The Blockchain Folk Theorem

Bruno Biais (TSE), Christophe Bisière (TSE),
Matthieu Bouvard (McGill) and Catherine Casamatta (TSE)

Swissquote Conference 2017 on Fintech
Swiss Finance Institute
Ecole Polytechnique Fédérale de Lausanne

Blockchain

Distributed ledger, records transactions and ownership, operated within a peer to peer network

Bitcoin blockchain: ownership of bitcoins.

Blockchain can be used for other assets & contracts (Ethereum)

If reliable and stable, new cost effective way to record transactions and ownership

Is it?

Nakamoto (2008): Steps in Blockchain

"The steps to run the network are as follows:

1. New transactions are broadcast to all nodes.

2. Each node collects new transactions into a block.

3. Each node works on finding a difficult proof-of-work for its block.

4. When a node finds a proof-of-work, it broadcasts the block to all nodes.

5. Nodes accept the block only if all transactions in it are valid and not already spent.

6. Nodes express their acceptance of the block by working on creating the next block in the chain, using the hash of the accepted block as the previous hash."

Blockchain

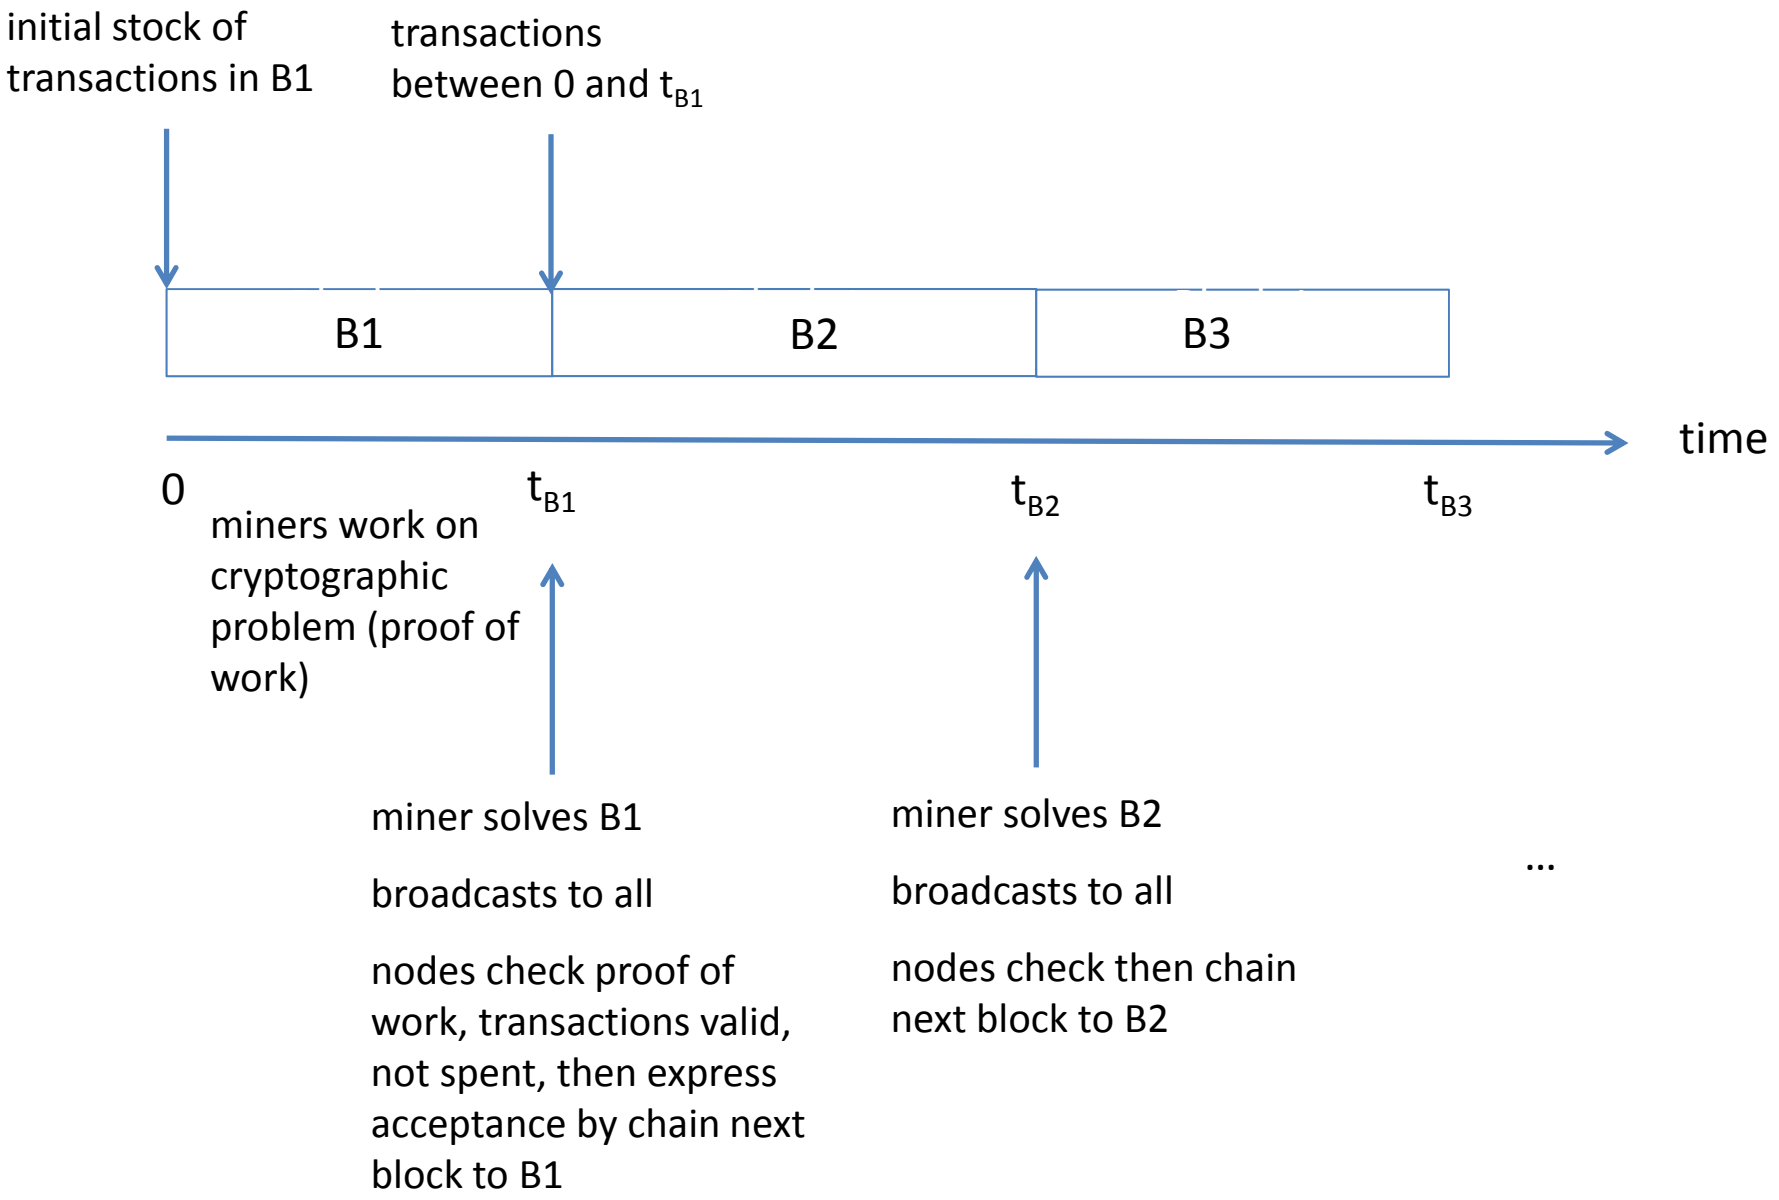
initial stock of transactions in B1
transactions between 0 and t_{B1}
B1
B2
B3
time
0
t_{B1}
t_{B2}
t_{B3}
miners work on cryptographic problem (proof of work)
miner solves B1
broadcasts to all
nodes check proof of work, transactions valid, not spent, then express acceptance by chain next block to B1
miner solves B2
broadcasts to all
nodes check then chain next block to B2
...

Nakamoto (2008): Longest chain rule (LCR)

To which previously solved block will miners chain their block?

> *"Nodes always consider the longest chain to be the correct one and will keep working on extending it."*

If miners follow LCR $\longrightarrow$ no fork, single chain

In practice forks sometimes occur: March 11, 2013, Bitcoin:

> *"Starting from block 225430, the blockchain literally split into two, with one half of the network adding blocks to one version of the chain, and the other half to the other... The split lasted for 24 blocks, or 6 hours."*

Fork

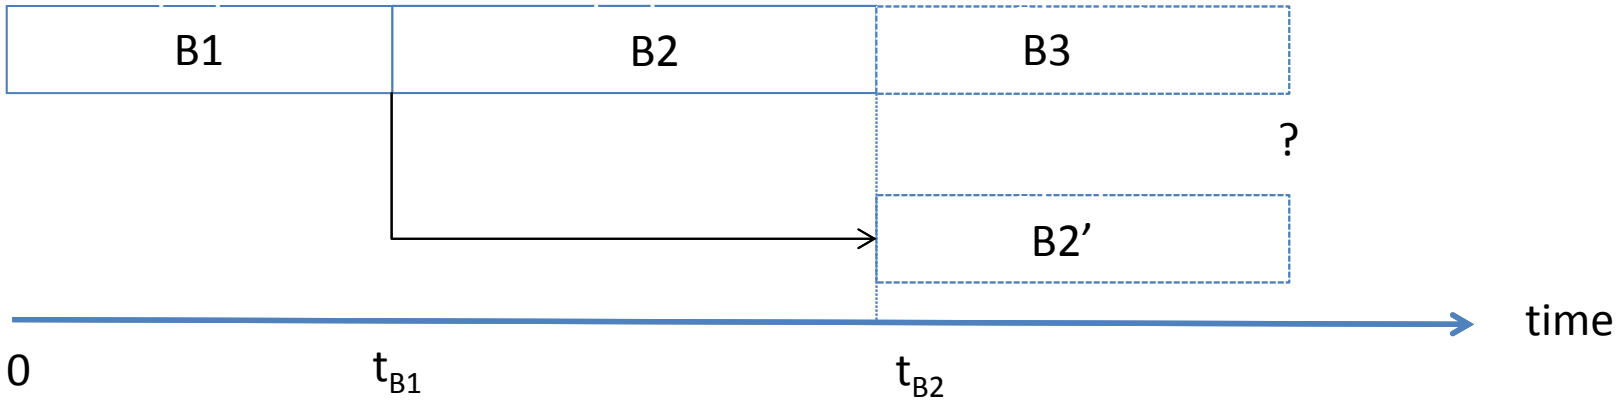

After B2 solved, all should mine B3 chained to B2

Instead some could mine B2’ chained to B1 (deviate from longest chain rule)

Can we expect miners to follow LCR?

We take a game-theoretic approach to study the strategies of miners: to which block do they chain their own block?

We consider a frictionless environment: all miners observe blocks and transactions simultaneously, no double spending

In this perfect environment, can we expect stable and reliable distributed consensus?

Is LCR an equilibrium of the game between miners? Yes

Is it the only equilibrium? No

What are miners mining for?

When miner solves block B_n in Bitcoin blockchain, he broadcasts this, including his reward, in bitcoins, in block B_n

This reward can't be spent immediately, only after sufficiently many blocks chained to B_n (k-blocks rule)

Miners want to

- mine on chain which they think will become consensus, so their rewards are valuable $\longrightarrow$ coordination
- protect value of reward they obtained and still hold $\longrightarrow$ vested interest

Mining is a coordination game

Miners want to attach their blocks to the chain to which they expect the others will attach their blocks

$\Longrightarrow$ Choice of which block to mine = coordination game

$\Longrightarrow$ Multiple equilibria, some without forks and others with forks

Forks $\rightarrow$ orphaned blocks, rewards for orphaned blocks lost, blockchain instability, uncertainty about consensus, undermines credibility/value of cryptocurrency

Coordination issues during March 2013 Bitcoin fork

Two competing chains (0.7 versus 0.8)

Miners didn't know which chain to coordinate on (reflecting uncertainty price of Bitcoin dropped 25% during incident!)

> *"Gavin Andresen: the 0.8 fork is longer, yes? so majority haspower is 0.8 ... first rule of bitcoin: majority haspower wins*
>
> *Luke Dashjr: if we go with 0.8 we are hard forking*
>
> *BTC Guild: I can single handedly put 0.7 back to the majority had power. I just need confirmation that that's what should be done.*
>
> *Pieter Wuille: that is what should be done, but we should have consensus first"*

Eventually, BTC guild chose 0.7, 0.8 orphaned, 24 blocks lost

Mining generates vested interests

Miner keeps rewards for solving block (k-block rule, preferred savings vehicle)

As long as keeps blocks earned on a chain, vested interest in that chain being consensus

$\longrightarrow$ keep mining on that chain to make sure it is consensus

$\longrightarrow$ earn block on that chain

$\longrightarrow$ strengthens vested interest

$\longrightarrow$ persistent competing chains?

Vested interests during March 2013 Bitcoin fork

"Luke Dashjr: it's either lose 6 blocks [mined on 0.8] or hardfork [to 0.8]

Pieter Wuille: all old miners will stick to their old chain regarless of the mining power behind the other

BTC Guild: I've lost so much money in the last 24 hours from 0.8"

"By switching BTC Guild loses the work they've done on 0.8 since teh fork started. On the other hand they are more or less assured that the 0;7 branch will win"
freedom-to-tinker.com/2015/07/28/analyzing-the-2013-bitcoin-fork

Miners and pools

Miners, $m \in \mathcal{M} = \{1, ...M\}$, risk neutral, rational

In practice miners pools coordinate miners $\longrightarrow$ M can be interpreted as number of pools

April 20^{th} 2017: 14 mining pools = 93% of total hash capacity $\longrightarrow$ 15 = reasonable order of magnitude for M

Relatively limited number of pools motivates game theoretic approach (but coordination effects also with competitive)

Hash rate distribution

Source: blockchain.info

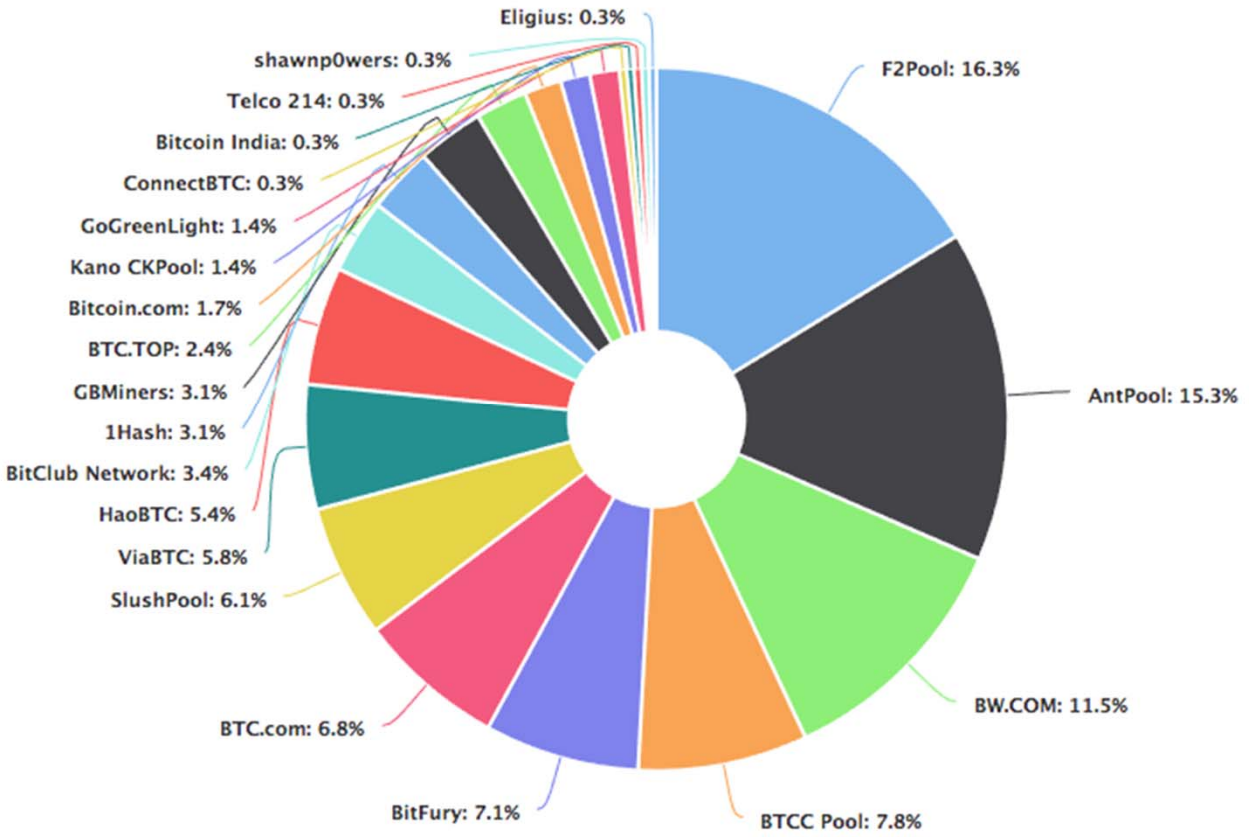

Transactions and blocks

We assume:

- Exogenous continuous flow of transactions sent for confirmation by end-users
- Instantaneously observed by miners (removes one of the potential causes for forks noted by Nakamoto (2008) and yet forks arise in equilibrium)

Exogenous initial state of ledger at time 0 = block B_0

Starting from B_0 participants start mining first block B_1

As time goes by, miners observe the set of solved blocks

At any time miners decide to which previously solved block they want to chain the block they currently mine (action space = set of previously solved blocks.)

Difficulty and speed of solving blocks

Time it takes miner to solve his block $\leftarrow$ difficulty of cryptographic problem and computing power

Difficulty set by protocol s.t. average duration between 2 blocks =10 minutes on Blockchain, around 20 seconds on ETH

If total computing power increases (new miners, new pools), difficulty scaled up to keep duration between blocks constant

Bitcoin: revision in difficulty = every 2,016 blocks, i.e., every 2 weeks

First consider a stationary environment: appropriate for short term horizon ($\leq$ 2 weeks)

Exponential interarrival time

Nakamoto (2008): time it takes m to solve block = exponential

$$\Pr(m \text{ solves block between } t \text{ and } t + dt) = \theta_m dt$$

independent of

- how long m has been mining the block
- which block m is mining
- which block the others are mining

$\longrightarrow$ Suppose m has been mining block B_n, and another miner solves: duration until next time m solves independent of whether m continues to mine B_n or any other block
$\longrightarrow$ Number blocks solved by m at t = exogenous random variable

$$N_m(t) = \int_{s=0}^{t} dN_m(s)$$

Times at which miner m solves blocks

Times of jumps of Poisson with intensity θ_m

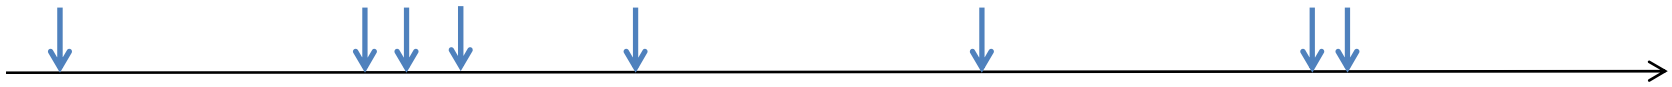

Times at which miner m' solves blocks

Times of jumps of Poisson with intensity $\theta_{m'}$

Times at which miner m” solves blocks

Times of jumps of Poisson with intensity $\theta_{m''}$

Stochastic maturity

At time z_m (exponential with intensity λ) miner m hit by liquidity shock $\longrightarrow$ must consume real goods $\longrightarrow$ sells cryptocurrency earned as reward to a new miner, who also inherits his beliefs

Exit compensated by entry $\longrightarrow$ stationary environment

Before z_m miner m keeps cryptocurrency obtained as reward (k-blocks rule)

Rewards for mining blocks

In block he mines, m includes his reward: G units of cryptocurrency (eg Bitcoins or ETH) (+ transaction fee, much smaller): miners earn *Seigneuriage* on cryptocurrency

Value of reward/value of seigneuriage $\leftarrow$ value of cryptocurrency in that chain $\leftarrow$ consensus on that chain

Value of reward $G(B_n)$ for mining block B_n in chain $\mathcal{B}$ depends on number of miners active in $\mathcal{B}$

Our key assumptions:

- G increasing in number of miners active in $\mathcal{B}$
- $G = 0$ if 0 or only 1 miner in $\mathcal{B}$ (orphaned)

State, strategy and equilibrium

At any time τ, state, ω_τ, includes

- tree of previously solved blocks
- number of miners active on different branches
- public randomisation device

Strategy: maps state ω_τ into action: where, in tree of previously solved blocks, to chain current block

We look for Markov Perfect Equilibria (MPE) (subgame perfect Nash in which action depends only on current state)

Maximum miner's gains

Upper bound on m's lifetime payoff

$$\mathcal{G}_m^{\max} = \left[\int_{s=0}^{z_m} dN_m(s)\right] G(M).$$

- total number of blocks solved by m before z_m: $\int_{t=0}^{z_m} dN_m(t)$, irrespective of mining strategy of m and $-m$
- m cannot earn more than $G(M)$ each time he solves a block.

At t, expectation of $\mathcal{G}_m^{\max}$, *conditional on* $z_m > t$

$$E_t\left[\int_{s=0}^{t} dN_m(s) + \int_{s=t}^{z_m} dN_m(s)\right] G(M).$$

LCR is Nash equilibrium

Proposition 1:
There exists a Markov Perfect Equilibrium in which, on the equilibrium path, there is a single chain and all miners follow the longest chain rule (LCR), thus obtaining the maximum expected payoff $E(\mathcal{G}_m^{\max})$

If m follows LCR, like the others
$\Longrightarrow$ at z_m only one chain (with M active miners)
$\Longrightarrow$ each block mined by m earns $G(M)$
$\Longrightarrow$ m obtains maximum possible gain: $\mathcal{G}_m^{\max}$
$\Longrightarrow$ no deviation yields strictly greater expected payoff
$\Longrightarrow$ optimal (at least weakly) for m to follow LCR

Coordination rather than competition

Miners are not really competing to solve their block before the others

That someone else solves his block before me, does not, in itself, reduce my gains

The only thing that matters is that we all coordinate well, and all work on the same chain, so that we all obtain maximum rewards for the blocks we solve

Proposition 1, entirely driven by coordination effects, does not depend on number of miners, also holds in large number of miners limit

Sunspot equilibrium fork

LCR not single equilibrium. Denote $B_{n(\tau)}$ last block solved by τ

Proposition 2:
There exists a MPE in which at a random time τ (sunspot time) all miners fork, chaining to $B_{n(\tau)-f}$, and following LCR on resulting chain.

Intuition: expect all to fork
$\longrightarrow$ expect only blocks on fork to be rewarded
$\longrightarrow$ also fork

Coordination game - strategic complementarity (again does not depend on number of miners)

Consequence: Fork becomes only active chain, blocks $B_{n(\tau)-f+1}$ to $B_{n(\tau)}$ orphaned, not rewarded $\implies$ fork equilibrium Pareto dominated by LCR equilibrium

Persistent forks

Candidate equilibrium: Some fork at τ^f to new chain whose parent is $B_{n(\tau^f)-f}$, while others remain on original chain

Vested interest on original chain at time τ^f : $v^o(m, \tau^f) =$ number of blocks solved by m on original after $B_{n(\tau^f)-f}$

Rank miners by vested interest in original chain

$$v^o(m+1, \tau) \geq v^o(m, \tau)$$

Proposition 3:
Under some conditions, $\exists$ integer K s.t $\exists$ a MPE in which, at random time τ^f, miners $m \leq K$ (with low $v^o(m, \tau)$) fork to new chain (and hereafter remain there), while miners $m > K$ (with large $v^o(m, \tau)$) forever remain on original chain.

Intuition for Proposition 3

$K \geq \frac{M}{2}$: Persistent forks can occur only if majority of miners fork

Benefit from forking = blocks mined on new chain more valuable (because majority mines it)

Cost of forking = reduces value of blocks already mined on old chain: large if $v^{old}(m, t_{B_n})$ large

Persistent fork equilibrium Pareto dominated

LCR equilibrium: each block rewarded by $G(M)$

Persistent fork equilibrium: blocks solved after $t_{B_{n-f}}$ rewarded by $G(M-K)$ or $G(K) < G(M)$

Information transmission delays

Suppose m does not immediately observe that B_n was solved, and continues to mine from B_{n-1}: If m solves his block before the others, there are now two competing chains of same length

Suppose that from that point on all observe all blocks solved, there are 3 possible equilibria:

- All ignore m and stick to original chain $\longrightarrow$ m's block quickly orphaned, only transient one-block fork
- All focus on m and abandon the original chain $\longrightarrow$ B_n orphaned, only transient one-block fork
- All but m, stick to original chain, while m sticks to his block: If m first to solve, then all chain to this blocks (original chain block B_n orphaned), otherwise all revert to original chain (m's block orphaned)

As in basic model: multiple equilibria. Information delay offers an interpretation for sunspot in Proposition 2.

Endogenous computing power

Mining = just a way to randomise who wins in decentralised manner

Planner would prefer all miners to acquire only very small computing power ε, so that randomisation can be achieved without wasting too much resources (electricity, hardware)

But, if I anticipate others to acquire only ε, I find it optimal to acquire larger computing power, to increase my chances to win

By doing so I exert an negative externality: I increase total computing power, hence difficulty, hence I reduce the frequency with which others solve

In equilibrium: all acquire computing power $>> \varepsilon$: equilibrium not socially optimal, due to negative externalities

Conclusion

LCR equilibrium and single persistent chain with no fork cannot be ruled out... but cannot be taken for granted

Number of miners/computing power (and end users) on a chain $\longrightarrow$ Credibility of chain $\longrightarrow$ Value of rewards for blocks mined on that chain $\longrightarrow$ Attractiveness of that chain

Coordination game $\longrightarrow$ Multiple equilibria

- Instability?
- Pareto dominated (waste of resources)