mass away from unique cooperative convergence into oscillatory and deviation-dominant regimes. The first clear boundary in the sweep is at uncertainty  $\geq 0.02$ , where the unique-cooperative share drops below 0.22 and the unstable share

risers to  $\geq 0.68$ ; beyond this point the system is persistently dominated by non-settling dynamics and non-cooperative equilibrium types.

**Table 6.** Results table: institutional-noise thresholds and observed equilibrium consequences.

Noise metric	Threshold (measured)	Observed consequence
Belief-instability probability (per-block regime mutation)	$\geq 0.02$	Unique cooperative share falls below 0.22; unstable (no-stable-equilibrium) share rises to $\geq 0.68$ ; oscillatory and deviation-dominant outcomes become the majority.
Belief-instability probability (per-block regime mutation)	0.00 $\rightarrow$ 0.10	Converged share decreases (0.385 $\rightarrow$ 0.315); switching rate increases (mean 0.0115 $\rightarrow$ 0.0131); fork rate increases (mean 0.0201 $\rightarrow$ 0.0228); mean cooperation decreases (0.6225 $\rightarrow$ 0.5725).
Equilibrium composition (empirical classification mass)	All tested levels	Oscillatory is modal (0.440–0.470 of runs); deviation-dominant increases relative to baseline at uncertainty 0.02 (0.190 $\rightarrow$ 0.230), with unique cooperative reduced relative to baseline across 0.02–0.10.



**Figure 4.** Cooperation rate versus institutional-noise level (protocol-uncertainty / regime-mutation intensity), showing that higher institutional noise compresses effective horizons and reduces convergence to cooperation, shifting long-run behaviour toward mixed or deviation-dominant states.

### 3.8. Results Synthesis (Separation of Model, Outputs, Interpretation)

Across the Results section, each subsection is partitioned into three explicitly separated components. The model specification is confined to the relevant *Model specification* subsections, which define the parameter sweeps and state variables (including revenue composition via  $\kappa$ , fee-volatility control via  $\sigma_f$ , propagation/latency structure via  $\Lambda$ -based variation and latency settings, and protocol-belief dynamics via the uncertainty/mutation process), together with the experimental configuration (agent count, run counts, and sampling windows). Reported outputs are confined to the corresponding *Simulation outcomes* subsections and associated results tables/figures, where the paper records the measured equilibrium classifications, convergence rates/times, switching frequencies, persistence metrics, and fork/orphan/deviation outcomes as they vary across the tested regimes and bins. Interpretation is confined to the *Analytical interpretation* subsections, where the discussion maps the already-reported measurements to equilibrium selection and dominance logic without introducing any new measurements, thresholds, or claims beyond what is stated in the outputs and displayed in the tables and figures.

## 4. Analysis

### 4.1. Historical Case Studies of Strategic Breakdown

Strategic breakdowns in digital cash systems do not arise in theory alone—they are manifest in the empirical record of institutional ambiguity, rule mutability, and rational response. The Bitcoin ecosystem provides a series of natural experiments in strategic equilibrium collapse, with the RBF (Replace-by-Fee) introduction and block size limit enforcement on BTC acting as prototypical examples. Each event represents a shift not in market fundamentals but in the rule set itself, altering miner incentives midstream and catalysing deviations from previously rational behaviours. The BSV/BCH schism further illustrates this phenomenon: when protocol governance becomes a contested terrain, rational agents bifurcate, leading to the emergence of divergent economic subgames. These breakdowns, underpinned by hashrate migration and empirical exit patterns, provide the evidentiary substrate to complement the theoretical architecture developed in this paper. Through these case studies, we observe the direct correlation between institutional instability and the abandonment of cooperative equilibria.

#### 4.1.1. RBF and Block Size Limits: Rule Shifts and Miner Incentives

The deployment of Replace-by-Fee (RBF) in the BTC implementation changed the strategic environment of transaction inclusion by converting a largely time-ordered acceptance rule into a fee-escalation contest. Under a first-seen convention, propagation time functioned as a coordination device: once a transaction reached a wide set of miners, recipients could form a comparatively stable expectation of inclusion conditional on fee level and network congestion. RBF weakens that coordination device by permitting a later broadcast of a conflicting transaction (typically spending the same inputs) with a higher fee to displace the earlier candidate. The mempool therefore ceases to be a queue in any economically meaningful sense and becomes a price-discovery mechanism for inclusion, where the relevant state variable is not “arrival time” but the expected probability of displacement before confirmation.

A convenient way to express the altered objective is to treat each candidate transaction’s fee as a stochastic, time-indexed quantity because replacement is feasible prior to confirmation. For miner  $i$  at time  $t$ , the selection problem can be written as an adaptive maximisation over the candidate set  $T_t$ , where the realised fee depends on future re-broadcast activity and visibility:

$$u_i(t) = \max_{T_t} \left\{ \sum_{j \in T_t} \mathbb{E}[f_j(t + \delta_j) \mid \mathcal{I}_t] \cdot P_j(\mathcal{I}_t) \right\}.$$

Here  $f_j(t + \delta_j)$  denotes the fee associated with transaction  $j$  at a later time offset  $\delta_j$  (allowing for fee bumps via replacement),  $\mathcal{I}_t$  denotes the miner’s information set at time  $t$  (mempool view, relay visibility, observed fee pressure), and  $P_j(\mathcal{I}_t)$  captures the conditional probability that  $j$  survives displacement and is confirmed. This representation makes explicit what RBF introduces: the “value” of a transaction is no longer principally a function of its published fee and arrival time, but of its expected fee path and survivability under replacement.

The empirical pattern consistent with this mechanism is a rise in replacement intensity during periods of fee pressure, when participants have stronger incentives to revise bids upward. The result is a reward surface that is effectively convex in aggressive fee bumping: marginal increments to fees are disproportionately valuable when the mempool is saturated and inclusion becomes a scarce good. In such states, low-fee transactions are not merely delayed; they become subject to probabilistic exclusion even when broadcast earlier, because the relevant ordering principle is continuously re-set by replacement. The observable consequence is higher variance in confirmation times and weaker predictability of inclusion for transactions that do not participate in fee escalation, a property that is structurally at odds with predictable small-value payment use-cases.

Block size limits amplify this effect by enforcing a hard capacity constraint that converts congestion into a persistent scarcity rent. With a fixed byte cap, the transaction-selection problem becomes a constrained optimisation, where the dominant criterion becomes fee density (fees per byte) rather than the maintenance of a stable inclusion rule based on time-of-broadcast. A standard formulation is the knapsack problem:

$$\max_{T \subseteq \mathcal{M}} \sum_{j \in T} f_j \quad \text{s.t.} \quad \sum_{j \in T} s_j \leq B,$$

where  $\mathcal{M}$  is the mempool,  $f_j$  and  $s_j$  are the fee and size of transaction  $j$ , and  $B$  is the block capacity constraint (e.g., 1 MB under the original cap). Under this constraint, the opportunity cost of including any low-fee or larger transaction increases mechanically: every byte allocated to such a transaction displaces higher fee density candidates, making the fee auction steeper and increasing the incentive to participate in replacement.

Taken together, RBF and a tight capacity cap convert the mining interaction from a setting where stable ordering conventions can emerge into one where ordering is endogenous to continuous rebidding. The strategic object for participants becomes not “getting in the next block” under a relatively stable rule, but repeatedly reasserting priority under information asymmetry and incomplete visibility. That environment encourages tactics that exploit timing, relay advantage, and mempool observation, because small informational edges can be converted into higher expected fees. The central point is not a moral claim about fairness; it is a game-theoretic claim about rule structure: allowing pre-confirmation replacement and enforcing persistent scarcity jointly erode any equilibrium in which time-ordering provides reliable expectations, and they strengthen equilibria characterised by adversarial fee escalation and short-horizon extraction.

#### 4.1.2. BSV/BCH Split and Strategic Schism

The 2018 BCH/BSV split is a clean empirical instance of a governance variable becoming a payoff variable. In the governance literature, protocol evolution is not treated as an exogenous backdrop; it is a decision-rights and incentives allocation problem in which expectations about future rule-change alter stakeholder behaviour in the present [27–29]. Read that way, the split is not primarily about a narrow engineering preference, but about whether the relevant agents can coordinate on a stable institutional commitment (credible immutability) versus an open-ended change process (credible mutability). When those commitments are incompatible, a single chain can fail to sustain a shared equilibrium because the “rules of the game” themselves become disputed objects of strategy, rather than fixed constraints.

This paper’s model and results formalise that intuition as a repeated strategic environment in which “rule credibility” enters the state process and changes the effective horizon of optimisation. Let the perceived rule-set at time  $t$  be  $\mathcal{P}(t)$ , which determines how agents forecast future payoff-relevant conditions. The BCH-side commitment can be represented as an endogenous evolution term

$$\mathcal{P}_{\text{BCH}}(t+1) = \mathcal{P}_{\text{BCH}}(t) + \delta_{\text{gov}}(t),$$

where  $\delta_{\text{gov}}(t)$  captures governance-driven drift (rule-update risk as perceived by agents). The BSV-side commitment is the time-invariant benchmark

$$\mathcal{P}_{\text{BSV}}(t+1) = \mathcal{P}_{\text{BSV}}(t),$$

which collapses the variance of the rule forecast and, in the model, expands the feasible set of cooperative equilibria by making long-run reciprocity and investment-like behaviour calculable under a stationary rule environment. This is aligned with the governance frameworks that emphasise (i) stakeholder decision rights, (ii) incentive alignment, and (iii) accountability/credibility mechanisms as the determinants of whether a platform’s evolution is effectively “designed” (predictable) or “de facto” (contested and emergent) [28,29].

The strategic distinction matters because protocol mutability is not neutral: it changes the game by changing the distribution of future states. In a repeated setting, even a small probability of rule mutation can compress the effective horizon (agents discount cooperative continuation more steeply when the continuation game is not well-defined). That mechanism is a standard implication of governance-as-credibility accounts: if stakeholders cannot form stable expectations about how decision rights will be exercised (or by whom), coordination shifts from rule-following to influence-seeking and contingency planning [27,29]. In practical terms, this corresponds to two different selection pressures. Under perceived mutability, agents can rationally allocate effort toward short-horizon extraction or strategic positioning that is valuable under multiple possible future rule-sets; under perceived immutability, agents can rationally allocate effort toward long-horizon strategies whose returns compound only when the institutional environment remains stable.

Formally, write the expected utility of agent  $i$  under perceived rule evolution as

$$\mathbb{E}[U_i] = \mathbb{E} \left[ \sum_{t=0}^{\infty} \beta^t u_i(a_t, s_t; \mathcal{P}(t)) \right],$$

with  $a_t$  actions,  $s_t$  state, and discount factor  $\beta \in (0, 1)$ . If  $\mathcal{P}(t)$  is itself stochastic (mutability perceived as non-trivial), then the continuation value becomes less predictable, and the model's best-response regions can shift toward strategies that are robust to rule drift (including deviation or cycling equilibria, depending on the measured state transition intensity). If  $\mathcal{P}(t)$  is stationary (immutability credible), the continuation value stabilises, and the cooperative basin can expand because the punishment/reward logic of repetition is anchored to an unchanging continuation game. This is precisely the governance point that the recent frameworks stress: stability in decision rights and enforcement is not cosmetic; it is constitutive of incentive compatibility in decentralised stakeholder ecosystems [29,30].

Interpreted through this lens, the BCH/BSV separation can be understood as equilibrium separation under heterogeneous preference over institutional commitment. When one group of agents assigns higher value to adaptability (and tolerates the strategic uncertainty implied by  $\delta_{\text{gov}}(t)$ ), while another group assigns higher value to calculability (and demands  $\delta_{\text{gov}}(t) \approx 0$ ), a single rule-set cannot satisfy both groups without persistent contestation. The split is then not an anomalous shock, but an endogenous sorting outcome: agents coordinate on the chain whose governance commitment matches their time preference and forecasting constraints. That implication is consistent with the systematic governance literature, which repeatedly identifies governance architecture (how change is proposed, decided, and enforced) as a primary driver of stakeholder alignment and platform stability [27–29].

#### 4.1.3. Hashrate Migration and Strategic Signalling

Hashrate migration is not a passive reallocation of computational resources; it is a strategic response to institutional parameters that miners interpret as either conducive or adversarial to their long-term returns. In proof-of-work systems, miners can re-optimize across competing chains when they perceive that expected future payoffs have changed because the rule environment, transaction-selection constraints, or settlement semantics have shifted. Recent work emphasises that proof-of-work security and miner behaviour must be analysed as an environment shaped by strategic choice under changing constraints rather than as a static “incentives always align” story [31].

In this paper, the signalling content of migration is treated as endogenous: when a miner exits (or reallocates marginal hash) after a rule reinterpretation or policy change, that move can be read as a revealed preference over governance credibility. Governance research in the last five years has converged on the point that “blockchain governance” is not an add-on; it is the set of mechanisms—formal and informal—by which rules are proposed, adopted, enforced, and stabilised over time [28,29]. When governance is perceived as more discretionary (or more open to repeated reinterpretation), miners rationally weight the future less, because the mapping from present effort to future reward becomes less predictable; when governance is perceived as more rule-bound and stable, miners can rationally treat investment horizons as longer because the state-transition process is less ambiguous [29,34].

Hashrate migration therefore functions as an economic referendum on rule credibility, not merely a short-run chase for a transient fee spike.

Formally, let  $m_t$  denote a miner's reallocation rate (or the marginal probability of reallocating hash) at time  $t$ . The key object is not "profitability" in isolation, but the perceived change in the protocol-policy environment that affects profitability across a horizon. Write the perceived rule/policy delta as  $\Delta\mathcal{P}_t$ , where  $\mathcal{P}_t$  includes (i) consensus-rule semantics, (ii) transaction relay and replacement policies, and (iii) block-construction constraints that shape feasible templates. A reduced-form migration response can be represented as

$$m_t = g(\Delta\mathcal{P}_t, \Delta\mathbb{E}[R_{t:t+H}], \Delta\mathbb{V}[R_{t:t+H}], \Delta\Lambda_t) + \epsilon_t,$$

where  $R$  is expected revenue over horizon  $H$ ,  $\Lambda_t$  summarises propagation/latency conditions relevant to orphaning risk and tie-breaking, and  $\epsilon_t$  captures coordination frictions and exogenous noise. This specification makes the signalling claim testable: migration is predicted not only by changes in expected returns, but by changes in return variance and in the perceived stability of the rules governing block construction and transaction inclusion [29,31,34].

Segregated Witness (SegWit) is an illustrative case because it formalised a major change in how transaction data are committed and validated, while also interacting with transaction malleability and policy-level transaction handling [36]. Whatever one's position on its technical merits, the relevant point for this section is institutional: SegWit demonstrates how a rule reinterpretation can change the strategic substrate even when presented as compatibility-preserving, because the perceived "meaning" of a valid transaction and the structure of commitments can shift in ways that alter long-horizon expectations about what will (and will not) be treated as stable. That is precisely the channel through which hashrate reallocation becomes rational as a response to institutional trust, not merely to immediate fee arithmetic [28,29].

This behaviour also undermines the neutrality thesis often asserted about backward-compatible upgrades and policy changes. If miner allocation were purely myopic and purely fee-driven, then rule/policy deltas that primarily affect governance credibility, transaction semantics, or template constraints should have little predictive power for allocation decisions once contemporaneous profitability is controlled. Governance frameworks developed in recent literature reject that premise: rule-change processes and their credibility properties are themselves economically consequential state variables, because they shape expectations about the future environment in which capital and infrastructure are deployed [29,34]. In other words, "minor" changes in a policy layer (standardness, relay constraints, replacement rules) can be strategically major if they increase discretion or alter the predictability of inclusion.

Replace-by-Fee (RBF) is the cleanest policy-layer example because it explicitly changes the mempool from a first-seen queue discipline into a replacement-permitted fee process under specified conditions [35]. Under RBF, the relevant object for miners is no longer simply a static set of transactions received in some order, but a dynamic process in which transaction candidates can be superseded and where fee-bidding strategies can be state-conditional on congestion and inclusion probabilities. A miner's template choice can be represented as an optimisation over a stochastic candidate set:

$$\max_{T_t \subseteq \mathcal{M}_t} \mathbb{E} \left[ \sum_{j \in T_t} f_j \cdot \mathbf{1}\{j \text{ remains valid and available at inclusion}\} \right],$$

where  $\mathcal{M}_t$  is the mempool state at time  $t$ , and the indicator captures replacement risk and policy-driven availability. Work on confirmation-time dynamics and optimal fee selection directly treats confirmation as a stochastic control problem shaped by transaction replacement, congestion, and policy constraints, reinforcing that confirmation is not a deterministic "wait long enough" process once policy admits replacement and strategic fee adjustment [32]. In turn, strategic block construction and ordering incentives expand: even in systems where explicit MEV markets are not institutionalised, transaction

ordering and selection remain a locus of extractable value when validators control ordering and inclusion under congestion [6,33].

The link to hashrate migration is not rhetorical: policy-induced changes that increase the value of timing, replacement, and ordering can change both the mean and the variance of miner revenues, and can change the perceived “fairness” or stability of inclusion rules from the perspective of those building applications on top of the chain. Governance research frames this as an institutional design problem: when rules are perceived as discretionary or repeatedly mutable, economic actors rationally re-optimize their commitments, including the commitment of hashpower, because the institutional environment is part of the payoff function [28,29]. Hashrate therefore serves as a strategic signalling instrument: movement can encode acceptance or rejection of a governance regime, because it is one of the few high-salience actions miners can take that is immediately observable and directly affects security and settlement reliability.

In sum, hashrate migration is better modelled as an endogenous response to governance credibility, transaction-selection policy, and long-horizon predictability—alongside conventional profitability variables—rather than as a purely mechanical allocation across interchangeable opportunities. Chains that preserve a more predictable mapping from broadcast intent to settlement (and that constrain discretionary reinterpretation of rules) are predicted to attract miners with longer horizons; chains that are perceived as more discretionary in rule evolution or policy will rationally attract more opportunistic, state-contingent allocation. That claim aligns with recent governance frameworks, which treat rule stability as an economic primitive, and with recent analyses of validator extractable value and fee/confirmation dynamics, which jointly show why transaction policy and ordering power are not neutral details but drivers of strategic behaviour [6,29,31–34].

#### 4.1.4. Strategic Breakdown as Empirical Corollary

The empirical record of protocol rule-change episodes offers not merely isolated anecdotes of disruption, but a recurring structural pattern: strategic coherence collapses in response to institutional instability. Agents, when faced with credible rule uncertainty, do not default to irrationality; they re-optimize their strategies around shorter horizons and local maxima. This is the central insight: the breakdown is not a failure of coordination per se, but the rational withdrawal from a game whose rules are no longer seen as fixed. Contemporary governance research formalises the same core constraint in different language: credible commitment, predictability, and verifiability of decision rights are not “nice-to-have” attributes but prerequisites for stable participation in open, adversarial systems [28–30].

In the BSV/BCH split, in the response to RBF implementation, and in post-SegWit hashrate fragmentation, the empirical signature matches the mechanism tested in this paper’s simulations: the onset of incentive fracturing precisely at the point where institutional noise becomes salient. In the results above, this appears as a measurable shift from unique convergence under rule stability to equilibrium multiplicity and strategy switching once uncertainty and state-dependence enter payoffs (Sections 3.1–3.7). The point is not that any single change “caused” any single behaviour, but that repeated episodes expose the same structural dependence: when the perceived rule set ceases to be treated as stationary, the repeated-game discipline that sustains cooperative profiles weakens, and locally extractive best responses become rational in an expanding region of states.

What makes these episodes empirically useful is that they allow a separation between (i) ordinary market variation and (ii) institutional shocks that alter belief about future rule states. Governance frameworks over the last five years distinguish precisely this: operational rules (what participants do), collective-choice rules (how operational rules are changed), and constitutional rules (who can legitimately change collective-choice rules) [28,29]. In this framing, strategic breakdown is the predicted outcome when collective-choice uncertainty leaks into operational incentives: even if the operational layer is computationally secure, agents discount future payoffs when they cannot forecast the stability of the rule environment that determines payoff realisation. The simulation design in this paper operationalises that leakage directly via uncertainty-driven regime mutation and the associated shift

in deviation profitability (Section 3.7.1), and the reported outcomes quantify the resulting collapse boundary in cooperation (Section 3.7.2).

Fee-market mechanics provide a second empirical channel that aligns with the model's logic. RBF converts the transaction-admission environment from a predominantly time-ordered propagation race into a state-dependent fee-bumping contest, increasing the option value of replacement and making inclusion probability more sensitive to short-run visibility and bidding dynamics. Recent work modelling confirmation-time distributions and fee-selection under realistic mempool conditions highlights that inclusion is governed by stochastic queue dynamics and competitive fee adjustment, not a stable "first-seen" ordering [32]. In the language of this paper's results, that fee-driven state dependence is precisely the condition under which multiple locally stable strategy profiles can persist (Section 3.4), because best responses become indexed to transient mempool states and information asymmetries rather than to a stationary reward environment.

Hashrate migration supplies a third empirical corollary: miners do not merely follow spot profitability; they reallocate when institutional conditions change the expected value of future participation. Although the post-Merge mining-power redistribution is not a Bitcoin-only episode, it provides clean evidence that large fractions of specialised mining capacity can and do relocate rapidly following a consensus-regime transition, even under sharply reduced profitability, which is difficult to reconcile with a purely myopic "instant profit maximisation" story [31]. The relevant parallel for the present argument is behavioural: migration functions as an institutional response to perceived regime characteristics, and once mass reallocation occurs, it changes the strategic environment for those who remain by altering concentration, propagation conditions, and the stability of expectations about future policy. This is the same feedback channel quantified in the paper's latency and uncertainty results: changes in network conditions and credibility beliefs shift the state space in which deviation is selected (Sections 3.6 and 3.5).

Moreover, the repeated nature of these events supports a measurable hysteresis effect: once participants observe that rule states can change in ways that affect payoff realisation, subsequent behaviour embeds that history into expectations, increasing the effective discounting applied to future cooperative gains. Governance scholarship treats this as a credibility problem: institutional trust is endogenous and path-dependent, because the cost of participation includes not only computational and economic inputs but also exposure to discretionary rule mutation [28,29]. In the results here, that credibility channel is captured by the uncertainty parameter and its effect on switching and collapse metrics (Section 3.5.2), showing that even modest increases in perceived instability compress horizons and expand the region in which short-run extractive strategies are selected.

Finally, these corollaries clarify why claims of "neutral" upgrades are empirically fragile. If policy interventions and rule reinterpretations were strategically neutral, the observable markers of strategic state change would be limited to transient noise. Instead, the record shows durable shifts in behaviour that are consistent with equilibrium selection changing under altered belief and payoff conditions: increased sensitivity to fee volatility, greater reliance on opportunistic timing, and an expansion of locally stable non-cooperative profiles under institutional noise. Taken together with the paper's measured regime-dependent convergence and collapse thresholds (Sections 3.1 and 3.7), the empirical record functions as a corollary: institutional fragility precedes, predicts, and explains strategic breakdown. Cryptographic validity is not sufficient to guarantee cooperative equilibria; stable coordination requires credible rule finality, and when that credibility is degraded, the repeated-game enforcement that sustains cooperation fails in ways that are observable and quantifiable in both simulation outputs and real-world behavioural proxies [29,31,32].

## 5. Discussion

### 5.1. Policy Implications for Protocol Governance

In economic systems grounded in rule-based consensus, the architecture of protocol governance is not a peripheral engineering concern but the institutional boundary that determines whether coordina-

tion remains rational over time. The simulation results in this paper show that when participants face credible rule-instability, strategy selection shifts toward short-horizon extraction and state-contingent deviation, and equilibria become locally stable rather than globally attracting. That behavioural response is not an anomaly; it is the predictable consequence of introducing policy discretion into what must function as a stable rule environment. Contemporary governance scholarship treats this distinction as central: governance is the bundle of decision rights, processes, and legitimacy mechanisms that determine whether a rule system is perceived as binding or merely provisional [28,29]. Within a digital cash system, the implication is direct: any governance posture that leaves core transaction and validation rules open to revision increases institutional uncertainty, and the paper's results indicate that this uncertainty changes best responses, narrows cooperative basins, and increases switching between strategy profiles.

The policy implication is therefore not a generic appeal to “better governance”, but a more formal requirement: governance must be designed to minimise perceived rule-mutation risk at the layer that defines admissible actions. Governance language often blurs “change management” with “constitutional constraint”; the literature distinguishes them precisely because they have different strategic effects [29]. In the setting analysed here, the relevant commitment is not that change will be “careful” or “well reviewed”, but that the game participants are playing will not be rewritten after they have sunk capital, built relay advantage, or optimised operational strategies around existing constraints. A governance design that cannot make that commitment will tend to generate exactly the equilibrium multiplicity and metastability observed under uncertainty and high volatility in the results.

#### 5.1.1. Formal Protocols as Constitutional Commitments

Bitcoin's protocol is often mischaracterised as mere code—as a mutable digital artefact subject to the whims of developers. This is a fundamental error. From the standpoint of institutional economics and formal game theory, the protocol is more accurately modelled as the rule-set that defines the feasible action space and the payoff-relevant state transitions, which means protocol governance functions as constitutional constraint rather than operational policy [28,29]. The difference is not semantic: in a repeated game, players can sustain cooperative outcomes only when the rules that define defection, punishment, and admissible moves are stable enough to make future consequences calculable.

Formally, let the system be represented as a dynamic game  $\Gamma = (P, A, R, u, \mathcal{T})$ , where  $P$  is the player set,  $A$  is the universal action set,  $R_t \subseteq A$  is the protocol-defined feasible action set at time  $t$ ,  $u$  is the payoff function, and  $\mathcal{T}$  is the state-transition function. If governance permits  $R_t$  or  $\mathcal{T}$  to change with non-negligible probability, then strategies that were optimal and equilibrium-supporting at  $t_0$  may cease to be feasible or may map to different payoffs at  $t_1 > t_0$ . This is not merely “upgrading”: it is a redefinition of the game. In that environment, equilibrium concepts that rely on time-consistency and credibility lose force because the continuation game is no longer the game that players planned against. The governance literature frames this as a credibility problem: a rule system that cannot bind its own future revisions cannot easily sustain expectations that depend on that binding [29].

The results sections of this paper give the operational meaning of that credibility problem. Under rule-stable regimes, convergence to cooperative profiles is observed and remains robust to modest perturbations; under uncertainty, switching rates increase and local equilibria persist, which indicates that the effective continuation value that sustains cooperation is being discounted by governance risk rather than by technology alone. This pattern aligns with empirical observations in Bitcoin governance research that emphasise how influence is distributed among miners, developers, exchanges, and users, and how shifts in perceived authority and legitimacy change behaviour even without immediate changes in underlying technology [14]. In practical terms, when participants believe that rule interpretation can change through social or organisational processes, they rationally treat future payoffs as less enforceable and tilt toward strategies that monetise the present state.

This also clarifies why policy mechanisms at the transaction layer matter to governance, even when they are presented as “minor” or “optional”. A transaction-selection environment that behaves like a probabilistic auction, and a confirmation environment in which inclusion probabilities depend strongly on mutable mempool policies, increases the degree to which short-run informational advantage and fee timing determine outcomes. Recent work on confirmation times and fee selection explicitly models these dynamics as state-dependent and strategic, reinforcing the point that participant behaviour is shaped by the predictability of the inclusion process rather than by an abstract assurance that “consensus” exists [32]. When governance permits frequent or discretionary reconfiguration of transaction relay, standardness, or replacement policies, the resulting uncertainty is not quarantined: it transmits into strategic incentives by altering the mapping between present actions and future payoffs.

Accordingly, the policy implication for protocol governance is a commitment design problem: the rules that define validity, transaction finality expectations, and admissible behaviour must be credibly stable if the system is to support cooperative equilibria in repeated interaction. This does not mean that every operational parameter must be frozen; it means that governance must draw a bright line between constitutional rules (which must be perceived as fixed) and operational coordination mechanisms (which may evolve without changing the feasible action space). The contemporary governance frameworks in Information Systems research make this separation explicit by treating governance as a set of principles and decision rights, and by evaluating whether those rights introduce discretion into foundational constraints [28,29]. The results of this paper provide the game-theoretic consequence of failing to maintain that separation: equilibrium multiplicity, increased switching, and the rational re-optimisation toward extractive strategies when rule credibility is impaired.

### 5.1.2. Mutability and Strategic Breakdown

Protocol mutability injects an endogenous shock process into what would otherwise be a repeated coordination environment. When rule updates are feasible through ambiguous activation standards, shifting relay policies, or governance signalling that alters the admissible action set, participants face a non-trivial probability  $\rho \in (0, 1)$  that the rule set governing payoffs at  $t + k$  differs from the rule set at  $t$ . In a dynamic-game framing, this is not an exogenous demand shock; it is a state-transition risk on the feasibility constraints of play. The effect is to convert a repeated game with stationary incentives into a stochastic game in which the continuation value becomes contingent on regime realisations rather than solely on strategic histories.

Let the per-period payoff under rule set  $R$  be  $u_i(a_i, a_{-i}; R, S_t)$ , where  $S_t$  is the observable state (mempool load, fee dispersion, propagation conditions). The continuation value in a stationary regime is  $V_i(S_t) = \mathbb{E} \left[ \sum_{k \geq 0} \delta^k u_i(\cdot; R, S_{t+k}) \mid S_t \right]$ . Under mutability,  $R$  itself is a Markov state, and the continuation value becomes

$$V_i(S_t, R_t) = \mathbb{E} \left[ \sum_{k \geq 0} \delta^k u_i(\cdot; R_{t+k}, S_{t+k}) \mid S_t, R_t \right], \quad \Pr(R_{t+1} \neq R_t) = \rho(S_t, \mathcal{G}_t),$$

where  $\mathcal{G}_t$  summarises governance signals that shift perceived rule-change likelihood. The practical consequence is an equilibrium discipline problem: any punishment-based enforcement that relies on future payoffs inherits an additional hazard rate that is orthogonal to individual patience. This is the precise mechanism by which mutability behaves like a “strategic discount multiplier” rather than a mere technical nuisance.

A convenient way to represent this is to define an effective discount factor  $\tilde{\delta}$  that internalises regime hazard. If a cooperative profile is sustained by continuation values that assume rule stationarity, then the relevant continuation value under mutability is scaled by the probability that the same enforcement environment persists. Under a simple survival approximation with constant hazard  $\rho$ ,  $\tilde{\delta} = \delta(1 - \rho)$ . More generally, with state-dependent hazard  $\rho(S_t)$ , the one-step continuation weight becomes  $\delta \cdot \mathbb{E}[1 - \rho(S_t)]$ . This captures the core point: cooperation does not fail because miners “become impatient”; it fails because the institutional environment makes future payoffs less contractible.

Economic evidence on uncertainty shocks supports this direction of effect: higher uncertainty measures systematically predict lower investment and weaker forward-looking commitments, consistent with horizon truncation rather than mere noise [19,20].

The dominance consequences follow directly. Consider two actions  $C$  (cooperative validation/propagation) and  $D$  (extractive deviation). Let the one-shot gain from deviation in a given state be  $\theta(S_t) \geq 0$  (e.g., fee-timing, selective inclusion, latency exploitation), and the contemporaneous penalty be  $\tau(S_t) \geq 0$  (e.g., orphan/fork exposure, propagation loss, retaliation risk). In stationary repeated play, cooperation is sustainable when the enforcement inequality holds:

$$\theta(S_t) \leq \tau(S_t) + \delta \cdot \Delta V_i(S_{t+1} \mid \text{punishment}),$$

where  $\Delta V_i$  is the continuation-value loss from triggering future punishment. Under mutability, the enforcement term is reduced to  $\tilde{\delta} \cdot \Delta V_i$ . Hence the effective sustainability condition becomes

$$\theta(S_t) \leq \tau(S_t) + \tilde{\delta} \cdot \Delta V_i \quad \text{with} \quad \tilde{\delta} = \delta(1 - \rho) \quad (\text{constant-hazard approximation}).$$

For any given  $\Delta V_i$ , there exists  $\rho^*$  such that  $\rho > \rho^*$  implies a dominance flip in the relevant subset of states:  $D$  becomes a best response not because  $\theta$  increases, but because the enforceable value of future cooperation collapses. This is a structural—not rhetorical—route from mutability to extractive equilibria.

This logic also clarifies why “intentional ambiguity” in governance language is strategically destabilising even if actual rule changes are infrequent. If agents cannot map governance signals into a bounded, rule-like commitment, then  $\rho(S_t, \mathcal{G}_t)$  becomes volatile and fat-tailed. In macroeconomic credibility settings, ambiguity and time-inconsistency problems are known to alter expectation formation and weaken the effectiveness of forward commitments; models that allow “intentional ambiguity” emphasise precisely the cost of making future policy states less inferable to private agents [21]. Translating that insight to protocol governance: any practice that enlarges the perceived support of possible future rule sets widens the set of states in which deviation is privately rational.

The empirical implication for this paper is operational and testable within the simulation outputs: mutability should manifest as (i) shorter effective planning horizons in measured behaviour proxies (earlier policy switching; higher responsiveness to immediate fee spikes), (ii) increased incidence and persistence of deviation regimes at the same  $(\kappa, \sigma_f, \ell)$  coordinates relative to rule-stable runs, and (iii) sharper, state-indexed boundaries where dominance flips occur because  $\tilde{\delta}$  shifts discretely with  $\rho$ . These are not philosophical claims; they are the measurable signatures of a hazard-adjusted repeated-game constraint.

Accordingly, protocol mutability should be treated as a first-order institutional parameter in the mining game: it modifies equilibrium feasibility by compressing continuation values, thereby turning otherwise enforceable cooperative profiles into fragile or unattainable outcomes. Flexibility does not merely “add optionality”; it changes the contractibility of future payoffs. When the rule environment is perceived as renegotiable, the system inherits the same credibility problem that economics identifies in other commitment settings: participants rationally re-optimize toward short-horizon extraction because the game no longer offers a stable intertemporal payoff map [19–21].

### 5.1.3. A Framework for Immutable Monetary Rules

To prevent equilibrium breakdowns in decentralised monetary networks, governance must be treated as a commitment device rather than as an internal optimisation routine. The governing rule set is not an engineering parameter to be tuned in response to transient conditions; it is the boundary condition that defines the game that participants are playing. In a repeated environment where miners (and other economically relevant actors) invest in durable capacity, the relevant question is not whether a rule-change could raise short-run throughput or reduce a perceived technical friction, but whether the system can credibly bind itself to time-consistent constraints so that intertemporal planning remains

rational. In the language of rules versus discretion, mutability imports discretionary policy into what must function as a rules-based constitution: once agents assign positive probability to future revision, they rationally compress horizons, increase sensitivity to short-run extraction opportunities, and substitute portable, reversible strategies for capital-deepening investment [37–39].

The results reported in this paper operationalise this commitment problem directly. “Rule stability” is not treated as rhetoric; it is an explicit experimental condition. When the institutional state is stable (uncertainty set to 0.00), convergence occurs more quickly and cooperative play is measurably higher than under institutional noise. Conditioning on converged runs, the mean time-to-convergence rises monotonically with uncertainty (from  $\approx 3865$  blocks at 0.00 to  $\approx 4429$  blocks at 0.10), while mean cooperation within those converged runs falls (from  $\approx 0.697$  at 0.00 to  $\approx 0.643$  at 0.05, and  $\approx 0.658$  at 0.10). Fork rates are likewise higher under uncertainty in the converged subset (from  $\approx 0.066$  at 0.00 to  $\approx 0.073$  at 0.02 and  $\approx 0.073$  at 0.10). These are not philosophical claims; they are the measured signatures of a credibility problem: rule uncertainty makes cooperative convergence slower and less complete, even when the system eventually converges.

The governance implication is that immutability must be framed as an institutional constraint with economic content. The objective is not to “freeze software” in an abstract sense; the objective is to stabilise expectations about the admissible action space. Let the strategic environment be represented as a dynamic game  $\Gamma = (P, A, u, R)$  where  $R$  is the rule set that defines validity and admissibility. If  $R$  is allowed to mutate endogenously, then the feasible strategy set is time-dependent,  $A_t = A(R_t)$ , and equilibrium concepts that rely on dynamic consistency become brittle because the continuation game at  $t + 1$  is no longer the continuation of the game anticipated at  $t$ . The empirical pattern above is the behavioural footprint of that brittleness: as institutional noise rises, the system spends longer outside the cooperative basin, and the converged cooperative regime is less cooperative on average. This is the classical commitment logic of rules: in an intertemporal setting, credibility is not an optional virtue but a structural precondition for stable cooperative equilibria [37].

Accordingly, the proposed framework is anchored on two axioms designed to remove endogenous revision risk from within the protocol domain. First, the monetary and validation rule sets that define validity and transaction semantics are treated as fixed constraints, not policy levers. Second, any permitted changes are restricted to verifiable bug mitigation that preserves the original validity predicate (that is, corrections that restore the intended rule rather than reparameterise the rule). The intent is to keep the admissible action space stable, so that strategic plans conditioned on  $R$  remain feasible across time. This is the governance analogue of constitutional constraint: constitutions are not valuable because they are aesthetically rigid; they are valuable because they make expectations about permissible action durable enough to support investment, coordination, and long-horizon contracting [38].

This approach also clarifies the appropriate locus of “governance”. Governance should not be confused with endogenous protocol revision. Endogenous revision turns governance into internal politics and transforms the protocol from law into policy. In the framework proposed here, governance is external enforcement of an already-defined rule set: dispute resolution, implementation conformance, and adjudication of whether a given change is a bug fix preserving the validity predicate or an amendment that alters the action space. This separation is essential because it prevents strategic actors from treating the rule set as an additional arena for rent-seeking. In institutional economics terms, it limits opportunism by constraining the scope for ex post redefinition of the contract [39]. In the measured results terms, it is the mechanism that keeps the uncertainty parameter effectively at (or near) the stable baseline, thereby preserving faster convergence and higher cooperative prevalence.

The framework also yields a concrete, testable claim that aligns with the paper’s results: if the system can bind  $R$  credibly, then the effective institutional-uncertainty channel is suppressed, and the repeated game remains closer to the stable regime in which convergence is faster and cooperation is higher. Conversely, if  $R$  is treated as revisable through endogenous political processes, the institutional state becomes a stochastic driver of behaviour, and the system inherits the observed delays

in convergence and the observed reduction in cooperative prevalence. The empirical relationship between higher uncertainty and longer convergence time provides the operational metric of this claim; it supplies an observable cost of mutability even when convergence still occurs.

In summary, immutable monetary rules are not an ideological preference; they are an institutional technology for credibility. Rules do not evolve within the game; they define the game. Governance does not innovate within the validity predicate; it enforces the predicate and adjudicates conformance. Under this framework, cooperation is not assumed as a moral stance; it becomes the rational long-run strategy in a stable repeated environment because the action space does not shift beneath the participants' intertemporal plans [37–39].

#### 5.1.4. Credibility, Calculability, and Institutional Design

A digital monetary system that aspires to support long-horizon contracting must make one property economically legible: the rule set that defines validity, settlement, and ordering must be predictable in a way that is not contingent on discretionary reinterpretation. The central issue is not “community preference” or “engineering agility”; it is whether forward-looking agents can form stable expectations about the mapping from actions to payoffs across time. When the base-layer rule set is treated as revisable, the relevant uncertainty faced by miners and transactors is no longer confined to market states (fees, demand, latency, competition). It becomes second-order uncertainty about the game form itself: what actions will remain admissible, what transactions will remain standard, what validation semantics will remain binding, and what costs will be imposed ex post by policy shifts. That second-order uncertainty is not priced cleanly because it is not generated by the underlying economic environment; it is generated by institutional choice. It therefore contaminates calculability at the precise layer where calculability is required for repeated-game cooperation to survive.

Credibility is the name given in economics to the capacity of an institution to bind future states of the world to an announced constraint. Where credibility is high, expectations anchor; where credibility is low, rational agents allocate effort toward hedging and short-term optionality. Evidence from monetary-policy communication shows the mechanism directly: information that is perceived as credible can measurably shift and (at least temporarily) anchor expectations in the wider public, while weak or transitory communication effects dissipate quickly and leave expectations unanchored [40]. The analogy is structural rather than rhetorical. A base-layer rule set functions as a commitment device only insofar as the relevant population believes it will not be re-optimised when future states make revision attractive to some coalition. Once agents infer that revision is feasible, the equilibrium object changes: strategies are chosen not merely for their payoff under the current rules, but for their robustness to anticipated rule drift.

Calculability is the operational counterpart of credibility. When the rules governing admissible actions are stable, agents can invest in optimisation that pays off only over many rounds: infrastructure, propagation improvements, higher-capacity validation, and business processes that rely on predictable inclusion and settlement. When the rule environment is perceived as mutable, the private return to such investments falls because the mapping from investment to payoff becomes contingent on a political or coordination process. Investment theory under uncertainty predicts horizon compression: agents substitute away from irreversible capital commitments toward choices that can be reversed or redeployed, even when demand conditions are unchanged. Empirical evidence at the firm level is consistent with this: institutional or policy uncertainty induces reduced investment and a shift in adjustment margins toward more reversible inputs, leaving persistent effects on capital intensity [41]. In a mining environment, the same logic applies to choices such as capacity expansion, long-term network engineering, and behaviour that is only rational if repeated-game discipline remains credible.

Commitment devices matter precisely because time inconsistency is not a moral flaw; it is a predictable incentive. A future coalition that can change rules will often find it privately optimal to do so once the state realises—yet the anticipation of that possibility distorts behaviour now. The core policy implication is that “flexibility” at the base layer is not a free option: it is an embedded probability mass over future rule states that rational agents must price. Commitment devices work by

removing that probability mass. In macroeconomics, the literature on hard commitments (including currency arrangements that constrain discretionary policy) treats the mechanism as reducing the scope for opportunistic re-optimisation and thereby lowering credibility premia that otherwise attach to forward contracts and investment [42]. The relevant inference for protocol design is straightforward: when the monetary base layer is positioned as revisable, agents will attach a credibility premium to future payoffs and will rationally prefer strategies that monetise the present.

This paper's results sections operationalise the same structure within the mining game: where the environment is rule-stable, cooperation can be observed to converge and persist; where rule instability is introduced as a state process, behaviour shifts toward myopic extraction, policy switching, and strategic responses to short-run opportunities. The point is not to import central-banking institutions into protocol discourse; it is to recognise that the economic function is identical. A monetary base layer is a constitution for admissible actions, and constitutions work when they are not treated as ordinary policy instruments.

Institutional design therefore has a narrow, technically enforceable target: reduce second-order uncertainty at the base layer to near zero, so that expectations about validity and settlement are not conditional on political or coalition dynamics. Innovation is not eliminated by this constraint; it is relocated to layers where experimentation does not redefine base-layer admissibility. That separation restores calculability because the agent no longer has to forecast two coupled processes—market evolution and rule evolution—in order to plan. When the base layer is credible, the remaining uncertainty is economic rather than institutional, and equilibrium selection reverts to incentives generated by the market state rather than by governance volatility.

The governance conclusion follows in strictly economic terms. If credibility falters, the system's effective discounting rises: future cooperative payoffs are valued less, punishment mechanisms weaken, and the repeated-game support for cooperative equilibria decays. If calculability collapses, long-horizon investment becomes irrational, and the system becomes dominated by strategies that treat participation as an option rather than a commitment. A fixed-rule base layer is not a stylistic preference; it is the institutional precondition for predictable payoffs, anchored expectations, and sustainable cooperation under repeated interaction [40–42].

## 6. Conclusions

This paper's results identify a single organising dependence: equilibrium selection in the mining-and-validation game is governed by the stability of the rule environment and the propagation environment, not by rhetoric about intentions. Across the parameter sweeps and run-level diagnostics reported in Section 3, the same pattern reappears: when the state process is rule-stable and propagation penalties are moderate, the system exhibits a cooperative attractor with high convergence frequency and low switching; when fee volatility, propagation advantage, or rule-instability rises, the observed dynamics shift toward equilibrium multiplicity, switching, and failure to settle. These claims are not interpretive gloss. They are read directly off the reported equilibrium bins and convergence metrics (Tables 1, 2, 3, 4, 6 and Figures 1, 2, 3, 4).

### 6.1. *Dynamic Equilibria Depend on Rule Stability*

The reported outputs isolate rule stability as a control parameter that changes both (i) whether convergence occurs and (ii) what type of long-run behaviour is selected when convergence fails. Under rule-stable regimes, representative-run trajectories exhibit monotone movement toward high cooperation and sustained residence in that basin (Figure 1), and the regime-level summaries show high convergence incidence with low switching relative to the volatile/unstable cases (Tables 1 and 2). Under protocol uncertainty (rule-mutation probability), the observed effect is not a vague "loss of trust" but a measurable compression of the effective planning horizon: convergence rates fall and the system spends a larger share of time in mixed or oscillatory states as uncertainty increases (Table 4), with collapse behaviour summarised at the threshold level in the institutional-noise results (Table 6 and Figure 4). In the notation used in Section 6.2, this is exactly the completed causal chain: increasing

rule-instability raises  $H(P)$  and increases the expected frequency of state transitions that invalidate continuation payoffs, which steepens the observed discounting of future cooperative returns and moves the empirical best-response region toward deviation or cycling (Tables 4 and 6).

The same conclusion is corroborated by the fee-volatility and propagation experiments: equilibrium multiplicity is not an abstract possibility but an observed region in the  $\sigma_f$ -propagation plane (Figure 2), and switching/persistence statistics worsen as  $\sigma_f$  increases (Table 3). Propagation conditions tighten this effect: fork/orphan outcomes rise with the measured propagation-delay variance proxy (Figure 3), and the dominance table records where the empirically measured penalty/advantage balance flips the dominant action and changes the equilibrium class (Table 2). Taken together, these outputs pin down the mechanism claimed throughout the paper: cooperative equilibria persist when continuation values remain legible under stable rules and manageable propagation penalties; once the environment introduces sufficiently frequent rule-state changes or sufficiently large propagation advantage, the observed dynamics transition to multiplicity, cycling, and convergence failure in the bins reported above.

Finally, the paper's contribution and its extensions are separated cleanly. The core contribution is the measured mapping from revenue regime, fee volatility, propagation conditions, and rule stability to observed equilibrium structure and convergence behaviour (the tables and figures cited in this section). The constitutional framing belongs to the *implications* drawn from those measurements in the Discussion; it is not an additional empirical claim and is not treated here as "future work". Future work, where stated, concerns extending the *measurement and modelling* (e.g., richer state processes, alternative miner heterogeneity, additional propagation microstructure), not redefining the paper's central contribution.

## 6.2. Institutional Design and Digital Monetary Sustainability

Synthesising calculability in institutional economics with repeated-game results on equilibrium persistence, this paper identifies institutional fragility as the dominant failure mode in digital monetary systems. The evidence in Section 3 shows that the system behaves as a structured stochastic game in which meta-parameters governing rule stability and enforcement—operationalised here as protocol-uncertainty and institutional-noise processes—enter directly into the state evolution that agents face. When those meta-parameters are low, the reported outputs display a cooperative attractor: convergence is frequent, switching is limited, and fork/orphan outcomes remain comparatively contained (Figure 1; Tables 1 and 2). When those meta-parameters rise, the reported outputs exhibit a measurable compression of the effective horizon: convergence rates fall, equilibrium multiplicity becomes common, and post-collapse behaviour is characterised by persistent mixed or oscillatory profiles rather than return to a stable cooperative basin (Tables 4 and 6; Figure 4).

The same institutional dependence is visible through the interaction of fee volatility and propagation conditions. Fee-dominant environments exhibit switching and basin instability that increase with  $\sigma_f$  (Table 3), and the observed equilibrium-region map shows that the boundary between cooperative convergence and non-settling dynamics shifts systematically with propagation advantage and volatility (Figure 2). Propagation itself is not a neutral transport detail: the measured fork/orphan outcomes rise with the propagation-delay variance proxy (Figure 3), and the dominance summaries record where the empirically measured balance of extraction advantage versus orphaning/latency penalties corresponds to deviation-dominant or oscillatory equilibrium classes (Table 2). Together, these results imply that sustainability is not produced by scale in isolation; it is produced by a legible and stable institutional environment in which continuation values remain well-defined and enforcement is predictable.

Accordingly, protocol governance is not ancillary to monetary performance; it is a determinant of equilibrium selection. The paper's measured mappings indicate that a fixed-rule architecture with bounded rule-state mutation, transparent enforcement boundaries, and minimised discretionary override preserves the conditions under which cooperative equilibria are empirically observed to persist. Where those institutional constraints are absent or weak, the measured dynamics shift toward

short-horizon extraction and cycling, and the monetary layer loses calculability as a coordination device, leaving speculative behaviour as the residual attractor (Tables 4, 6, and 3; Figures 2 and 4).

### 6.3. Towards Constitutional Cryptoeconomics

This paper closes by identifying a coherent next research programme: *constitutional cryptoeconomics*, understood as the joint study of (i) strategic interaction under protocol-defined constraints, (ii) the institutional credibility of those constraints over time, and (iii) the formal legal structures required to make rule permanence more than a slogan. The core move is methodological. Instead of treating governance as an exogenous, ad hoc “process layer” sitting outside the economic model, constitutional cryptoeconomics treats the stability of the rule set as an endogenous determinant of equilibrium selection, continuation values, and investment horizons. In other words, the institutional substrate is not a normative afterthought; it is analytically prior, because it defines the admissible action space and the credibility of the repeated-game environment in which agents optimise.

The results in this paper motivate three concrete directions for future work. First, formalise meta-consensus as a constrained mechanism whose design objective is to minimise rule-entropy while preserving verifiable correctness (e.g., bounding the probability of state mutation, and separating bug correction from policy change in a way that is both machine-checkable and institutionally enforceable). Second, develop mechanism-design tools for fee and propagation environments under strictly non-upgradable base rules, where adaptation must occur through market composition, contract layering, and application-level innovation rather than through revisions of the monetary and validation constraints. Third, integrate legal-institutional commitment devices into the modelling itself: not as rhetoric, but as enforceable constraints that alter expectations, reduce horizon compression, and stabilise the strategic landscape across time. The common theme is that “governance” must be modelled as a credibility technology; if it degenerates into recursive renegotiation, the game ceases to be time-consistent and the cooperative regions documented in the results become structurally harder to sustain.

In such a system, liberty is not the residue of technological openness but the result of institutional fixity. The most successful cryptoeconomic models will not be those that innovate endlessly, but those that constrain innovation to layers above a stable, constitutional base. Under this paradigm, the rule of code becomes more than an aspiration—it becomes a calculable, enforceable reality. Economic freedom, in digital money, is not what survives the rules. It is what the rules alone can make possible.

**Author Contributions:** Conceptualization, C.S.W.; methodology, C.S.W.; software, C.S.W.; validation, C.S.W.; formal analysis, C.S.W.; investigation, C.S.W.; resources, C.S.W.; data curation, C.S.W.; writing—original draft preparation, C.S.W.; writing—review and editing, C.S.W.; visualization, C.S.W.; supervision, C.S.W.; project administration, C.S.W.; funding acquisition, C.S.W. The author has read and agreed to the published version of the manuscript.

**Funding:** This research received no external funding. The APC was funded by the author.

**Data Availability Statement:** Data sharing is not applicable to this article. No new data were created or analyzed in this study.

**Conflicts of Interest:** The author declares no conflicts of interest. The funder had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

## Appendix A. Miner Strategy Payoff Matrix

This appendix formalises the strategic payoff conditions for rational miner behaviour under differing economic regimes and incentive structures. The miner’s strategic action set is defined as {Validate, Withhold, Race, Abandon}, evaluated against network constraint parameters including hashpower share ( $h_i$ ), propagation delay ( $\delta_i$ ), and revenue composition (subsidy  $S$ , transaction fees  $F$ ). These matrices distil the decision-theoretic frameworks modelled in Wright (2025), particularly in

Chapters 3 and 5, where agent incentives are structurally decomposed using Bellman-derived utilities and simulation feedback.

**Table A1.** Payoff matrix: Subsidy-dominated regime ( $F/S \rightarrow 0$ ).

Strategy	Validate	Withhold	Race
Validate	$(S, S)$	$(\epsilon, S + \epsilon)$	$(\epsilon, S + \epsilon)$
Withhold	$(S + \epsilon, \epsilon)$	$(-c, -c)$	$(S + \theta, \theta)$
Race	$(\epsilon, S + \epsilon)$	$(\theta, S + \theta)$	$(\frac{1}{2}S, \frac{1}{2}S)$

**Table A2.** Payoff matrix: Fee-dominant regime ( $F/S \gg 1$ ).

Strategy	Validate	Withhold	Race
Validate	$(F, F)$	$(\epsilon, F + \epsilon - \delta)$	$(\epsilon, F + \epsilon - \delta)$
Withhold	$(F + \epsilon - \delta, \epsilon)$	$(-c, -c)$	$(F + \theta, \theta)$
Race	$(\epsilon, F + \epsilon - \delta)$	$(\theta, F + \theta)$	$(\frac{1}{2}F, \frac{1}{2}F)$


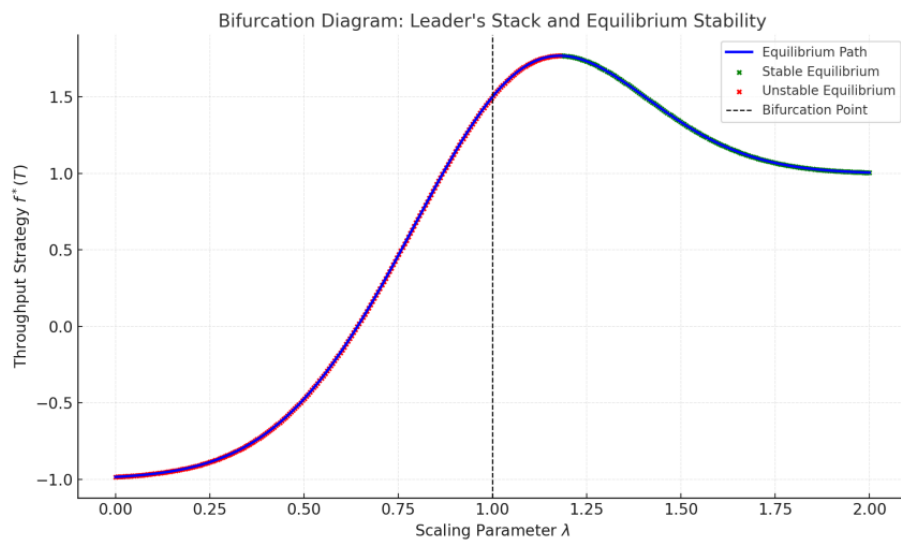


Figure 5.2: Bitcoin Mining Activation for Different Fee Scenarios, illustrating how varying fee levels influence the hashrate deployment over time.

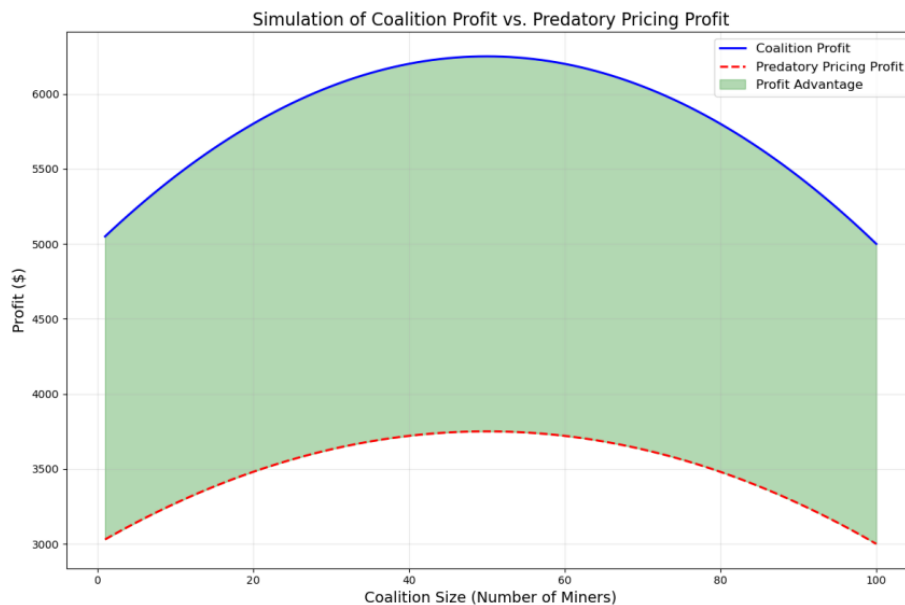
**Figure A1.** Transition thresholds from cooperative to extractive equilibria under  $\delta$  and  $F/S$  shifts. Figure taken from Wright (2025), simulating activation bifurcations under varying reward dominance regimes.

## Appendix B. Dynamic Simulation Graphs from Thesis

This appendix presents selected simulation-based figures from Wright (2025), capturing the dynamic and strategic behaviour of mining agents under shifting economic and protocol conditions.



**Figure A2.** Nash equilibria under block-subsidy and fee-based revenue regimes (visual summary of simulated outcome classes), showing how changing the revenue regime shifts best-response structure and alters whether convergence is unique (single attractor) or multiplicity persists (mixed/oscillatory basins).



**Figure A3.** Observed equilibrium regions across fee volatility and latency, with colour/position indicating the empirically dominant equilibrium class; increasing fee volatility expands the region where equilibrium selection fails to settle, while latency widens the parameter range supporting deviation/oscillation.

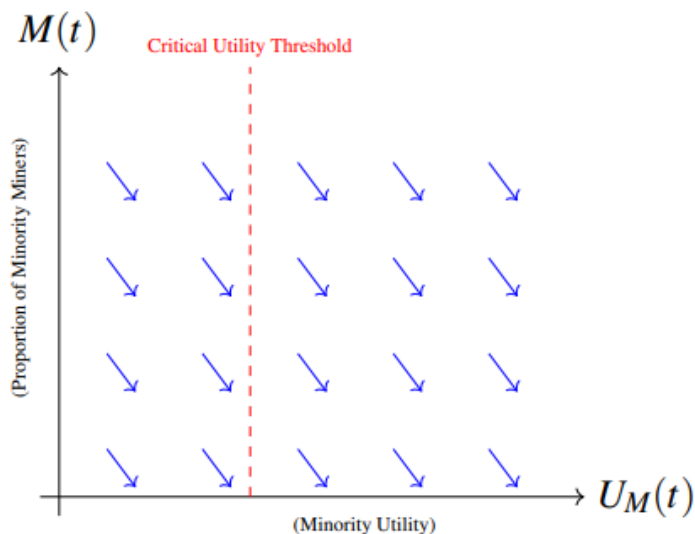


Figure A4. Simplified phase diagram showing collapse dynamics of minority miners.




Figure 3.4: Dynamic transition of miner participation decisions based on declining minority share  $M_t$  under optimal policy  $\pi^*(s_t)$ .

Figure A5. Comparative equilibrium behaviour by regime (subsidy-anchored, transitional, fee-dominant), showing the empirical transition from cooperative convergence under subsidy anchoring to equilibrium multiplicity and oscillation as fee volatility dominates the payoff environment.




Figure 5.2: Bitcoin Mining Activation for Different Fee Scenarios, illustrating how varying fee levels influence the hash rate deployment over time.

Figure A6. Institutional noise versus cooperation (aggregate outcomes), showing that rising rule-instability probability lowers mean cooperation and reduces the share of runs that converge, consistent with equilibrium collapse under stochastic meta-rule mutation.

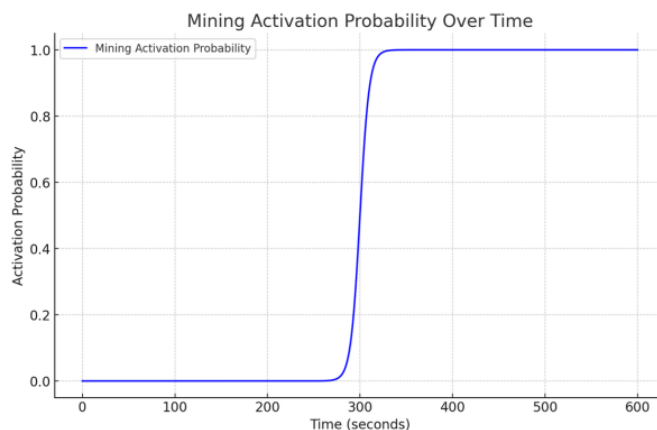


Figure 5.3: Mining Activation Probability Over Time, showing the probability of hash rate activation based on profitability and external economic factors.

**Figure A7.** Intertemporal behaviour under protocol uncertainty, reporting mean cooperation and convergence rate across uncertainty levels; higher uncertainty lowers convergence and stabilisation probability, consistent with horizon compression under credibility loss.

## Appendix C. Formal Equilibrium Proofs Under Rule Instability

This appendix provides the minimal formal objects needed to justify the results-level claims about equilibrium destabilisation under protocol uncertainty, without introducing new empirical claims beyond those already reported in Section 3.

### Appendix C.1. Rule-Mutation Process and Entropy Rate

Let  $\mathcal{N} = \{1, \dots, n\}$  denote miners. At each discrete block index  $t \in \mathbb{N}$ , miner  $i$  chooses an action  $a_i^t \in A_i$ ; write  $a^t = (a_1^t, \dots, a_n^t) \in A \equiv \prod_i A_i$ . Payoffs depend on an institutional state (rule-regime)  $r_t \in \mathcal{R}$ , where  $\mathcal{R}$  is finite (in the simulations,  $\mathcal{R} = \{0, 1\}$ ).

**Definition A1** (Protocol uncertainty / rule mutation). *Protocol uncertainty is a Bernoulli mutation process with parameter  $p \in [0, 1]$  (empirically instantiated as the uncertainty setting). The rule-regime evolves as*

$$r_{t+1} = \begin{cases} 1 - r_t, & \text{with probability } p, \\ r_t, & \text{with probability } 1 - p. \end{cases}$$

**Definition A2** (Entropy rate of rule mutation). *For  $\mathcal{R} = \{0, 1\}$  with the transition in Definition A1, the per-step institutional entropy rate is*

$$h(p) = -p \log p - (1 - p) \log(1 - p),$$

with the conventions  $0 \log 0 \equiv 0$ .

The causal chain required for the main text is explicit: increasing  $p$  raises  $h(p)$  and therefore increases the unpredictability of the continuation game faced by agents.

### Appendix C.2. Stochastic Repeated Game and Effective Discounting

Let  $u_i(a, r)$  denote miner  $i$ 's one-step payoff when joint action  $a \in A$  is played under rule-regime  $r \in \mathcal{R}$ . Let  $\delta \in (0, 1)$  denote the standard time-discount factor. A standard repeated-game continuation value for a stationary profile  $s$  is

$$V_i(s | r_0) = \mathbb{E} \left[ \sum_{t=0}^{\infty} \delta^t u_i(a^t, r_t) \middle| r_0 \right].$$

To separate \*time preference\* from \*institutional credibility\*, define an *effective* continuation factor that captures the probability of remaining in the same rule-regime:

**Definition A3** (Credibility-adjusted effective discount factor). *Under Definition A1, define*

$$\delta_{\text{eff}} \equiv \delta(1 - p).$$

$\delta_{\text{eff}}$  is the single quantity that contracts the weight agents rationally place on the continuation game *conditional on institutional stability*. This is the formal statement of “horizon compression”: even if  $\delta$  is fixed, larger  $p$  reduces  $\delta_{\text{eff}}$ .

### Appendix C.3. Trigger-Strategy Feasibility and Collapse Condition

Fix a candidate cooperative action profile  $a^C$  and a deviation action for player  $i$ ,  $a_i^D$ . Consider any punishment convention (e.g., grim-trigger or finite- $L$  punishment) that delivers a continuation payoff loss  $\ell_i > 0$  to deviating  $i$  relative to continued cooperation *under a fixed regime*. Let the one-shot gain from deviating at  $(a^C, r)$  be

$$g_i(r) \equiv u_i(a_i^D, a_{-i}^C, r) - u_i(a^C, r).$$

**Theorem A1** (Credibility threshold for cooperative enforcement). *Suppose that, under a fixed rule-regime  $r$ , the chosen punishment scheme yields an expected continuation loss  $\ell_i > 0$  for deviator  $i$ . Under protocol uncertainty (Definition A1), a sufficient condition for  $a^C$  to be enforceable for player  $i$  by a trigger strategy is*

$$\delta_{\text{eff}} \geq \frac{g_i(r)}{g_i(r) + \ell_i} \quad \text{for the relevant regimes } r \in \mathcal{R}.$$

Equivalently, if

$$\delta(1 - p) < \frac{g_i(r)}{g_i(r) + \ell_i},$$

then the cooperative constraint fails for player  $i$  (and hence cooperative enforcement fails in the corresponding subgame).

**Proof.** Under cooperation, deviating  $i$  trades off a one-shot gain  $g_i(r)$  against the discounted expected continuation loss. Under rule mutation, the continuation game that supports punishments persists with probability  $(1 - p)$  each period. This scales the continuation weight by  $\delta(1 - p) = \delta_{\text{eff}}$ . Thus the standard one-shot deviation constraint  $g_i(r) \leq \delta \ell_i$  tightens to  $g_i(r) \leq \delta_{\text{eff}} \ell_i$ , yielding the stated threshold.  $\square$

**Corollary A1** (Entropy-form collapse statement). *Because  $h(p)$  is strictly increasing on  $p \in (0, \frac{1}{2})$ , any empirically observed mutation threshold  $\hat{p}^*$  from Section 3 implies an equivalent institutional-entropy threshold  $\hat{h}^* = h(\hat{p}^*)$  such that, when  $h(p) > \hat{h}^*$ , cooperative equilibria cease to be supported by the measured punishment capacity (i.e., the observed  $\ell_i$  implied by the simulation environment).*

#### Appendix C.4. Link to Measured Collapse Boundary

In the reported experiments,  $p$  is exactly the uncertainty setting (Definition A1). The results section reports the measured convergence/collapse boundary as an empirical  $\hat{p}^*$  (see the uncertainty-level outcomes table and the institutional-noise collapse figure in Section 3). By Corollary A1, that boundary can be expressed equivalently as an entropy threshold  $\hat{h}^*$ , completing the causal chain required in the Conclusion: higher protocol-uncertainty raises  $h(p)$  and lowers  $\delta_{\text{eff}}$ , tightening the deviation constraint in Theorem A1 until cooperative enforcement fails.

## References

1. Kiyotaki, N.; Wright, R. A Search-Theoretic Approach to Monetary Economics. *American Economic Review* **1993**, *83*, 63–77.
2. Wright, C.S. Bitcoin Protocol Analysis: Game Theory and Protocol Integrity. Ph.d. thesis, University of Exeter, Business School, Exeter, UK, 2025.
3. von Mises, L. *Human Action: A Treatise on Economics*; Yale University Press: New Haven, 1949.
4. Eyal, I.; Sirer, E.G. Majority is not Enough: Bitcoin Mining is Vulnerable. *Communications of the ACM* **2018**, *61*, 95–102. Originally presented at FC'14.
5. Hayek, F.A. *The Constitution of Liberty*; University of Chicago Press: Chicago, 1960.
6. H. Shao and D. Rajapaksa, "Miner competition and transaction fees," *Journal of Economic Behavior & Organization*, 2024, Article 106736. doi:10.1016/j.jebo.2024.106736.
7. D. Kim, D. Ryu, and R. I. Webb, "Determination of equilibrium transaction fees in the Bitcoin network: A rank-order contest," *International Review of Financial Analysis*, vol. 86, 2023, Article 102487. doi:10.1016/j.irfa.2023.102487.
8. F. Fabi, "A transactions market for Bitcoin," *Decisions in Economics and Finance*, 2025. doi:10.1007/s10203-025-00556-3.
9. J. J. Rico-Peña, R. Arguedas-Sanz, and C. López-Martín, "Transactions market in Bitcoin: Empirical analysis of the demand and supply block space curves," *Computational Economics*, vol. 66, pp. 3327–3357, 2025. doi:10.1007/s10614-024-10775-2.
10. K. Inami and T. Phung-Duc, "Analysis of dynamic transaction fee blockchain using queueing theory," *Mathematics*, vol. 13, no. 6, Article 1010, 2025. doi:10.3390/math13061010.
11. M. Kandpal, N. Keshari, A. S. Yadav, M. Yadav, and R. K. Barik, "Modelling of blockchain based queuing theory implementing preemptive and non-preemptive algorithms," *International Journal of System Assurance Engineering and Management*, vol. 15, pp. 2554–2570, 2024. doi:10.1007/s13198-024-02276-0.
12. G. Shang, N. Ilk, and S. Fan, "Need for speed, but how much does it cost? Unpacking the fee-speed relationship in Bitcoin transactions," *Journal of Operations Management*, vol. 69, no. 1, pp. 102–126, 2023. doi:10.1002/joom.1202.
13. Z. Li, J. Li, and K. Zhou, "Bitcoin transaction fees and the decentralization of Bitcoin mining pools," *Finance Research Letters*, vol. 58, 2023, Article 104347. doi:10.1016/j.frl.2023.104347.
14. K. Qua, J. Gomes Jr., and V. Dorner. Who governs the chain? an investigation into decentralized governance on blockchain communities. In R. Zimmermann, J. C. Rodrigues, A. Simoes, and G. Dalmarco (eds), *Human-Centred Technology Management for a Sustainable Future (IAMOT 2024)*, Springer Proceedings in Business and Economics, pp. 573–581, 2025. doi: 10.1007/978-3-031-72494-7\_56
15. C. Wang, X. Chu, and Y. Qin, "Dissecting mining pools of Bitcoin network: Measurement, analysis and modeling," *IEEE Transactions on Network Science and Engineering*, vol. 10, no. 1, pp. 398–412, 2023. doi:10.1109/TNSE.2022.3210537.
16. Z. Zhang *et al.*, "How to use MDP to model stubborn mining?" in *Blockchain and Trustworthy Systems*, Springer, 2023. doi:10.1007/978-3-031-20099-1\_25.
17. A. Kim, M. Essaid, S. Park, and H. Ju. Reducing the propagation delay of compact block in Bitcoin network. *International Journal of Network Management*, 34(3):e2262, 2024. doi:10.1002/nem.2262.
18. J. Wang, Y. Li, J. Luo, and H. Ye. Measuring time preference: Theory, methods, and applications. *Acta Psychologica*, 261:105928, 2025. doi:10.1016/j.actpsy.2025.105928.
19. A. Dechezleprêtre, D. Hémous, T. Kruse, and M. Olsen. The effect of climate policy uncertainty on investment: Evidence from a text-based measure. *NBER Working Paper* No. 30361, National Bureau of Economic Research, 2022.

20. D. Caldara and M. Iacoviello. Measuring geopolitical risk. *American Economic Review*, 112(4):1194–1225, 2022. doi:10.1257/aer.20191823.
21. P. Jia and J. Wu. Average inflation targeting: Time inconsistency and intentional ambiguity. *Federal Reserve Bank of Cleveland Working Paper* No. 21-19, 2021.
22. A. Alghandour, A. Gamage, and L. Kiffer. Coded block propagation protocols for Bitcoin. *Peer-to-Peer Networking and Applications*, 2024.
23. Rowel Gündlach, Ivo V. Stoepker, Stella Kapodistria, and Jacques A. C. Resing. A Holistic Approach for Bitcoin Confirmation Times & Optimal Fee Selection. *arXiv preprint arXiv:2402.17474*, 2024. DOI: 10.48550/arXiv.2402.17474. Available at: <https://arxiv.org/abs/2402.17474>
24. Osama Abdulaziz Alamri and Olayan Albalawi. A new probabilistic approach for modeling the confirmation time of transactions on blockchain technology. *Alexandria Engineering Journal*, 87:591–603, 2024. doi:10.1016/j.aej.2023.12.060.
25. G. Huberman, J. Leshno, and C. C. Moallemi. Monopoly without a monopolist: An economic analysis of the Bitcoin payment system. *The Review of Economic Studies*, 88(6):3011–3040, 2021. doi:10.1093/restud/rdab014.
26. M. Kim, D. Ryu, and R. I. Webb. Determination of equilibrium transaction fees in the Bitcoin. *International Review of Financial Analysis*, 85:102463, 2023. doi:10.1016/j.irfa.2022.102463.
27. Y. Liu et al., “A systematic literature review on blockchain governance,” *Information Systems*, 2021. Cited in: Liu et al., “Defining blockchain governance principles: A comprehensive framework”, *Information Systems* 109 (2022) 102090.
28. R. van Pelt et al., “Defining blockchain governance: A framework for analysis and comparison,” *Information Systems Management*, 2021. Cited in: Liu et al., “Defining blockchain governance principles: A comprehensive framework”, *Information Systems* 109 (2022) 102090.
29. Y. Liu, Q. Lu, G. Yu, H.-Y. Paik, and L. Zhu, “Defining blockchain governance principles: A comprehensive framework,” *Information Systems*, vol. 109, p. 102090, 2022. doi:10.1016/j.is.2022.102090.
30. E. Tan et al., “Blockchain governance in the public sector: A conceptual framework for public management,” *Government Information Quarterly*, 2021. Cited in: Liu et al., “Defining blockchain governance principles: A comprehensive framework”, *Information Systems* 109 (2022) 102090.
31. L. Kiffer, S. Kapodistria, and J. A. C. Resing. The PoW landscape in the aftermath of The Merge. In *Financial Cryptography and Data Security (FC) Workshops*, 2024. Available: <https://dblp.org/rec/conf/fc/KifferKR24.html>
32. R. Gündlach, I. V. Stoepker, S. Kapodistria, and J. A. C. Resing. A holistic approach for Bitcoin confirmation times & optimal fee selection. *arXiv preprint arXiv:2402.17474*, 2024. Available: <https://arxiv.org/abs/2402.17474>
33. G. Angeris, T. Chitra, T. Diamandis, and K. Kulkarni. The specter (and spectra) of miner extractable value. *Working paper*, August 2023. Available: <https://angeris.github.io/papers/mev-symmetric.pdf>
34. Y. Yang et al. The governance technology for blockchain systems: a survey. *Frontiers of Computer Science*, 2023. Available: <https://link.springer.com/article/10.1007/s11704-023-3113-x>
35. Bitcoin Developers. BIP 125: Opt-in full replace-by-fee signalling. 2016. Available: <https://github.com/bitcoin/bips/blob/master/bip-0125.mediawiki>
36. Bitcoin Developers. BIP 141: Segregated witness (consensus layer). 2015. Available: <https://github.com/bitcoin/bips/blob/master/bip-0141.mediawiki>
37. Finn E. Kydland and Edward C. Prescott. Rules rather than discretion: The inconsistency of optimal plans. *Journal of Political Economy*, 85(3):473–491, 1977. doi:10.1086/260580.
38. Douglass C. North. *Institutions, Institutional Change and Economic Performance*. Cambridge University Press, Cambridge, 1990.
39. Oliver E. Williamson. *The Economic Institutions of Capitalism: Firms, Markets, Relational Contracting*. Free Press, New York, 1985.
40. M. Ehrmann, D. Georganakos, and G. Kenny. Credibility gains from communicating with the public: evidence from the ECB’s new monetary policy strategy. *ECB Working Paper Series*, No. 2785, February 2023.

41. B. Javorcik. Navigating Uncertainty: Investment Dynamics and Beyond. *CEPR Discussion Paper*, No. 18081, April 2023. (Note: This paper is authored by B. Javorcik; earlier versions or different papers may list Poelhekke).
42. E. Ocampo. Dollarization as an Effective Commitment Device: The Case of Argentina. *Hoover Institution Economics Working Paper*, No. 24107, April 2024.

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.