The Blockchain Folk Theorem

by B. Biais, C. Bisière, M. Bouvard and C. Casamatta

Discussion by Emiliano S. Pagnotta

Imperial College London

Cambridge U., September 14, 2018

Outlook

1. Why blockchains?
    1. Private? “security” and “cost efficiency”
    2. Public? Permissionless access and censorship resistance
2. Decentralized Consensus: Proof-of-Work competition
3. Potential issues: strategic analysis of them is this paper’s main contribution
4. Model and Results in perspective

Blockchain's Value: Why bother?

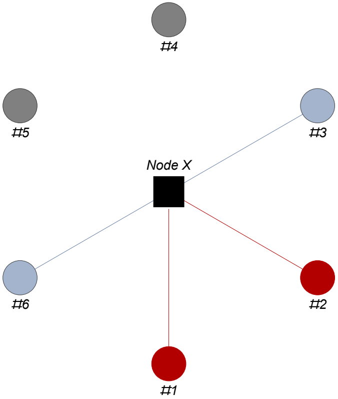

Blockchain's Value: Why bother?

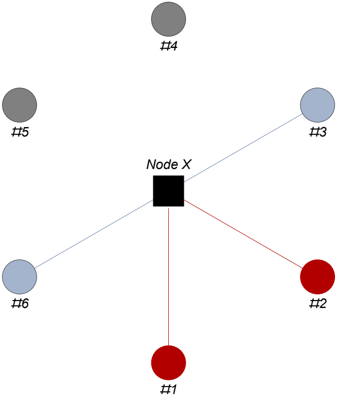

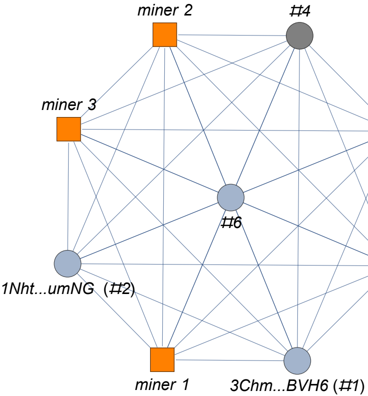

Proof of Work Competition

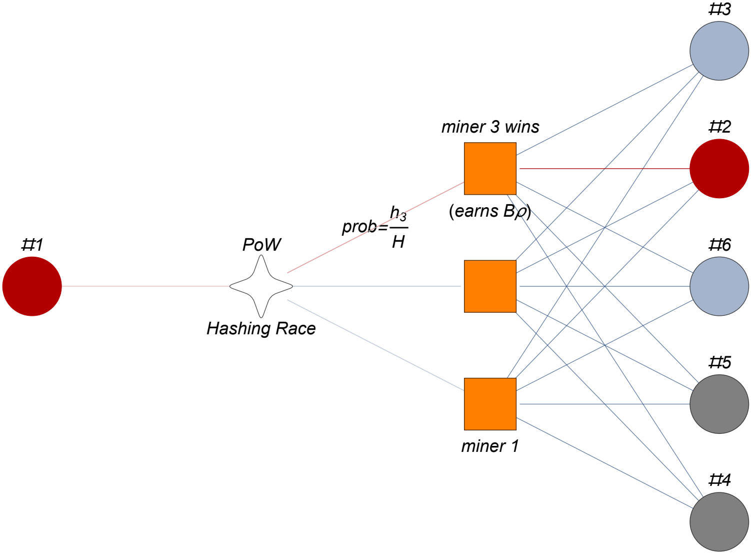

Proof of Work Competition

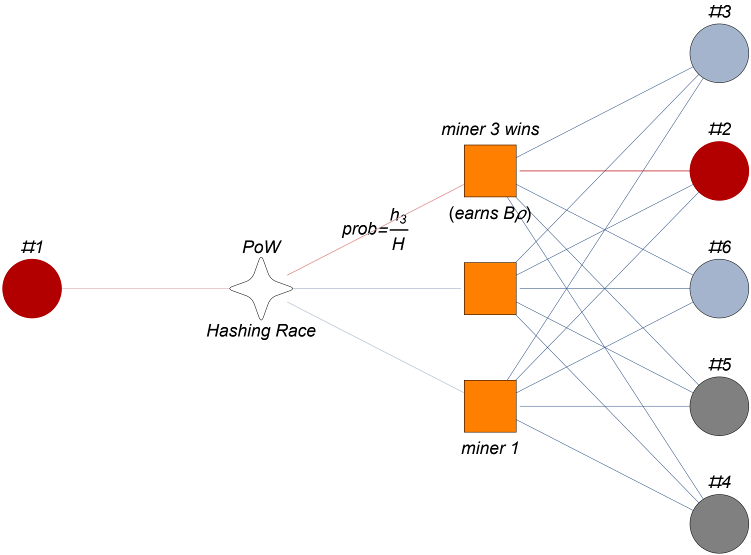

- Since Jan 2009, 541,000 blocks have been mined uninterruptedly

How does Bitcoin's Blockchain work? (1)

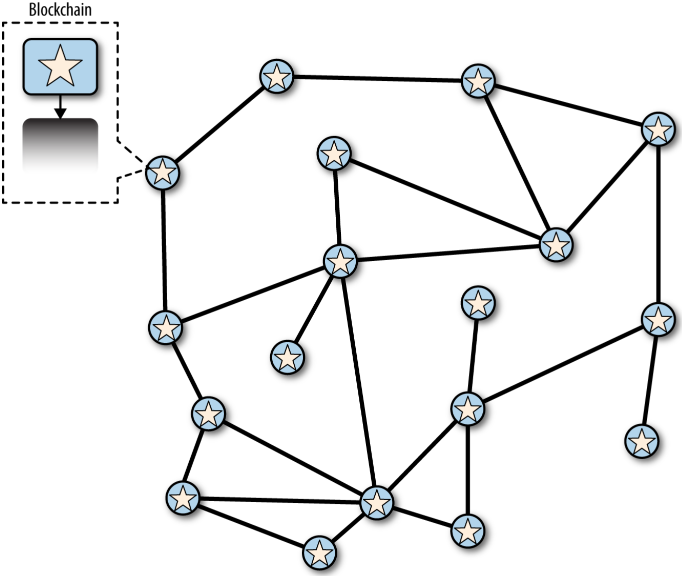

How does Bitcoin's Blockchain work? (2)

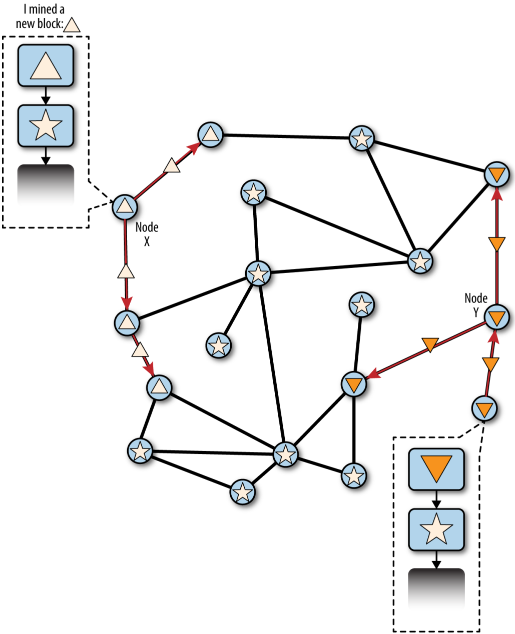

How does Bitcoin's Blockchain work? (3)

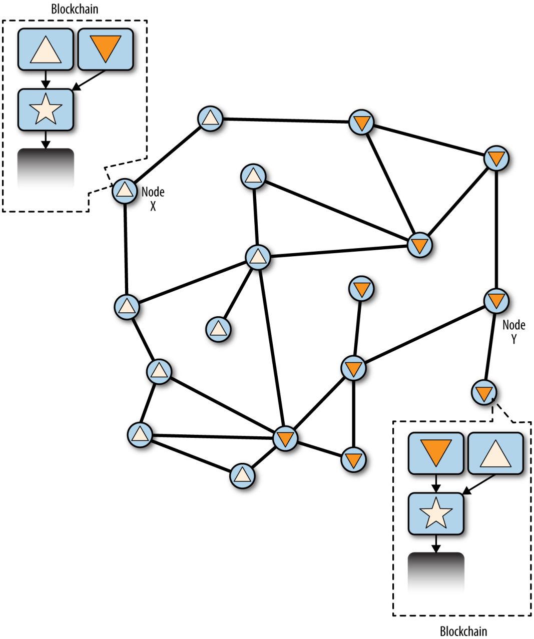

How does Bitcoin's Blockchain work? (4)

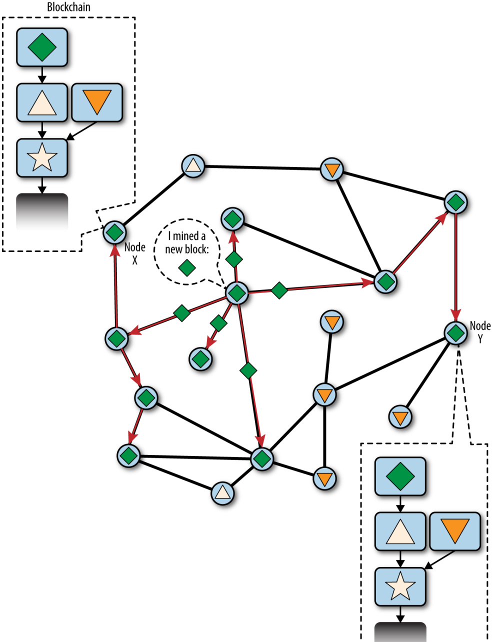

How does Bitcoin's Blockchain work? (5)

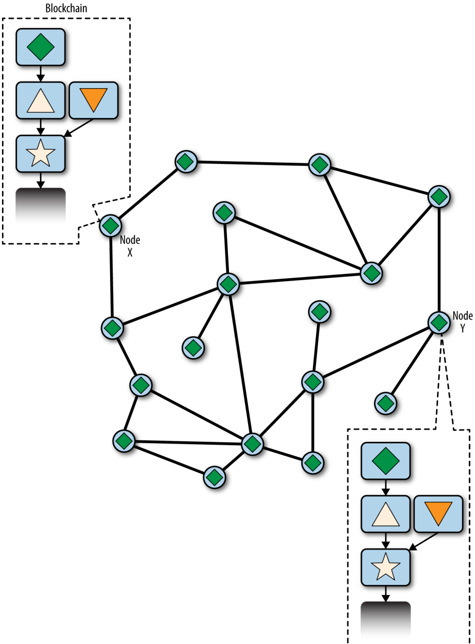

Main Results in BBBC: Equilibrium Character of the Longest Chain Rule

- **Proposition 1**: The LCR is a Markov Perfect Equilibrium of the miners' game
- **Proposition 4**: Robust to realistic informational delays (multiple equilibria is possible)

Main Results in BBBC: Possibility of Persistent Forks

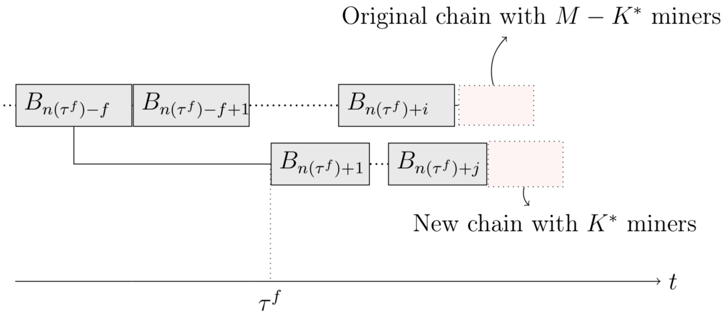

- **Proposition 3**: Persistent forks can occur if, using sunspots to coordinate, a majority of miners decide to fork and this is expected by all.
  - Those forking benefit from a larger expected value of the reward $G(K^*) > G(M - K^*)$
  - Those remaining benefit from continue cumulating vest interest on the original chain

Evidence: March 2014–December 2017

- Average number of blocks per day 144
- Blocks within March 2014 and December 2017= 218, 477
- Number of orphaned blocks = 935
- Number of one-block forks = 928
- Probability of a one-block fork (within period)= 928/218,477= **0.425%**
- Probability of a two-block fork (within period) = **0.0032%**
- Ex post consensus seems very reliable (like in Propositions 1 and 4)
- Orphaned blocks most likely due to random delays. No well-known instances of attacks like double spending, denial of service (censorship), or 'selfish mining' (Eyan and Sirer, 2014)

Why is Bitcoin consensus so reliable? (1)
Assessing Vested Interest

Quantifying Vested Interest

- Vested interest in rewards is generally low
- Currently not feasible for miners to reverse tens of block. Rewards received for mining on the main chain two days ago (even two hours ago) are almost impossible to be reversed (short reward vested horizon)
- The target block finding size of 10 minutes is much larger than typical information delays of a few seconds
- Even if it were possible, if a given pool were acting in a way that undermines Bitcoin reliability, mining nodes can shift hashrate to a different pool
- If a fork did not resolve quickly, the network is likely to stop transacting until there is clarity 'on the truth.' This typically happens around hard forks

Why is Bitcoin consensus so reliable? (2)
General Equilibrium Effects

General Equilibrium effects

- BBBC: When difficulty adjusts, higher hashrate creates a negative impact on others
- Also: Undermining confidence in the integrity of the ledger will affect the market value of the token
- Even is miner has no vested interest in rewards, vested interest is hardware (ASIC) give strong incentive not to undermine confidence and the market price

Explicit Equilibrium Prices

Technology: Network Security Mapping

- Probability of network survival over next period $\tau(H) = \frac{H}{\phi^{-1}+H}, \phi > 0$
  - ϕ: Price-insensitive factors. Examples: quality of the open source code, number of non-mining full nodes

Technology: Hashrate Cost

- $C(h) = \frac{c}{2}h^2, c > 0$

Explicit Equilibrium Prices

Technology: Network Security Mapping

- Probability of network survival over next period $\tau\left(H\right) = \frac{H}{\phi^{-1}+H}, \phi > 0$
  - ϕ: Price-insensitive factors. Examples: quality of the open source code, number of non-mining full nodes

Technology: Hashrate Cost

- $C\left(h\right) = \frac{c}{2}h^2, c > 0$

$$p_t^B = \left(\frac{\sqrt{B_{t-1}\rho p_t^B \left(\frac{m-1}{c}\right)}}{\phi^{-1} + \sqrt{B_{t-1}\rho p_t^B \left(\frac{m-1}{c}\right)} \left(1 - \delta r^B\right)} \right)^{\frac{1}{\sigma}} \frac{n_t^{\frac{\alpha(1-\sigma)}{\sigma}+1}}{B_{t-1}\left(1+\rho\right)}.$$

Calibration Bitcoin Network: June 30, 2018

Data Points

- Price =6,381 USD
- Hashrate = 35.88 Ex/sec
- Supply: 17.1M. Implies yearly $\rho = 3.8\%$
- Network: 20M users. 10 mining pools

Unobservables

- Prob. of Successful attack: $1 - \tau(35.88) = 0.01$

Table: Parameter Values

	Supply and Mining		Beliefs	Preferences		
Parameter	c	ϕ	r	δ	σ	α
Value	29.7E+6	2.78	1	0.95	0.5	0.18

General Equilibrium Effects: Mining Decentralization and Price Reaction

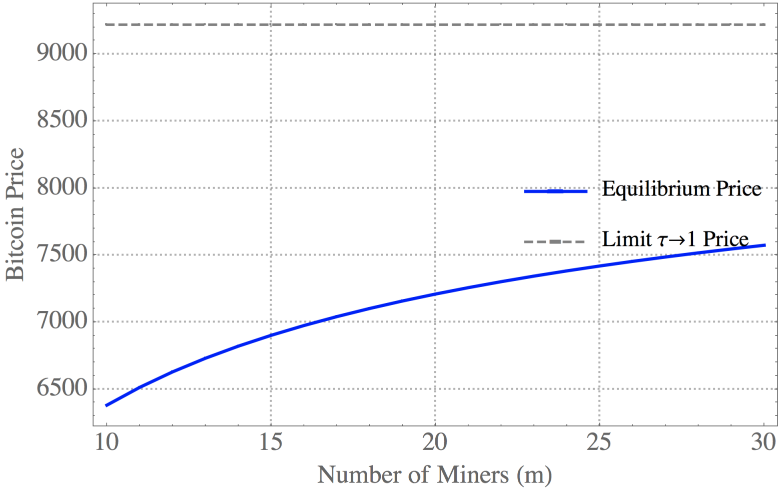

Mining Decentralization and Profits

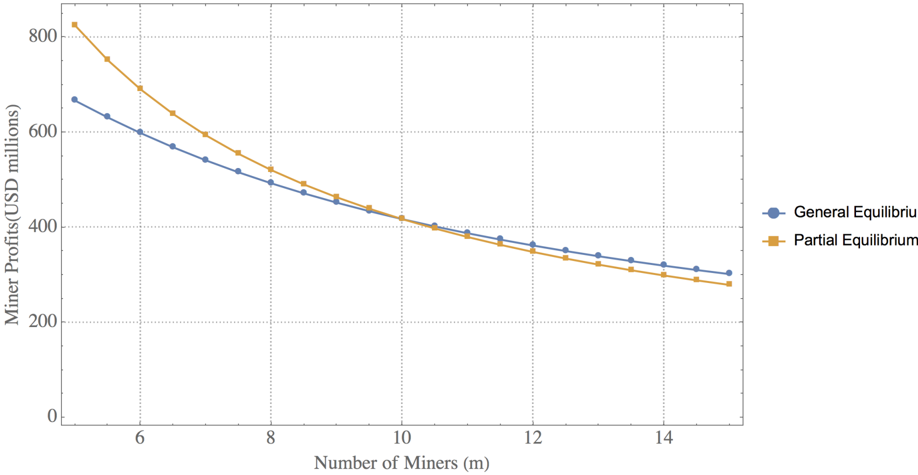

AND THERE CAN BE MULTIPLE EQUILIBRIUM PRICES

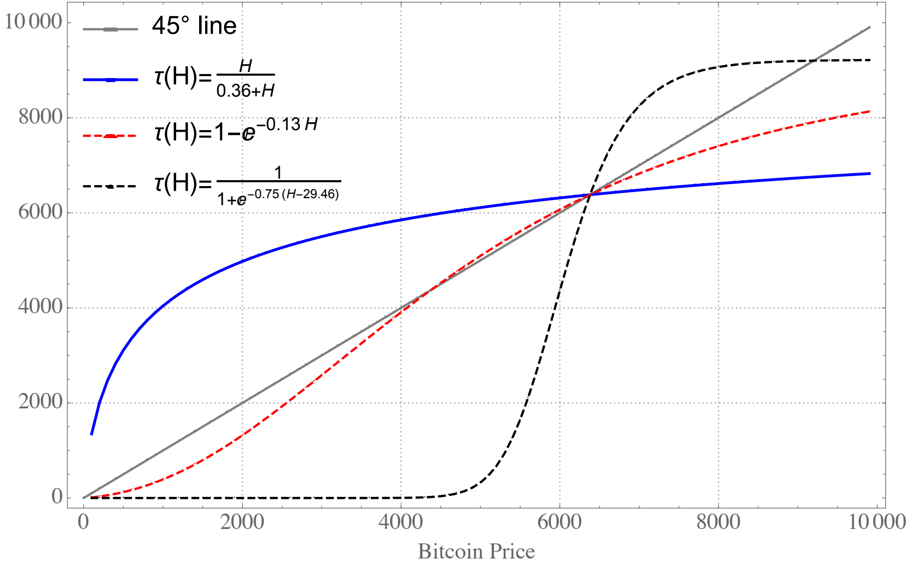

Quantifying the Model (1)

Mining Reward function (*G*)

- Miner *m* computing power is θ_m. Unless $\theta_m = \theta$ for all *m*, it is reasonable to make G a function of hashrate in that chain instead of the number of miners
- Opening the *G* blackbox:
  - With risk-neutral miners, one should have *G* = monetary value of reward times the probability of that reward being confirmed.
  - While one can maintain some exogeneity on the price process, relating the probability of fork survival explicitly to miners hashrate allocation can help assessing the likelihood of each equilibria
  - Important to create bridges between blockchain game theory and investment

Quantifying the Model (2)

Miners life span, risk aversion, and investment

“To capture the 100-block delay before the mining reward can be spent, we assume miner m keeps the units of cryptocurrency he earned until z_m, and then consumes at z_m the rewards earned throughout the game”

- 100 blocks$\approx$16.66 hours. What is then a reasonable z_m? Mining pools are long-lived, so λ_m should be low enough such that $E(z_m)$ is measured in years
- If z_m is expected to be large, the combination of risk neutrality and reward accumulation seems odd in an environment where bitcoin prices are very volatile and electricity is priced in fiat currency.
    - Miners are risk averse, what rationalizes the proliferation of mining pools that allow lowering the variance of the mining reward. Also, typical mining pools must distribute rewards regularly (daily) to members.
- Alternative: $z_m = \infty$, liquidity shock forces to sell stock of mining rewards at unpredictable times (preserving stationarity). Track distribution of hash power as a function of unconfirmed rewards (0, 1, 2 may suffice). Do we get the same outcomes?

EXTENDING THE MODEL

Vested Interest: Additional Dimensions

- Rewards are most likely not hoarded for long, which reduces the chance of persistent strategic forks. BBBC Also analyze private benefits. Intuitive results.
- But vested interest could matter a great deal. Example: Mining hardware
    1. Signaling support for protocol updates (e.g., BIP 9) either soft (SegWit) or hard forks (changes in block size, mining algorithm/rules, etc.)
    2. Economic hashrate wars. Example Bitcoin Cash hard fork
        1. Mining at a loss to encourage use of a chain? Manipulation of hashrate to play difficulty adjustment gains?
        2. Subsidy developers to support a profitable fork
        3. Jihan Wu of Bitmain only accepting Bitcoin Cash for ASIC equipment
    3. Geographical location
        1. Complementarity (lobby power on local authorities)
        2. Substitutability: legal risk (e.g., Chinese miners in 2017)

Extending the Model

Malicious Attackers

- Some attackers may act in a non-profit driven fashion (governments, central banks,...)
- Key: How likely is that the system operates as it should and survives depending on the attacker’s budget and the honest miners budget?
- When is the atomic option superior?

Final Remarks

- Seminal contribution to the understanding of equilibrium consensus in proof-of-work distributed systems.
- Likely to provide a benchmark for the coming years in Econ and computing science
- Analysis suggests that there is much to learn about the long-term reliability of proof-of-work based systems
- Helpful to evaluate plausibility of alternative consensus algorithms (e.g., Cardano's Ouroboros or EOS' Delegated POS)