

Blockchain as Game Theory

Incentives, Equilibria, and Security in Decentralized Networks

Amy O. Khaldoun

December 2025

Abstract

A permissionless blockchain is a distributed protocol that must remain stable without a central enforcer. Its security and liveness therefore depend not only on cryptography and networking, but also on strategic behavior. This research note frames a blockchain as a hierarchy of games, from consensus participation to transaction propagation and token holding decisions. We formalize the main player types and equilibrium concepts used in blockchain analysis, discuss how strategic deviations arise as equilibrium responses, and connect protocol security to aggregate equilibrium quantities such as total hashrate. We then summarize two mean field perspectives: one for Proof of Work mining and one for token economics, emphasizing how microscopic incentives shape macroscopic outcomes.

Contents

| | |
|--|----------|
| 1 Why blockchain is fundamentally game theoretic | 2 |
| 2 Layered view: where the games live | 2 |
| 3 Game models used in blockchain research | 2 |
| 3.1 Static non cooperative games | 2 |
| 3.2 Repeated games | 3 |
| 3.3 Stochastic games and Markov equilibria | 3 |
| 3.4 Mean field games and mean field type games | 3 |
| 4 Formal setup and equilibrium concepts | 3 |
| 5 Consensus and incentives as mechanism design | 3 |
| 6 Security as equilibrium behavior | 4 |
| 6.1 Selfish mining and withholding intuition | 4 |
| 6.2 Propagation incentives and message timing | 4 |
| 6.3 Fee driven forking and undercutting intuition | 4 |
| 7 Proof of Work mining as a mean field game | 4 |
| 7.1 Mining as an aggregate contest | 4 |
| 7.2 Continuum model and objective | 5 |
| 7.3 Real hashrate, technological progress, and stationary security | 5 |
| 7.4 Comparative statics and security margin | 5 |
| 8 Token economics as a mean field type game | 5 |
| 8.1 Variance aware utility and distribution dependence | 5 |
| 8.2 Qualitative equilibrium implications | 6 |
| 8.3 Equilibrium price and token velocity | 6 |

| | |
|---|----------|
| 9 Design checklist for protocol builders | 6 |
| 9.1 Step one: define players | 6 |
| 9.2 Step two: enumerate actions including deviations | 6 |
| 9.3 Step three: define payoffs including costs and risk | 6 |
| 9.4 Step four: solve or approximate equilibrium | 6 |
| 9.5 Step five: connect equilibrium to security margin | 6 |
| 10 Open research directions | 6 |
| 11 Conclusion | 7 |

1 Why blockchain is fundamentally game theoretic

A public blockchain invites open participation. Miners, validators, pools, users, and application agents can choose actions that affect network outcomes and each other. This makes the protocol a strategic environment. Game theory provides the language to describe this environment through players, strategies, payoffs, and equilibrium.

The design goal is simple to state: honest participation should be a best response under realistic assumptions about costs, information, and objectives. When this holds, protocol compliance emerges from incentives rather than from trust or enforcement.

2 Layered view: where the games live

A convenient way to structure analysis is the layered stack often used in blockchain architecture:

- **Data layer:** cryptographic commitments, signatures, timestamps, and ledger structure.
- **Network layer:** peer to peer propagation, verification messaging, connectivity and latency.
- **Consensus layer:** leader selection, block validity rules, chain selection, finality.
- **Incentive layer:** issuance, rewards, penalties, and allocation mechanisms.
- **Contract layer:** scripts and smart contracts that implement additional rules.
- **Application layer:** markets and products that introduce new strategic incentives.

Strategic behavior appears everywhere, but it becomes unavoidable in consensus and incentives. Once rewards and penalties exist, participants optimize around them. Applications can then add a second game on top of the base game.

3 Game models used in blockchain research

Different protocol components map naturally to different game models.

3.1 Static non cooperative games

A static game models a one shot interaction: each player chooses an action once, payoffs are realized, and the game ends. This is useful for settings such as immediate pool choice, parameter choice under a fixed environment, or one time deviation decisions.

3.2 Repeated games

Protocols run continuously. Many decisions repeat: include transactions, propagate messages, build on one chain or another, vote, and so on. Repeated games model how future consequences shape current behavior. Reputation and punishment mechanisms are naturally expressed in repeated game form.

3.3 Stochastic games and Markov equilibria

Blockchain state evolves over time and decisions depend on state. Forks, difficulty adjustments, mempool congestion, and network conditions are state variables. A stochastic game captures state dependent actions with probabilistic transitions, and Markov style equilibria capture best responses that depend on current state rather than full history.

3.4 Mean field games and mean field type games

When the number of participants is huge, it can be useful to model each participant as negligible in isolation but influential in aggregate. Mean field methods formalize this idea and connect individual best responses to aggregate equilibrium quantities.

Two objects often show up:

- **Mean field game:** each agent interacts with the aggregate through a mean field term such as total hashrate, average congestion, or population distribution.
- **Mean field type game:** each agent payoff can depend nonlinearly on the population distribution, for example through variance aware terms.

4 Formal setup and equilibrium concepts

Let there be N strategic players indexed by $i \in \{1, \dots, N\}$. Each player chooses an action $a_i \in \mathcal{A}_i$. Let $a = (a_1, \dots, a_N)$ and a_{-i} denote the action profile of all players other than i . A payoff function is

$$u_i : \mathcal{A}_1 \times \dots \times \mathcal{A}_N \rightarrow \mathbb{R}.$$

A Nash equilibrium a^* satisfies

$$u_i(a_i^*, a_{-i}^*) \geq u_i(a_i, a_{-i}^*) \quad \text{for all } a_i \in \mathcal{A}_i \text{ and all } i.$$

For blockchain, the core modeling step is identifying:

- **Players:** miners, validators, pools, users, builders, relayers, governance voters.
- **Actions:** resource allocation, timing, message propagation, chain choice, pooling, voting.
- **Payoffs:** rewards, fees, penalties, hardware cost, energy cost, risk, and external utility.

5 Consensus and incentives as mechanism design

A consensus protocol defines the admissible state transitions and who is eligible to propose them. Incentives then determine why a participant should contribute to consensus.

A practical security principle is **incentive compatibility**: deviations should not improve expected payoff after accounting for costs and risk. This is not a moral claim. It is a statement about best responses.

Incentive compatibility can fail in subtle ways:

- A deviation may increase expected reward share by affecting competitors.
- A deviation may exploit information timing, such as withholding or delayed release.
- A deviation may rely on coordination among a subset of players.

When incentive compatibility fails, deviations become rational strategies that can persist at equilibrium.

6 Security as equilibrium behavior

Security questions can be split into two categories:

- **Cryptographic security:** can an adversary forge signatures or violate hash properties.
- **Economic security:** can an adversary profit from strategic deviation given incentives.

The second category is game theoretic. Many real attacks are not about breaking cryptography, but about exploiting payoff structures.

6.1 Selfish mining and withholding intuition

A canonical strategic deviation in Proof of Work is selfish mining: a miner or pool may withhold newly found blocks and release them strategically to increase relative success rates, forcing honest miners to waste computation.

One important pattern is pool infiltration and block withholding: an attacking pool allocates some mining power inside another pool and withholds valid blocks found there. This changes reward distribution and creates a competitive game between pools.

6.2 Propagation incentives and message timing

Even if mining itself is incentivized, propagation is a separate action. If transaction fees matter, miners may have incentives to delay information flow to increase expected profit. A protocol can introduce explicit propagation rewards, but then participants may attempt to game the reward mechanism through fake identities or strategic relaying choices. This becomes a strategic interaction best modeled as a non cooperative game.

6.3 Fee driven forking and undercutting intuition

If fees dominate rewards, miners may prefer forks that maximize fee capture rather than simply extending the longest chain. One can view this as a repeated interaction: miners repeatedly choose whether to mine honestly or to attempt a fee driven fork strategy. Under learning dynamics, the network can converge to equilibria that are harmful, such as persistent unclaimed transactions or unstable fork competition.

7 Proof of Work mining as a mean field game

This section provides a quantitative intuition for why equilibrium hashrate is a natural security proxy.

7.1 Mining as an aggregate contest

Mining is a competition for a reward stream. Each miner chooses a hashrate level, increasing their share of rewards while increasing costs and reducing the share available to others. A central security insight is that higher total hashrate makes majority control more expensive, and therefore raises the cost of attacks that require controlling a large fraction of mining power.

7.2 Continuum model and objective

Consider a continuum of miners indexed by $i \in [0, 1]$. Miner i chooses hashrate $h_i(t) \geq 0$. Define aggregate hashrate

$$H(t) = \int_0^1 h_i(t) di.$$

A stylized reward share rule is proportional allocation. If $R(t)$ is the total reward flow in native units and $P(t)$ is an exchange rate into a numeraire, then miner i receives expected flow

$$\frac{h_i(t)}{H(t)} R(t) P(t).$$

Let $c_i(t)$ be a marginal cost per unit hashrate. A basic objective is

$$\max_{h_i(\cdot)} \mathbb{E} \left[\int_0^T \left(\frac{h_i(t)}{H(t)} R(t) P(t) - c_i(t) h_i(t) \right) dt \right].$$

An equilibrium is a fixed point where each h_i is optimal given the induced aggregate H .

7.3 Real hashrate, technological progress, and stationary security

A key modeling distinction is between nominal hashrate, which is what one observes, and real hashrate, which can be interpreted as the number of effective machines needed to reproduce the nominal hashrate under current technology. Under technological progress, nominal hashrate can rise even if the real economic cost of reproducing it remains constant.

An important long run result of this class of models is the existence of a stationary state for real hashrate and convergence toward it under broad conditions. The interpretation is that equilibrium security against external majority attacks can be constant in the long run in deterministic settings, because real resource commitment stabilizes even as nominal measurements trend upward.

7.4 Comparative statics and security margin

In this view, a protocol security margin is an economic quantity linked to the cost of reproducing a large fraction of equilibrium hashrate. If the equilibrium real hashrate is stable, then the attack cost barrier is stable. In stochastic settings, equilibrium security can become positively related to demand because reward value variation affects optimal investment incentives.

8 Token economics as a mean field type game

Consensus is only one game. Tokens introduce a second strategic environment: adoption, holding, and participation.

8.1 Variance aware utility and distribution dependence

A variance aware approach models agents that trade off expected payoff and risk. Let X_i be the state of agent i , and let μ be the population distribution of states. A generic variance aware utility can be expressed as

$$J_i(a_i; \mu) = \mathbb{E} [V(X_i, a_i; \mu)] - \lambda_i \text{Var} (V(X_i, a_i; \mu)),$$

where $\lambda_i \geq 0$ captures risk sensitivity and μ enters because network value depends on adoption, security, and aggregate behavior. The presence of variance terms makes the payoff depend on the distribution in a nonlinear way, which is characteristic of mean field type games.

8.2 Qualitative equilibrium implications

A useful qualitative prediction from this family of models is that an agent optimal token position tends to increase with participation and productivity, and decrease with attack probability and token price, with additional dependence on the dynamics of token price. This formalizes a widely observed narrative: perceived security and adoption jointly shape token demand and sustainability.

8.3 Equilibrium price and token velocity

In equilibrium, demand and supply meet. A model can yield a token price expression as a function of population size, token user base, and supply. This enables mechanism design insights: parameters that increase velocity can limit value appreciation under scarce supply objectives, so token velocity enters as a design constraint rather than as an afterthought.

9 Design checklist for protocol builders

A protocol is a game. A practical evaluation can follow a disciplined checklist.

9.1 Step one: define players

List who can change outcomes, including coalitions and intermediaries such as pools and builders.

9.2 Step two: enumerate actions including deviations

Do not only list intended behavior. Include feasible deviations such as withholding, delayed release, duplication for rewards, fee driven forks, and strategic pooling.

9.3 Step three: define payoffs including costs and risk

Reward flow alone is not enough. Include energy, capital lockup, opportunity cost, risk of variance, and market exposure.

9.4 Step four: solve or approximate equilibrium

Identify whether honest participation is a best response. If not, locate the deviation region and modify payoffs, information structure, or constraints.

9.5 Step five: connect equilibrium to security margin

For Proof of Work, the equilibrium hashrate and its real cost interpretation links directly to the economic barrier to majority attacks. For token economies, equilibrium adoption and demand link to sustainability and the ability of incentives to fund security.

10 Open research directions

- **Heterogeneous utilities:** participants may value censorship power, ideology, or external hedges.
- **Coordination and coalitions:** equilibrium changes when coordination is feasible.
- **Latency and information:** network layer affects feasible strategies and payoffs.
- **Joint modeling:** security, adoption, and price form a coupled equilibrium system.

11 Conclusion

Blockchains survive in adversarial environments by aligning incentives so that protocol compliance is a best response. Game theory provides the right primitives to define players and deviations, and equilibrium analysis explains why some attacks persist without breaking cryptography. Mean field methods then connect individual optimization to aggregate security quantities such as equilibrium hashrate, and mean field type models connect adoption, risk, and token price through distribution dependent payoffs. The practical message is that robust protocols require designing equilibria, not only writing code.

References

- [1] Ziyao Liu, Nguyen Cong Luong, Wenbo Wang, Dusit Niyato, Ping Wang, Ying Chang Liang, and Dong In Kim. A Survey on Applications of Game Theory in Blockchain. arXiv:1902.10865, 2019.
- [2] Julian Barreiro Gomez and Hamidou Tembine. Blockchain Token Economics: A Mean Field Type Game Perspective. IEEE Access, 2019. DOI: 10.1109/ACCESS.2019.2917517.
- [3] Charles Bertucci, Louis Bertucci, Jean Michel Lasry, and Pierre Louis Lions. Mean Field Game Approach to Bitcoin Mining. arXiv:2004.08167, 2020.
- [4] Satoshi Nakamoto. Bitcoin: A peer to peer electronic cash system. 2008.