Sidechain Risks – The Case Against Worrying

Paul Sztorc

Sept 2nd, 2016

Bloq's Paul Sztorc on the 4 Main Benefits of Sidechains

Concerns

0 PM by Kyle Torpey

Although the sidechains concept has received large amounts of interest and praise from the Bitcoin community (both technical and otherwise), doubts still exist. Most notably, Bitcoin Core contributor Peter Todd has voiced his concerns over the security of merge-mined sidechains.

As recently as this month, Todd has said, "[Merge-mined sidechains] are a broken idea that fundamentally means miners have the ability to steal coins, and makes the scaling problem a lot worse."

Concerns

143

Peter Todd explainins why side-chains are insecure and bad for decentralization

(soundcloud.com)

submitted 2 years ago by Chakra_Scientist

75 comments share save hide give gold report

[–] **giszmo** 13 points 2 years ago

I don't agree with this merged-mining == centralization. Sure, pools provide the service of handling merged mining but why would not new tools emerge that are multi-full-nodes?

permalink embed save give gold

[–] **petertodd** 4 points 2 years ago

Bandwidth and disk space aren't free, and it's easy to see how "multi-full-node" tools themselves will be the point of centralization - either you have people vetting the lsit of chains to mine, or the simple cost of researching and installing is your barrier. Those tools also don't solve the problem that unless a merge-mined chain has a majority of hashing power it can be attacked for free - easy to imagine something like Zerocoin running into problems there if governments start trying to ban it.

Ultimately it's really the arguments about blocksize all over again, but with an even higher administration overhead.

permalink embed save parent give gold

[–] **giszmo** 5 points 2 years ago*

As I don't get your full argument, allow me to go step by step: A multi-full-node tool could have check-boxes

Problem?

https://blockstream.com/sidechains.pdf

Automatic Zoom

4.3 Risk of centralisation of mining

An important concern is whether the introduction of sidechains with mining fees places resource pressure on miners, creating Bitcoin centralisation risks.

Because miners receive compensation from the block subsidy and fees of each chain they provide work for, it is in their economic interest to switch between providing DMMSes for different similarly-valued blockchains following changes in difficulty and movements in market value.

One response is that some blockchains have tweaked their blockheader definition such that it includes a part of Bitcoin's DMMS, thus enabling miners to provide a single DMMS that commits to Bitcoin as well as one or more other blockchains — this is called *merged mining*. Since merged mining enables re-use of work for multiple blockchains, miners are able to claim compensation from each blockchain that they provide DMMSes for.

As miners provide work for more blockchains, more resources are needed to track and validate them all. Miners that provide work for a subset of blockchains are compensated less than those which provide work for every possible blockchain. Smaller-scale miners may be unable to afford the full costs to mine every blockchain, and could thus be put at a disadvantage compared to larger, established miners who are able to claim greater compensation from a larger set of blockchains.

What is Drivechain?

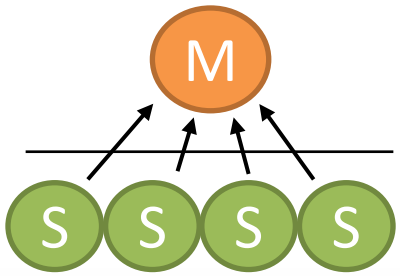

www.truthcoin.info/blog/drivechain/

Drivechain – The Simple Two Way Peg

24 Nov 2015

Main-to-Side

(Much Earlier)

Transaction 1: 96e08e333304dad95d10...

17bAY2JM37he7t.. OP_RETURN 1MFWbAqkcA6Pjz... → 4HftNSm282A3rG..

Transaction 2: ff44d9cdda857a902cf698...

14e4o4b5yibp4w.. OP_RETURN 1QAtLEPrxddW9D... → 4HftNSm282A3rG..

Transaction 3: 653cbe17844eba9df26ad...

1WEZNFsw5trRJU.. OP_RETURN 1teQsVXAXxrCtTT... → 4HftNSm282A3rG..

Three deposits, from Bitcoin (orange) to an address (green) on a Bitcoin sidechain (bold green). The Mainchain BTC are trapped "in 4HftNSm282A3rG.."

Sidechain

(Very Recent Past)

17bAY2JM37he7t.. OP_WITHDRAW 96e08e333304dad95d10... OP_RETURN 1ExTxQbQrdXXJM... → 1BitcoinEater000..

14e4o4b5yibp4w.. OP_WITHDRAW ff44d9cdda857a902cf698... OP_RETURN 1BDRQhGK16GZbs... → 1BitcoinEater000..

1WEZNFsw5trRJU.. OP_WITHDRAW 653cbe17844eba9df26ad... OP_RETURN 1KvQ2xQAiB6W7k... → 1BitcoinEater000..

...dividual withdrawals (WTs), from a Sidechain ...to the Mainchain (orange). Software selects ...txns (blue) to match the desired amounts.

Contracts, Payments, Services

Side-to-Main

(At Present)

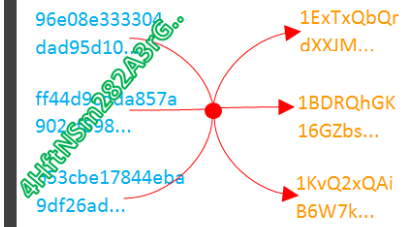

The assembled WT^, whose ID will be included in the Sidechain header for a while. On the Bitcoin Mainchain, it eventually moves BTC "from 4HftNSm282A3rG.." to their new owners.

For simplicity, I assume that all addresses/transactions contain exactly 1 BTC (except for the WT^ which contains 3 BTC).

What is Dr...

04B4697D5319B8B0E461BE624EAD61331
CA613216F061D2533490ABBB71616A0

www.truthcoin.info/blog/drivechain/

...o Way Peg

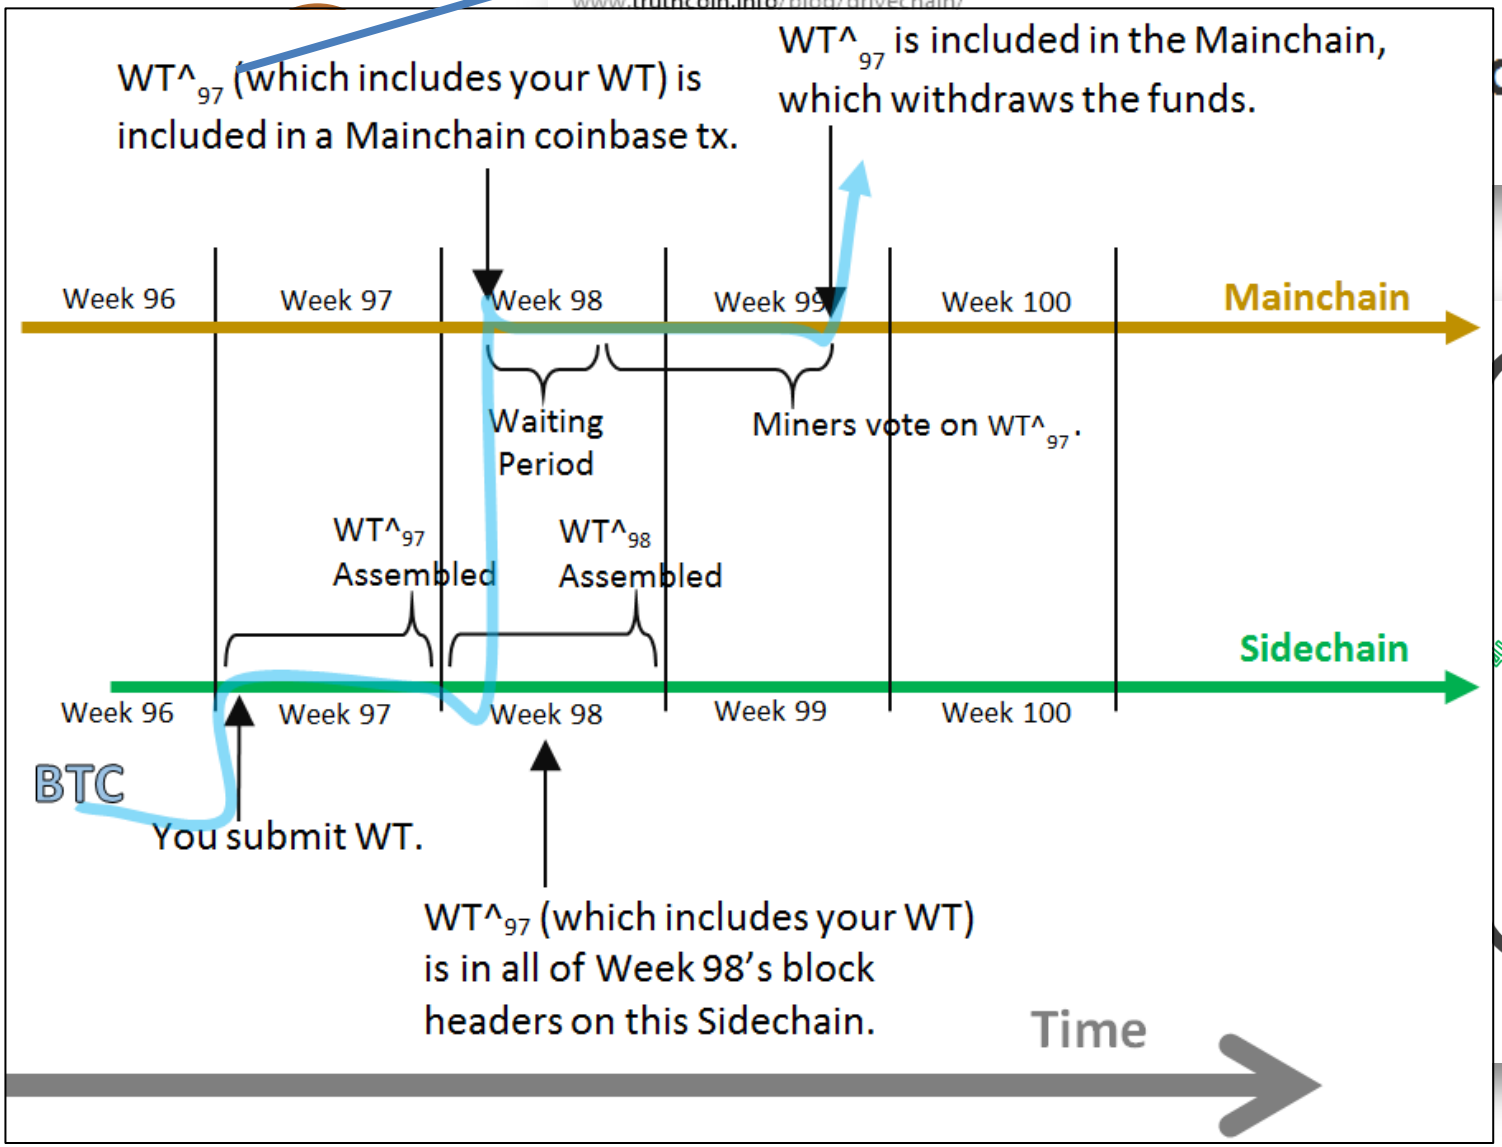

Side-to-Main

(At Present)

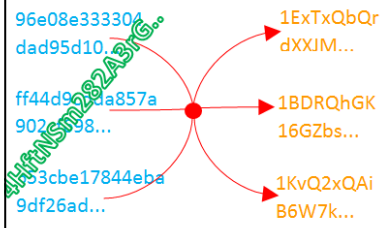

The assembled **WT^**, whose ID will be included in the Sidechain header for a while. On the Bitcoin Mainchain, it eventually moves BTC "from 4HftNSm282A3rG.." to their new owners.

Presentation Overview

Imagine a user who *hates* the idea of sidechains.

1. [*a*] To what extent can Bitcoin sidechains affect the Bitcoin Mainchain? (Limited to Mining)

   [*b*] To what extent might the Mining Network be affected, by sidechains? (Data Transmission Capacity 'Bandwidth', Transaction Fees)

2. Interlude: The Docile Miner
3. Discuss Bandwidth in Detail (Propagation)
4. Discuss Fees in Detail (Competition, Calculus)
5. Light Commentary on Orphaning, & Conclusion

Conclusions -- **Preview**

1. Node Costs
2. Relative to other SF/MM
3. Docile Miners
4. Tame vs. Aggressive
5. Fees and Bandwidth

Just don't run the software.

Node Costs
- Opt-In
- Internalized
- Anti-Fragile

Mining
- "There's only one SH
- How can different ch common network?

namecoin

No one cares!

My Mining Experience?

Interlude: The Docile Miner

I'm still worried

"Tame" Sidechains
1. To Trespass, Need:
1. 51% Miners (else
2. Interested, Fee-Payi
3. Code to have been
- Contracts are firewalled – opt in, and don't affect each other.
- Contracts are managed, to maximi BTC value.

"Aggressive" Sidechains
affected. Bitcoin
other for the dura
3. Fundamental Q:
mining activities?
- Permiss ss Innovation.
- No debate er precise "split" of validation resources (MB, SigOps).

Stop stealing my space!

Instant Block Propagation
High, Stable Demand

- **SPV Mining**
- Bloom Filters
- FIBRE
- UTXO Commitments
- **Block Scheduling***
- Clever Transaction Strategies

A block header contains these fields:

Field	Purpose
Version	Block version number
hashPrevBlock	256-bit hash of the previous block header
hashMerkleRoot	256-bit hash based on all of the transactions
Time	Current timestamp as seconds since 1970-01-0
Bits	Current target in compact format
Nonce	32-bit number (starts at 0)

Bandwidth

Fees

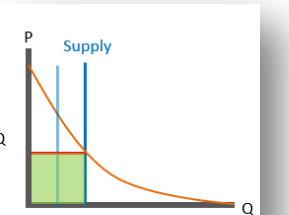

Cartel

Maximized

Part 1 – Blockchain Interactivity

1. How do sidechains affect Bitcoin?
2. How do we reduce / eliminate this?

Agenda

1. The Problem (11)
2. One General Solution (3)
3. Is this GS Robust? (5)
4. Beyond The Limits (5)

The Problem

- **More Stuff = More Resources**

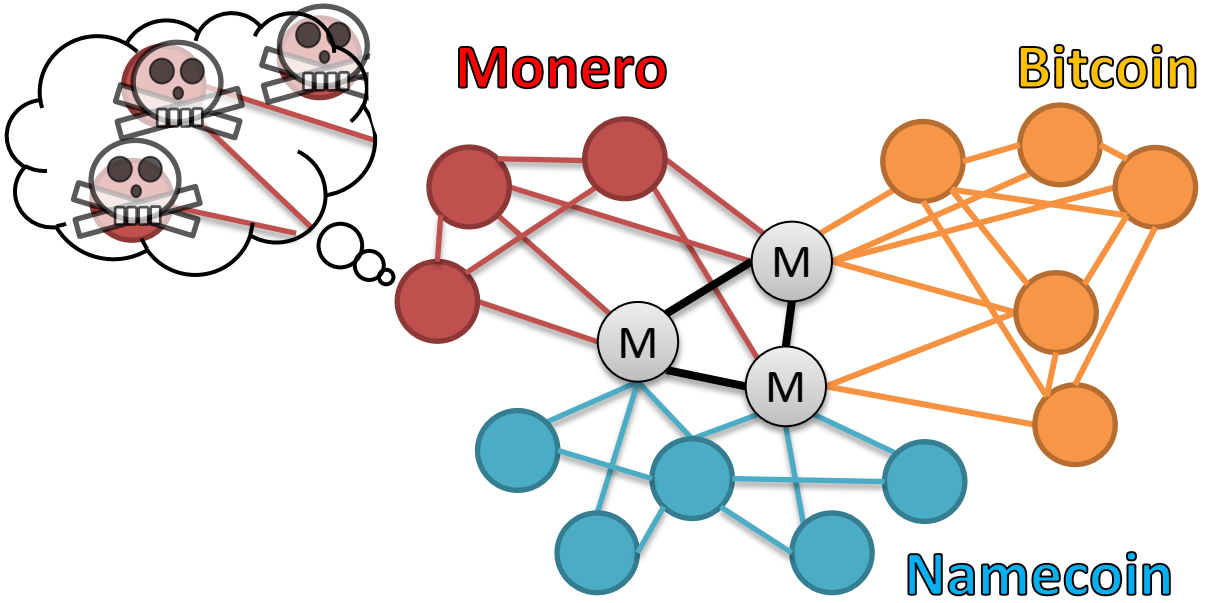

- Node Costs
  - Opt-In
  - Internalized
  - Anti-Fragile
- Mining
  - "There's only one SHA256²( ) network."
  - How can different chains affect each other, *through* the common network?

Damaging the (BTC) Miners

bloq

How are *Bitcoin* blocks found?

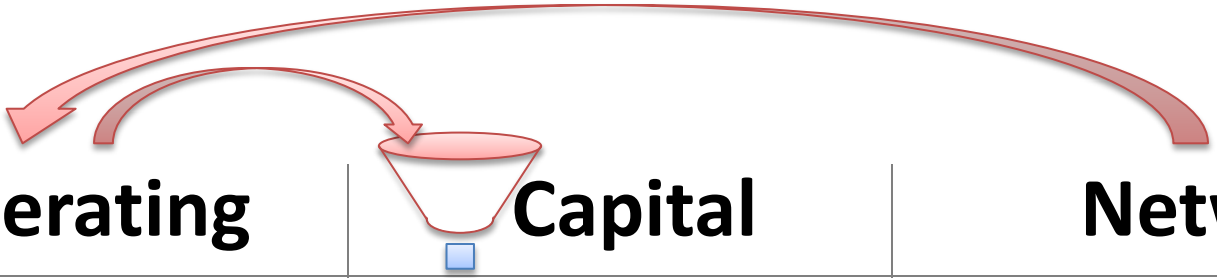

	Operating	Capital	Network
Needs:	Power	Hardware	Current Block (header) Blockreward (tx fees, sub.)
Ideal:	*arg min (Power_Costs) would be Geographically Distributed*	*Equal Access to Best Hardware/Practices*	*Instant Block Propagation* *High, Stable Demand*

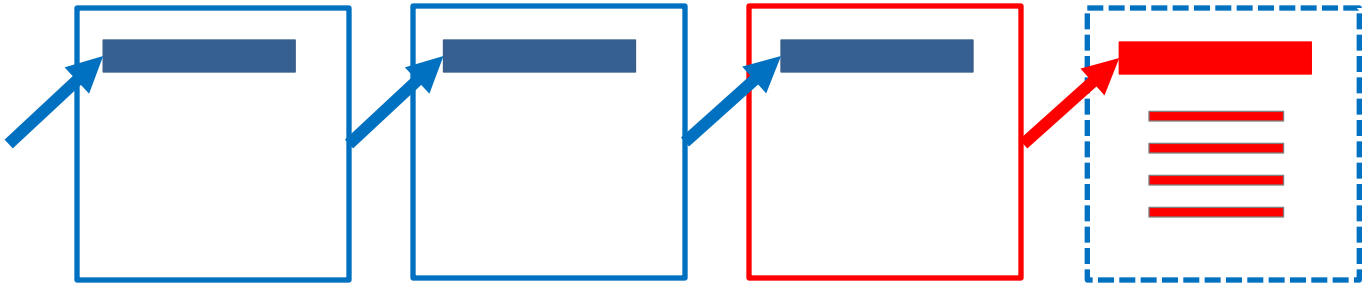

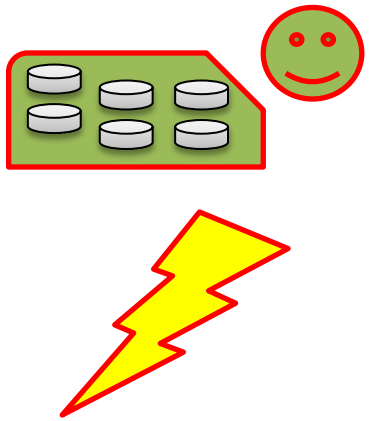

What's Changing?

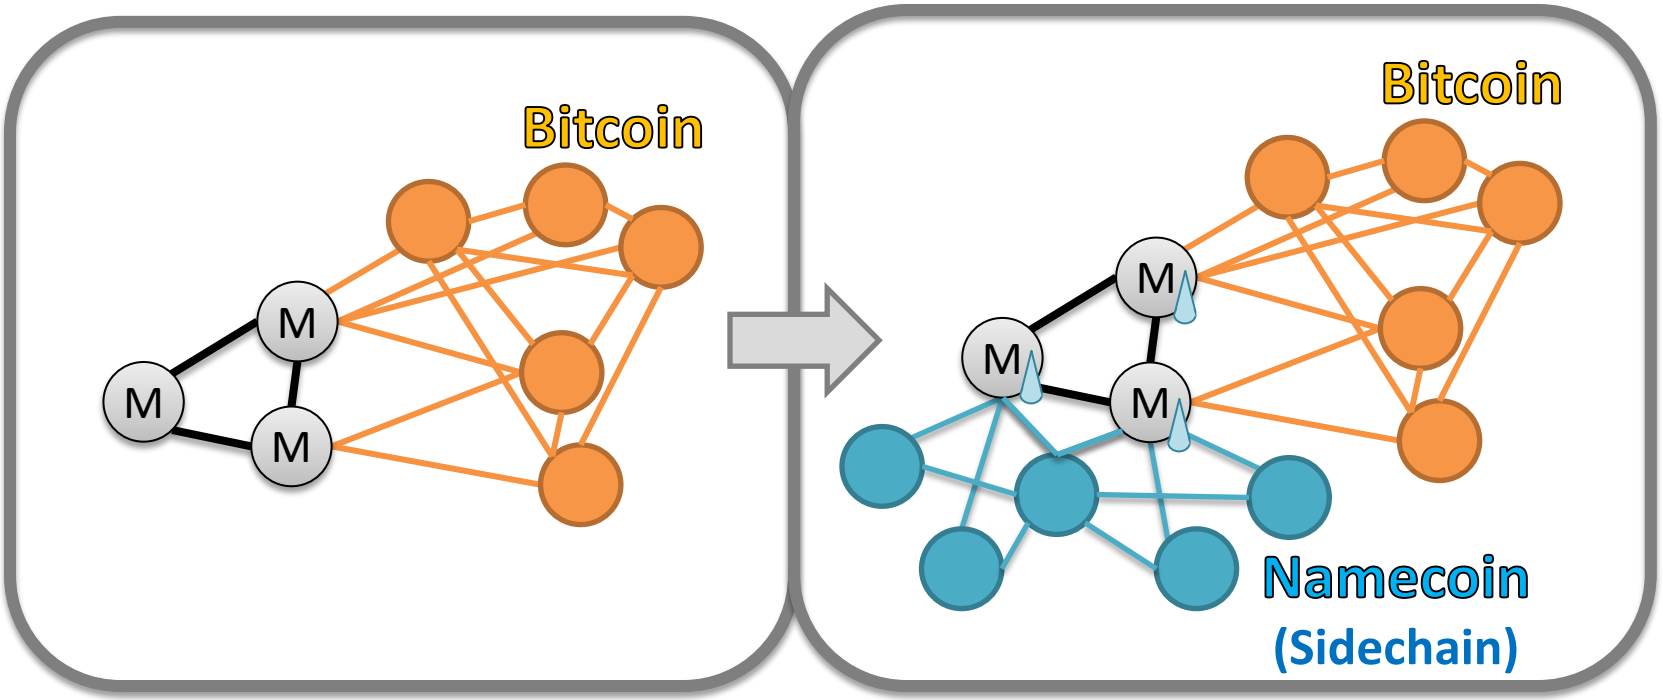

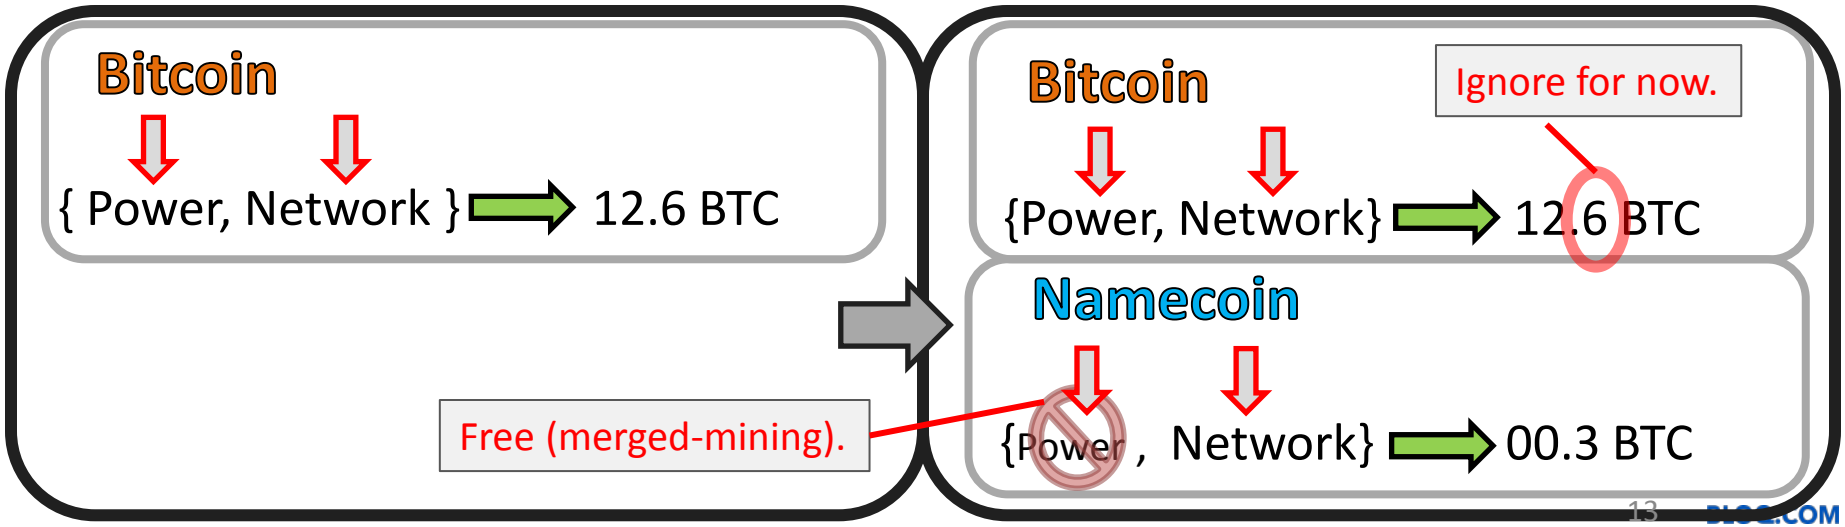

Define Value of Sidechain Option

$$\alpha = E(\text{ New tx fees }) - E(\text{ New Node Costs })$$

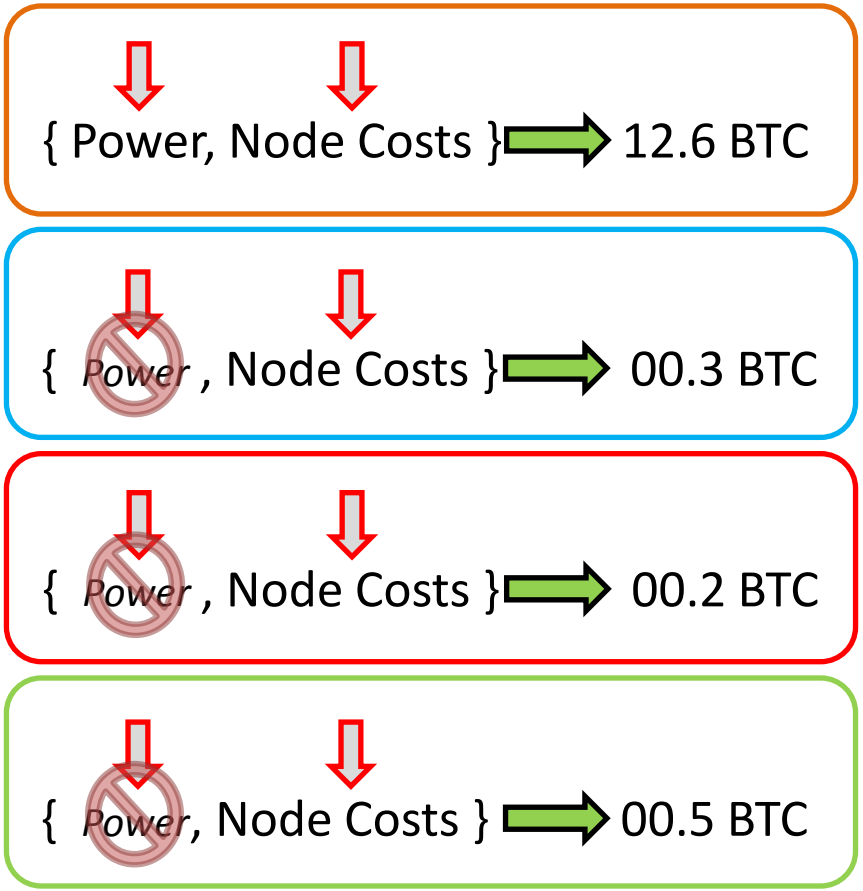

Bitcoin (N/A)

$\alpha_{\text{Namecoin}} = 0.296$ BTC

$\alpha_{\text{Zerocoin}} = 0.197$ BTC

$\alpha_{\text{Rootstock}} = 0.480$ BTC

But is it *really* an "option"?

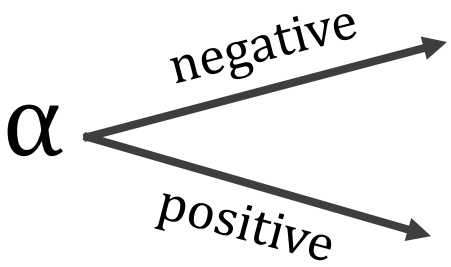

The sidechain is unprofitable, and will be ignored.

The sidechain is profitable...

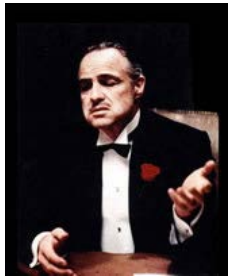

...and *must* be mined.

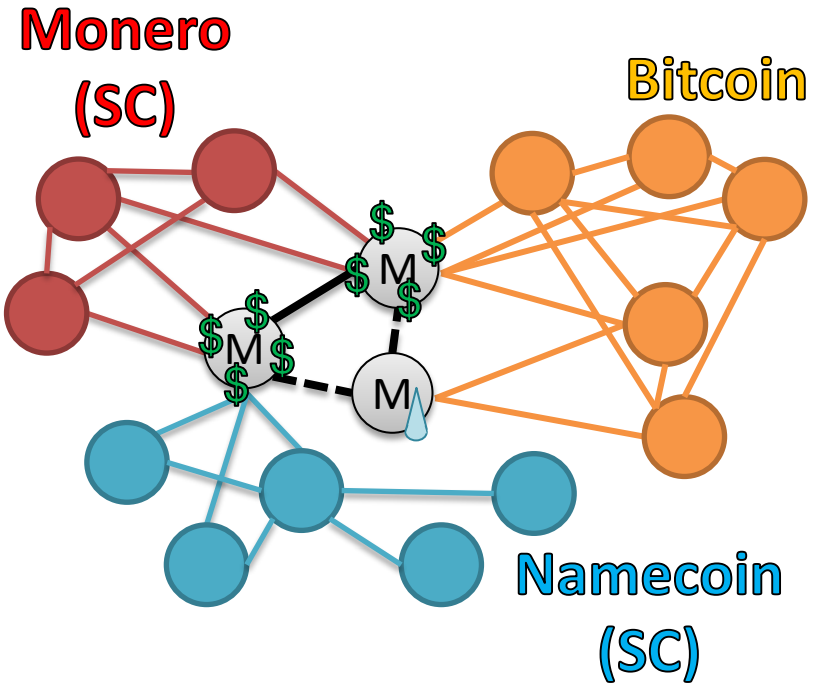

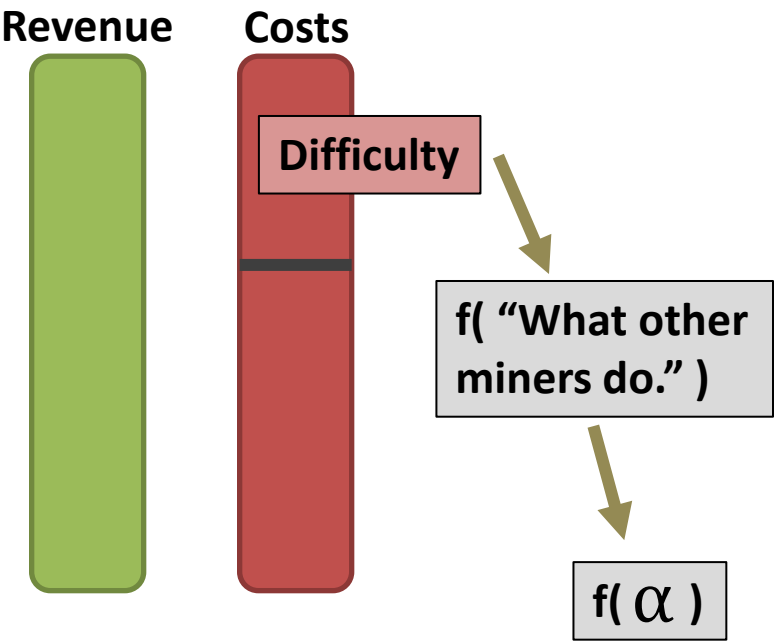

Sidechains: Not An Option

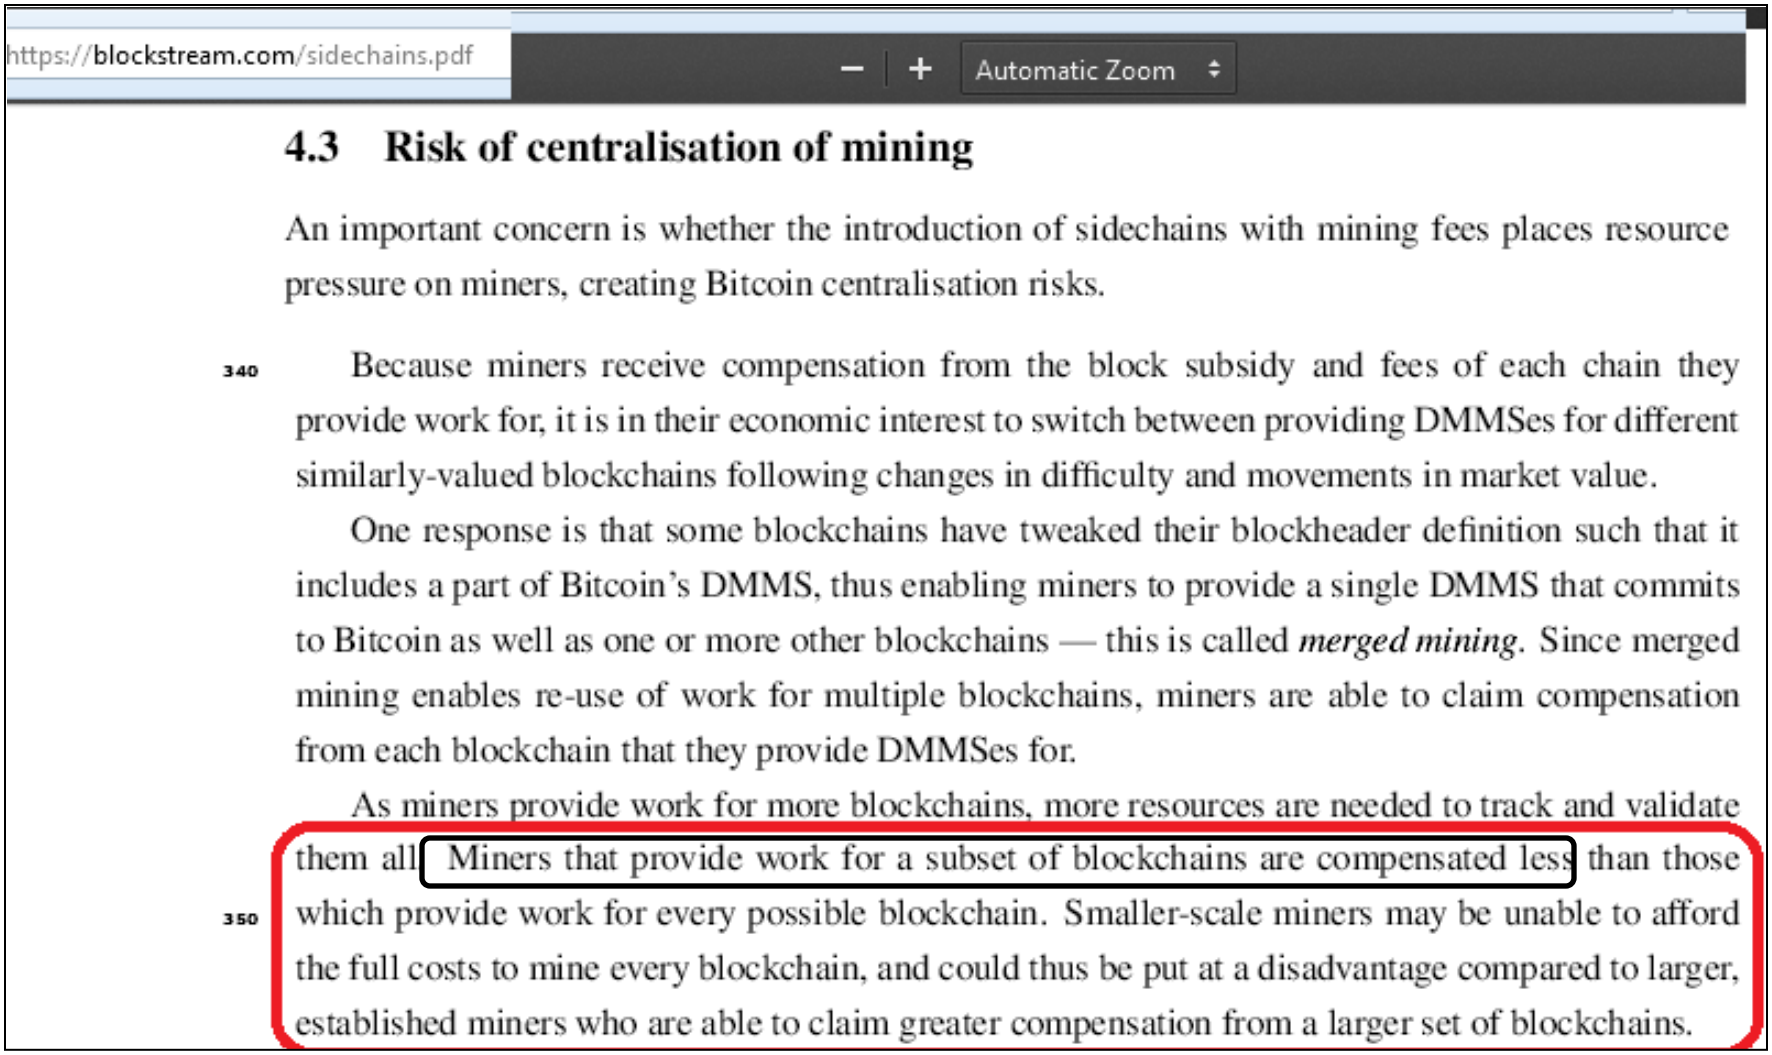
https://blockstream.com/sidechains.pdf

Automatic Zoom

4.3 Risk of centralisation of mining

An important concern is whether the introduction of sidechains with mining fees places resource pressure on miners, creating Bitcoin centralisation risks.

Because miners receive compensation from the block subsidy and fees of each chain they provide work for, it is in their economic interest to switch between providing DMMSes for different similarly-valued blockchains following changes in difficulty and movements in market value.

One response is that some blockchains have tweaked their blockheader definition such that it includes a part of Bitcoin's DMMS, thus enabling miners to provide a single DMMS that commits to Bitcoin as well as one or more other blockchains — this is called *merged mining*. Since merged mining enables re-use of work for multiple blockchains, miners are able to claim compensation from each blockchain that they provide DMMSes for.

As miners provide work for more blockchains, more resources are needed to track and validate them all. Miners that provide work for a subset of blockchains are compensated less than those which provide work for every possible blockchain. Smaller-scale miners may be unable to afford the full costs to mine every blockchain, and could thus be put at a disadvantage compared to larger, established miners who are able to claim greater compensation from a larger set of blockchains.

Miners Must Run All Viable Sidechains

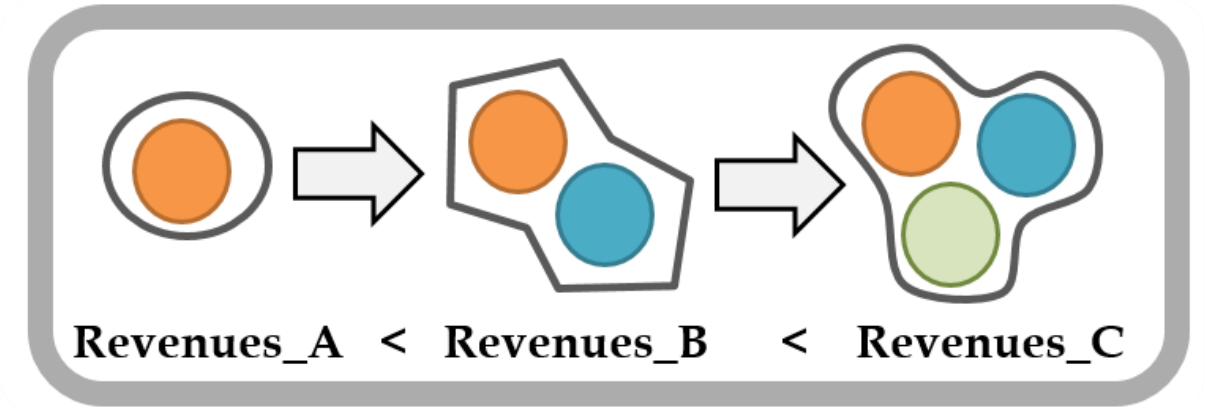

Sidechains are affecting ***the Bitcoin miners***.

1. What are the effects of this?
2. Can we minimize these effects?
3. What recourse do we have, if something bad happens?

Recall: Drivechain Philosophy

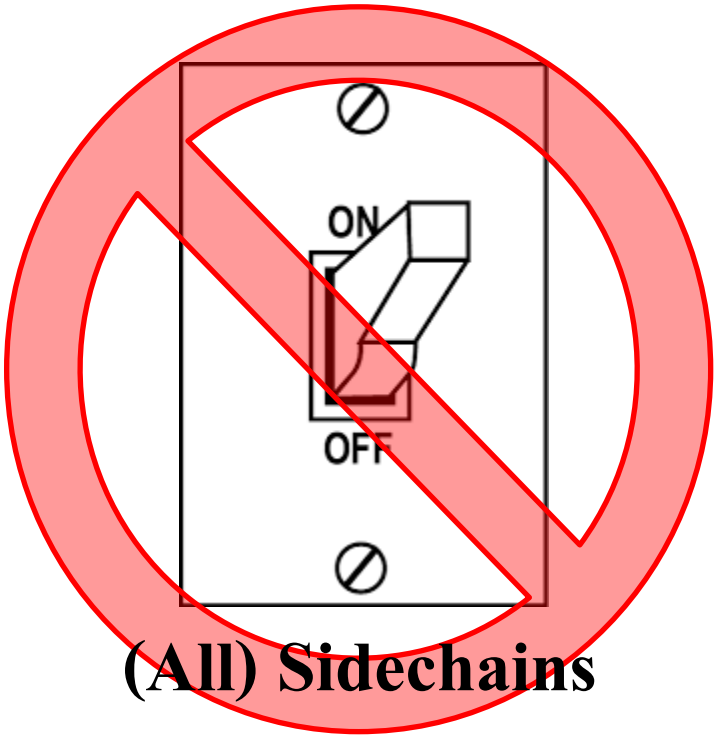

(All) Sidechains

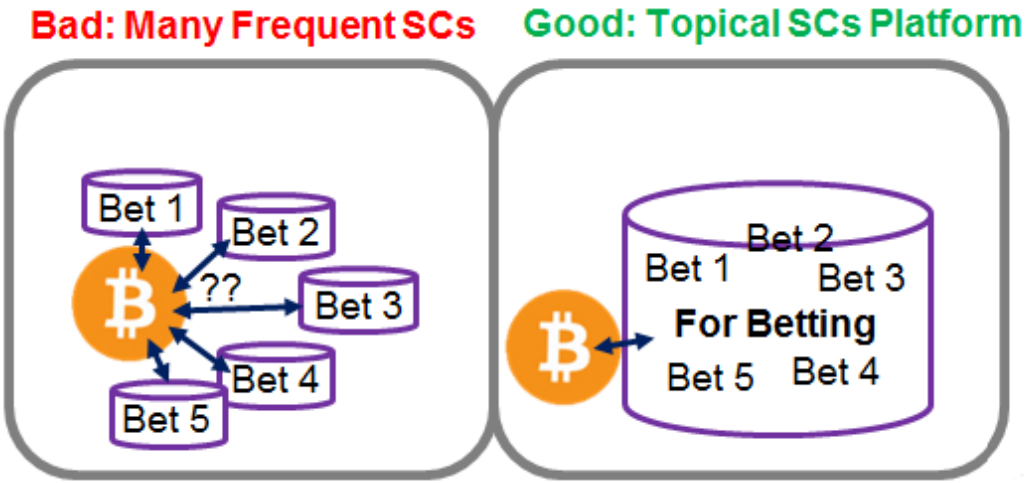

- Soft fork to add each *optional* SC.
- SCs are Rare, Topical, ***Profitable***, Slow, ***Deliberate***.

(Trivial) Protection: No SCs added, until market is okay with it.

Interactivity (Previous presentation.)

Restatement – Internalize the Externalities

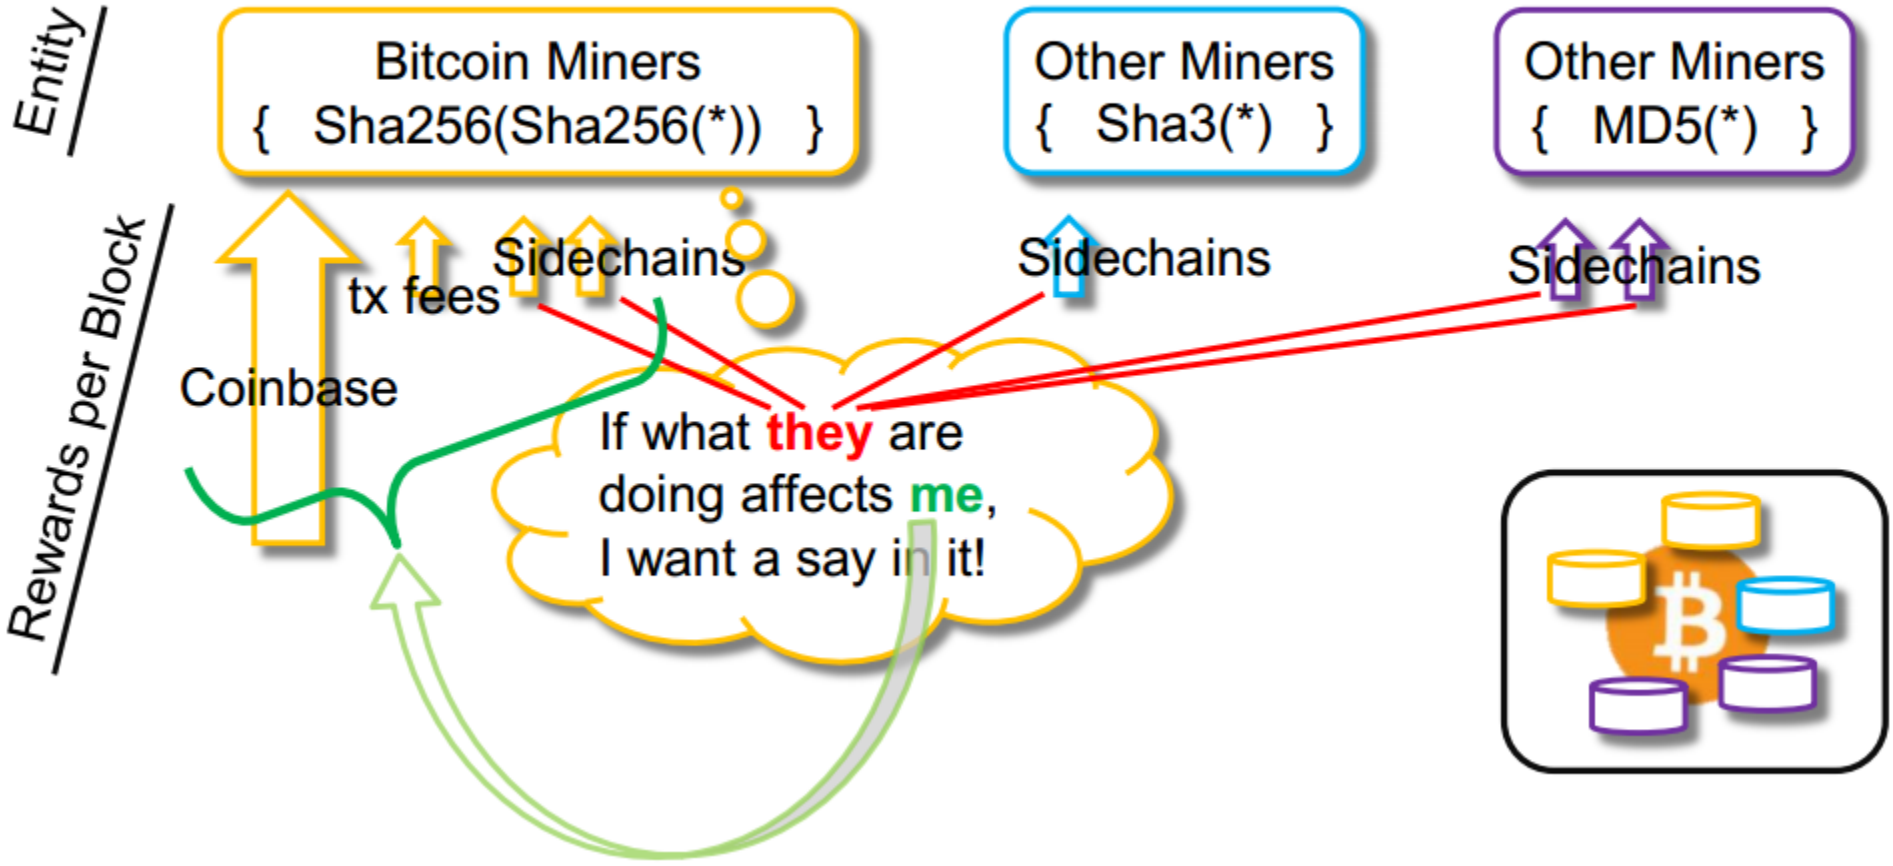

Solves Many Problems

- Nodes “just validate”.
- Miners “just mine”.
  - Drop sidechains they don’t like.
    - That have burdensome node-requirements.
    - That threaten privacy or fungibility.
  - Keep sidechains they like.
- Precedent

Merged Mining

+ Implications of 21 L for MM.
+ Impl. of MM on relative HP.
+ Exploit. of LN / AS, a-SC for Sec.

Before discussing how to handle this, we observe that this risk can be made arbitrarily small by simply increasing the contest period for transfers. Better, the duration of the contest period could be made a function of the relative hashpower of the two chains: the recipient chain might only unlock coins given an SPV proof of one day’s worth of *its own* proof-of-work, which might correspond to several days of the sending chain’s proof-of-work. Security parameters like these are properties of the particular sidechain and can be optimised for each sidechain’s application.

Sidechain Interactivity

- What we're talking about.
- Bitcoin Core, being *affected* by a sidechain.

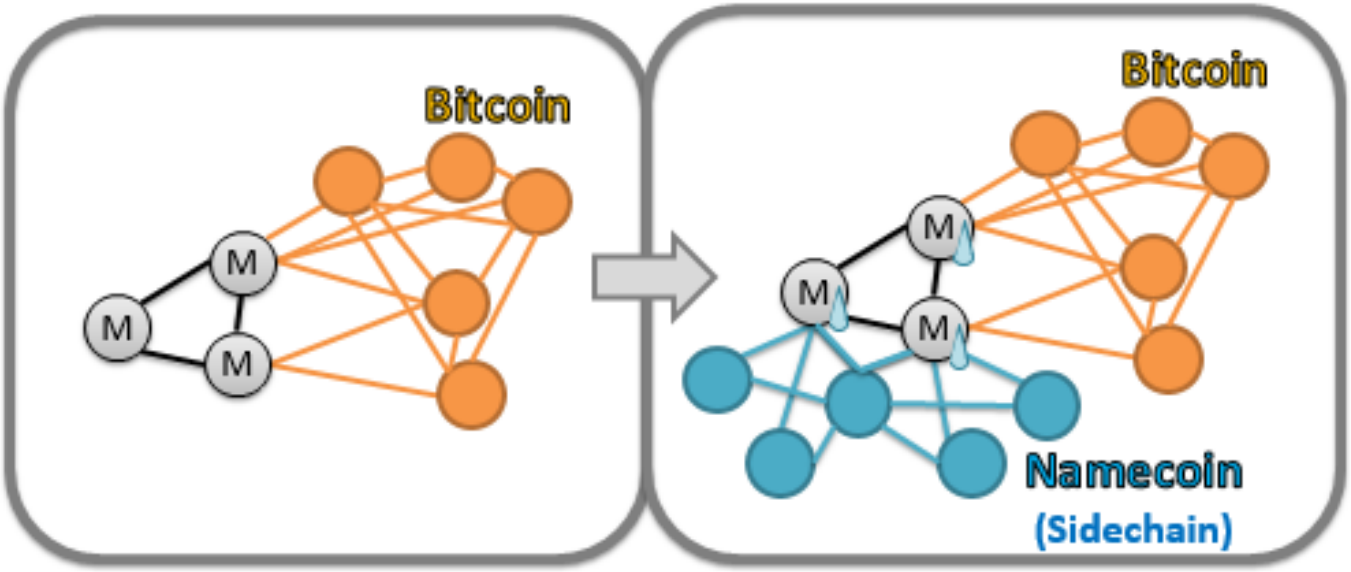

General Solution

Rules for Bitcoin

- Blocks = 1 MB
- SIGOPS = 20,000
- …

Rules for All SCs

- Total Blocks = 1 MB
- Total SIGOPS = 20,000

- Blocks = **0.9** MB
- SIGOPS = 18,000
- …

- Blocks = **0.1** MB
- SIGOPS = 2,000
- …

- This does NOT ends the compulsion for miners to "run all profitable sidechains".
- But it DOES, necessarily, limit the total burden to *exactly* what it was before (ie, pre-sidechain).

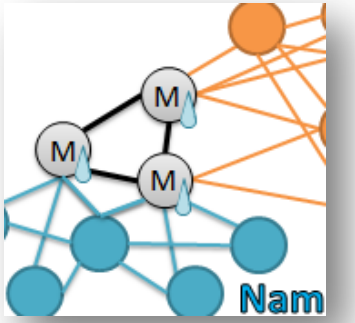

Bad News: Bitcoin Affected

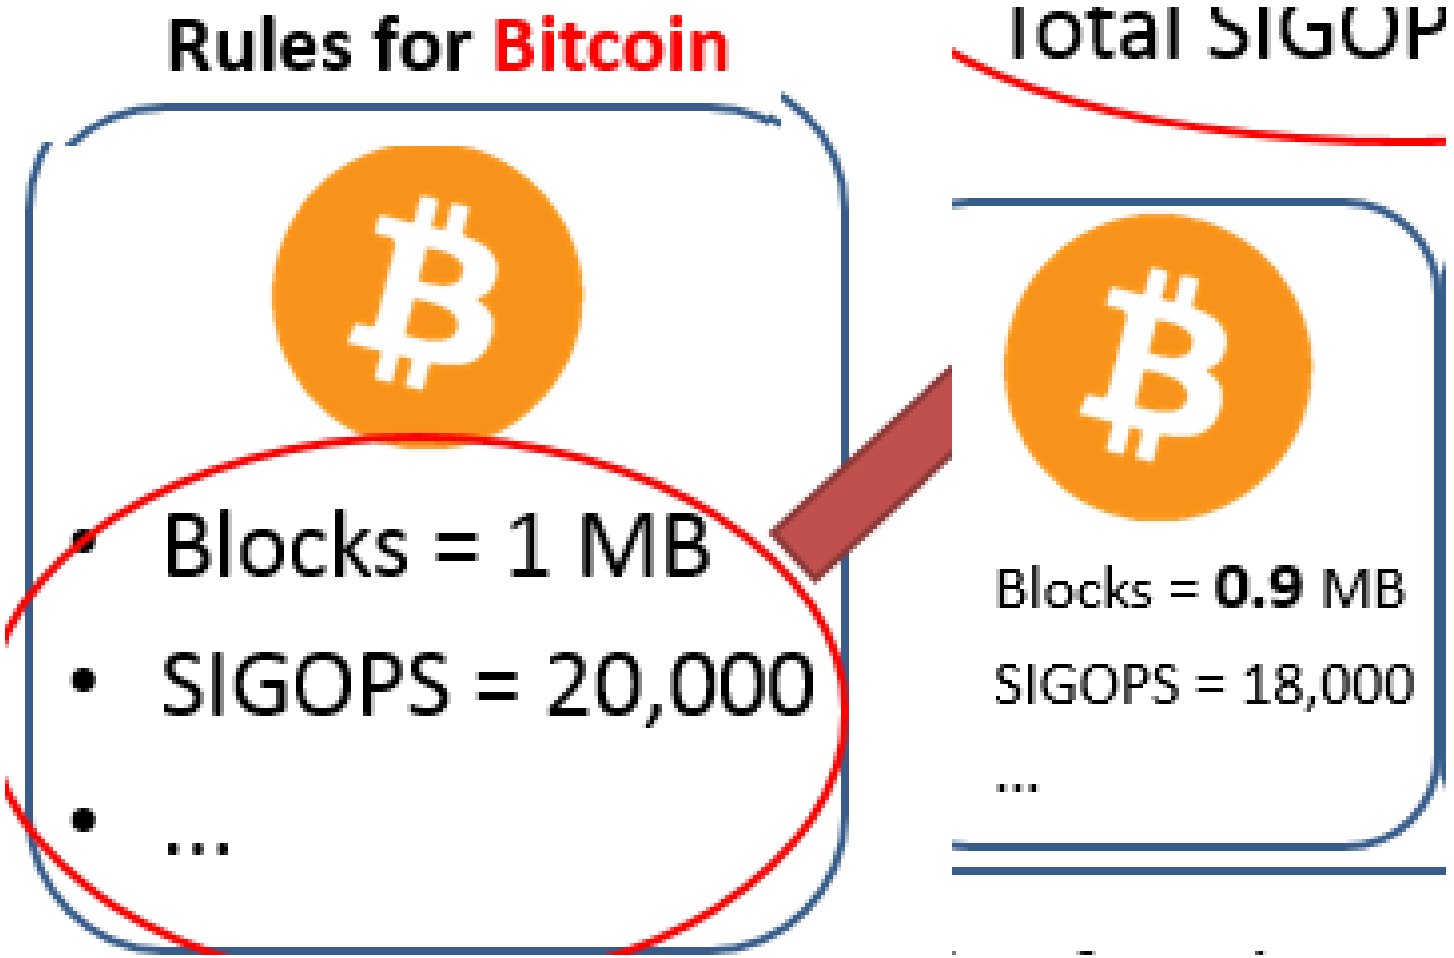

- We'll fix this later.

General Solution

Rules for Bitcoin

- Blocks = 1 MB
- SIGOPS = 20,000
- …

Rules for All SCs

- Total Blocks = 1 MB
- Total SIGOPS = 20,000

- Blocks = **0.9** MB
- SIGOPS = 18,000

- Blocks = **0.1** MB
- SIGOPS = 2,000

Rules enforced by soft fork(s) & “soft fork”-like agreements.

- If users/miners try to add a sidechain in **violation** of this policy, then: (miners should) steal the sidechain’s BTC - deposits.
- Hence, incentive to maintain policy. (+Incentive for miner to pretend to violate policy, then backstab.)

General Solution

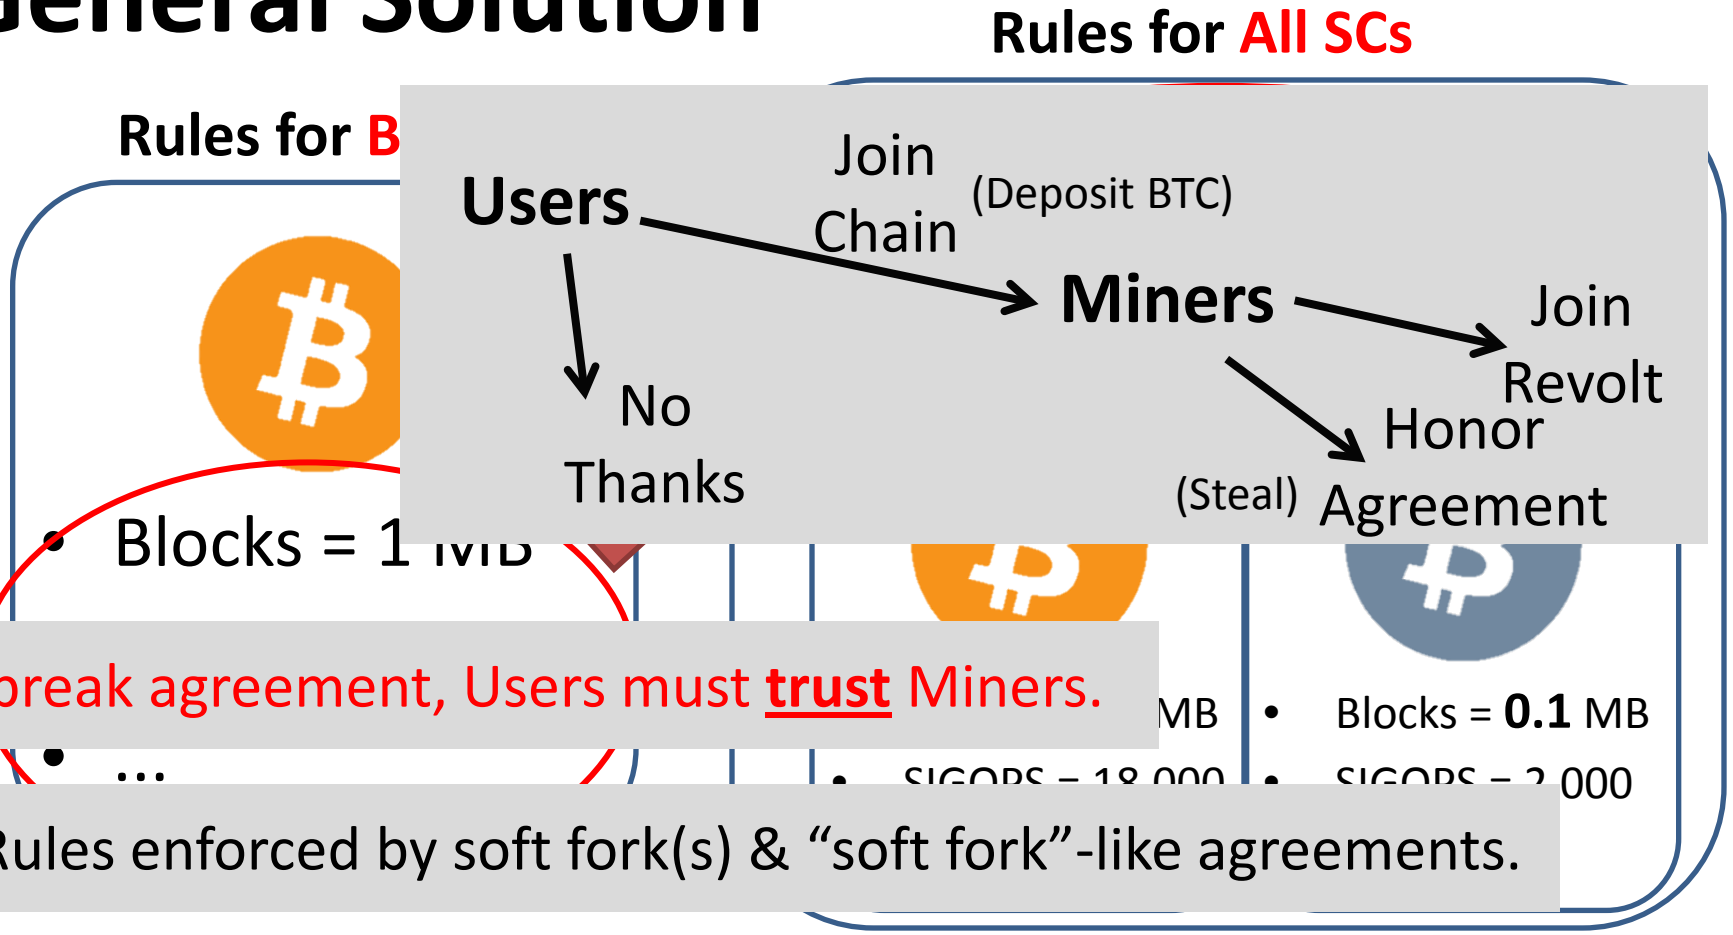

- If users/miners try to add a sidechain in **violation** of this policy, then: (miners should) steal the sidechain’s BTC - deposits.
- Hence, incentive to maintain policy. (+Incentive for miner to pretend to violate policy, then backstab.)

General Solution

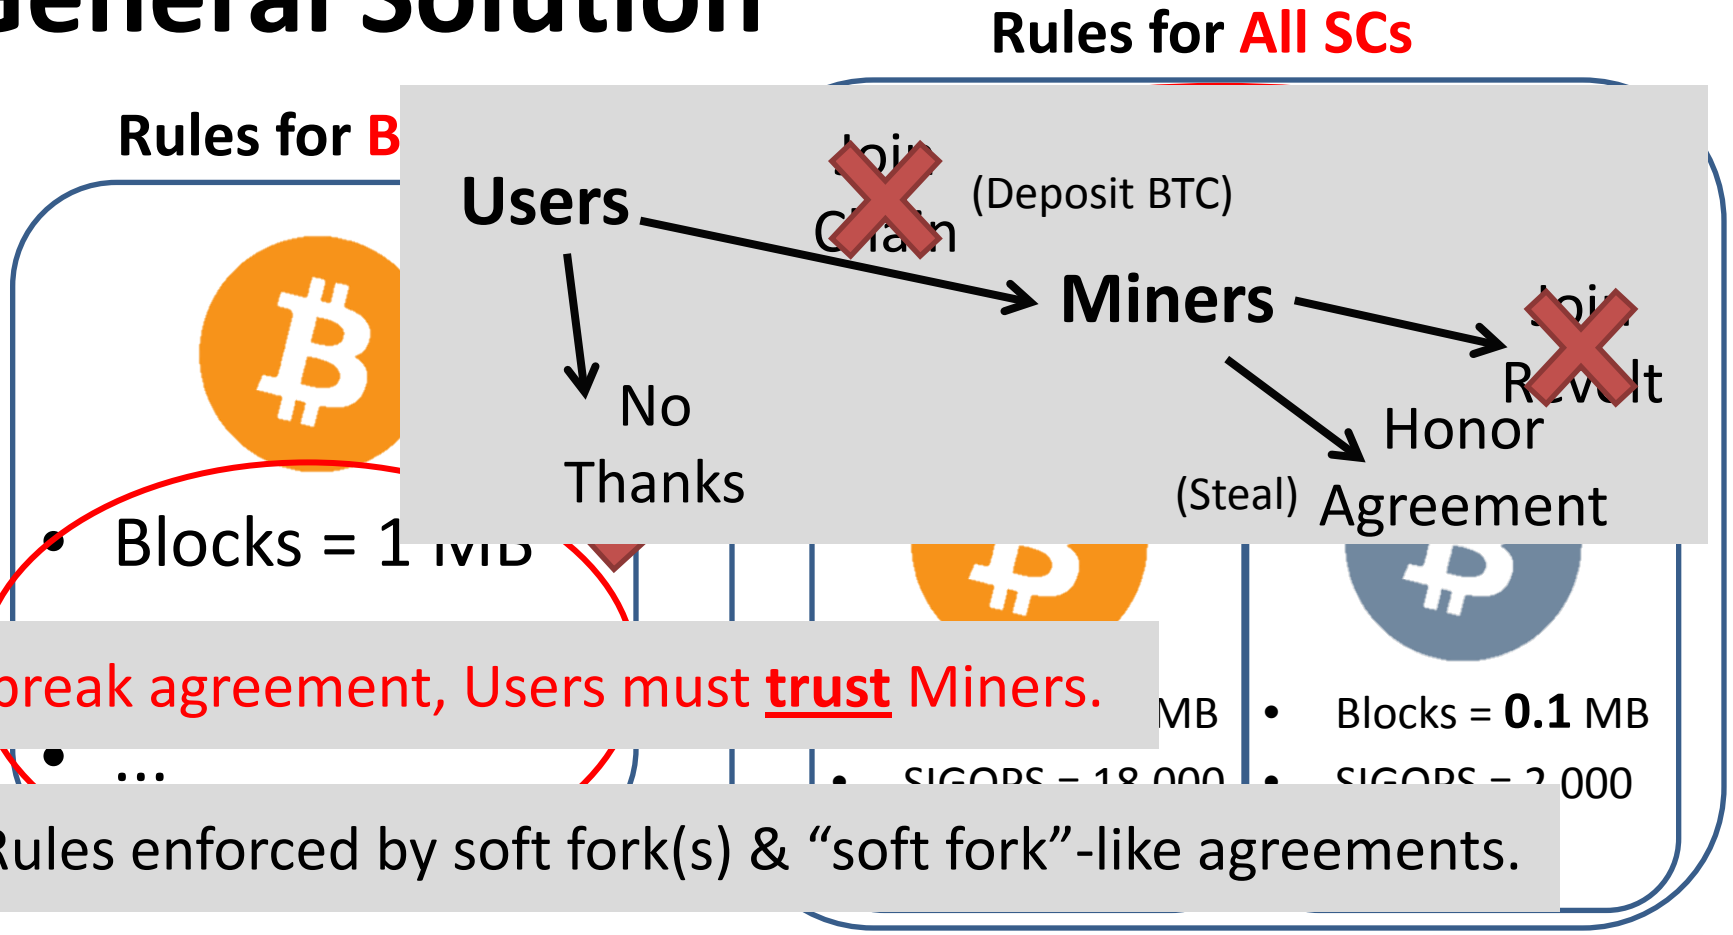

- If users/miners try to add a sidechain in **violation** of this policy, then: (miners should) steal the sidechain’s BTC - deposits.
- Hence, incentive to maintain policy. (+Incentive for miner to pretend to violate policy, then backstab.)

Evil Miners? (Vs Users)

Remember: Users Always Have Option to Ignore New SCs

- Node Costs
  - Opt-In
  - Internalized
  - Anti-Fragile
- Mining
  - "There's only

~~(All) Sidechains~~

- Soft fork to add each *optional* SC.
- SCs are Rare, Topical, ***Profitable***, Sl

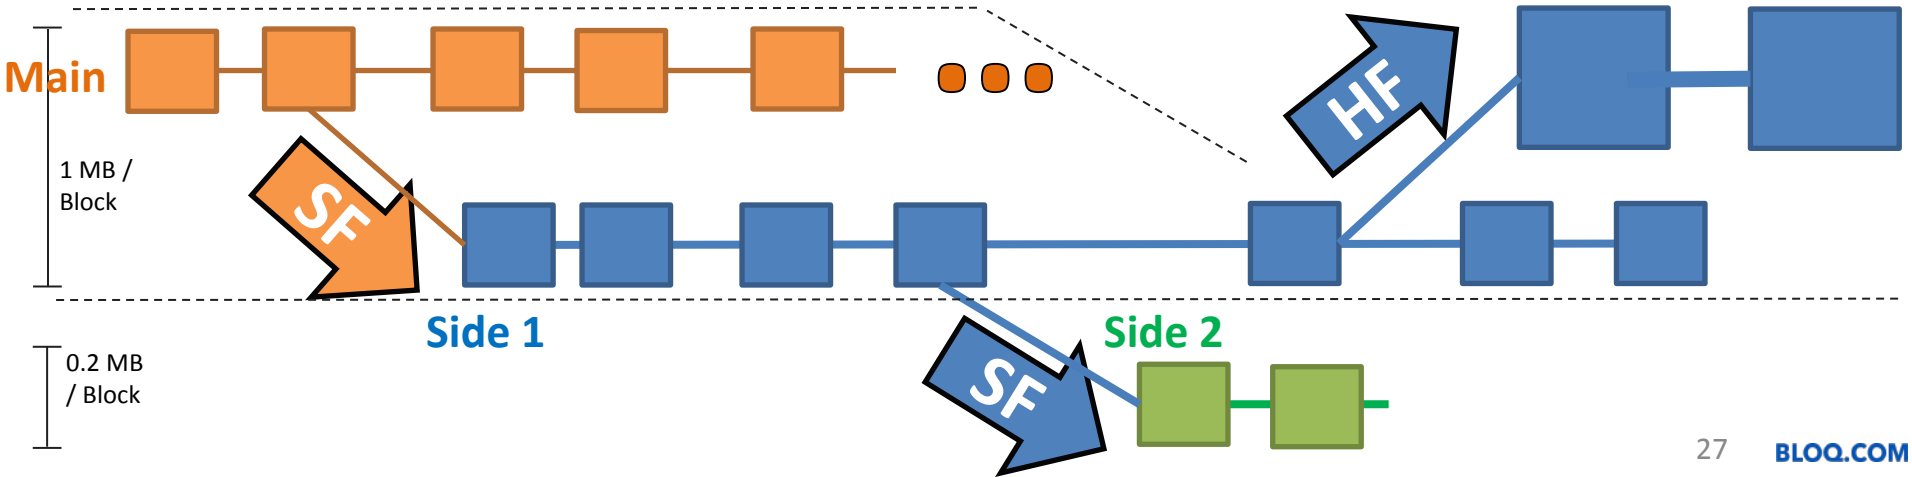

Are these rules enforceable?

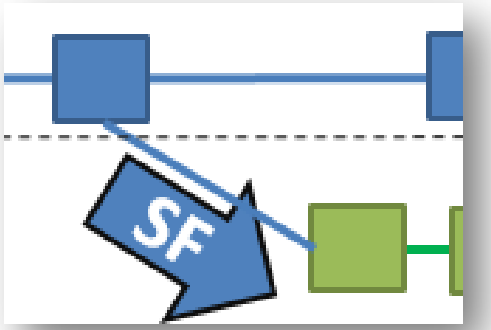

1. To Trespass, Need:
    1. 51% **Miners** (else, funds stolen)
    2. Interested, Fee-Paying **Users**
    3. Code to have been written by **Developers**

2. Consequences of Trespass: Bitcoin *users* not affected. Bitcoin *miners* will be **affecting each other** for the duration of the “transgression”.*
3. *Fundamental Q: How interdependent are mining activities?*

Are these rules enforceable?

"Tame" Sidechains

- **Add new features to Bitcoin.**
- **Contracts are firewalled – opt in, and don't affect each other.**
- **Contracts are managed, to maximize BTC value.**

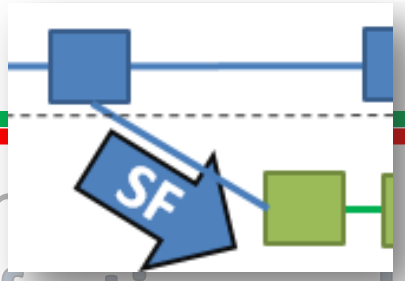

"Aggressive" Sidechains

- **All the benefits of "Tame"...and:**
- ***Permissionless Innovation*.**
- **No debate over precise "split" of validation resources (MB, SigOps).**

Agenda: Part 1

1. The Problem (11)
2. One General Solution (3)
3. Is this GS Robust? (5)
4. Beyond The Limits (5)

Mining Interaction

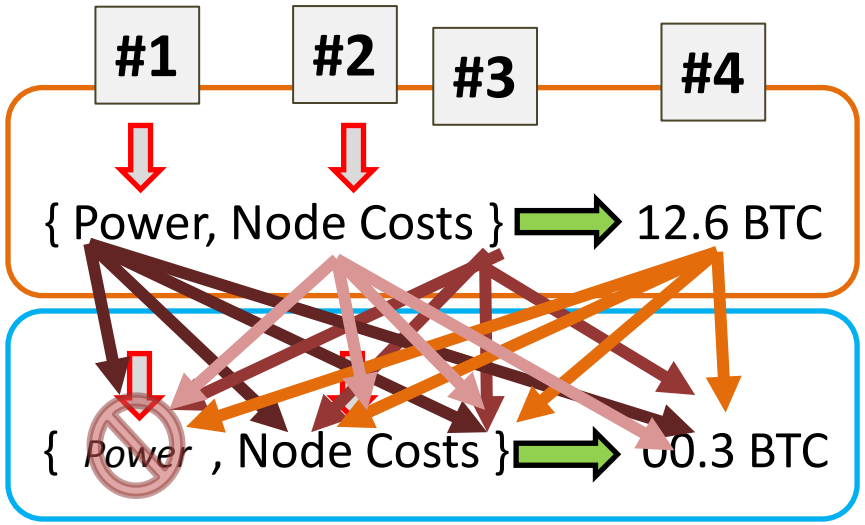

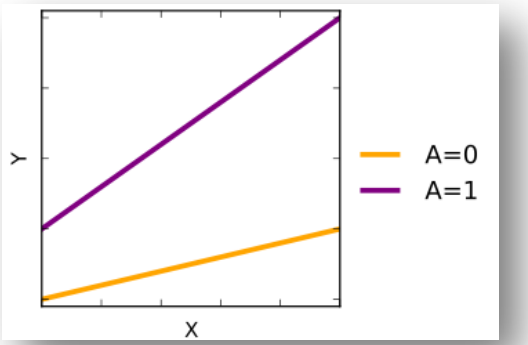

		Capital			Operating	Network	
		Hardware	Software	Know How	Power	Bandwidth	CPU / RAM / Storage
Capital	...						
	...						
	...						
Operating	...						

Costs -- (Non) Interaction

Don't Interact With Anything

		Capital			Operating	Network	
		Hardware (excl. BW)	Software	Know How	Power	Bandwidth	CPU / RAM / Storage
Capital	...	*Indep.*	*Unlikely*	Sub-Additive	*Indep.*		
	...	*Indep.*	*Unlikely*	Sub-Additive	*Indep.*		
	...	*Indep.*	*Unlikely*	Sub-Additive	*Indep.*		
Operating	...	*Indep.*	*Unlikely*	Sub-Additive	*Indep.*		

1. Hashrate is shared, for free.
2. How costly is it to run (more) software? Given from devs.
3. How much *marginal* knowledge is req'ed to add sidechain.

Interaction(?)

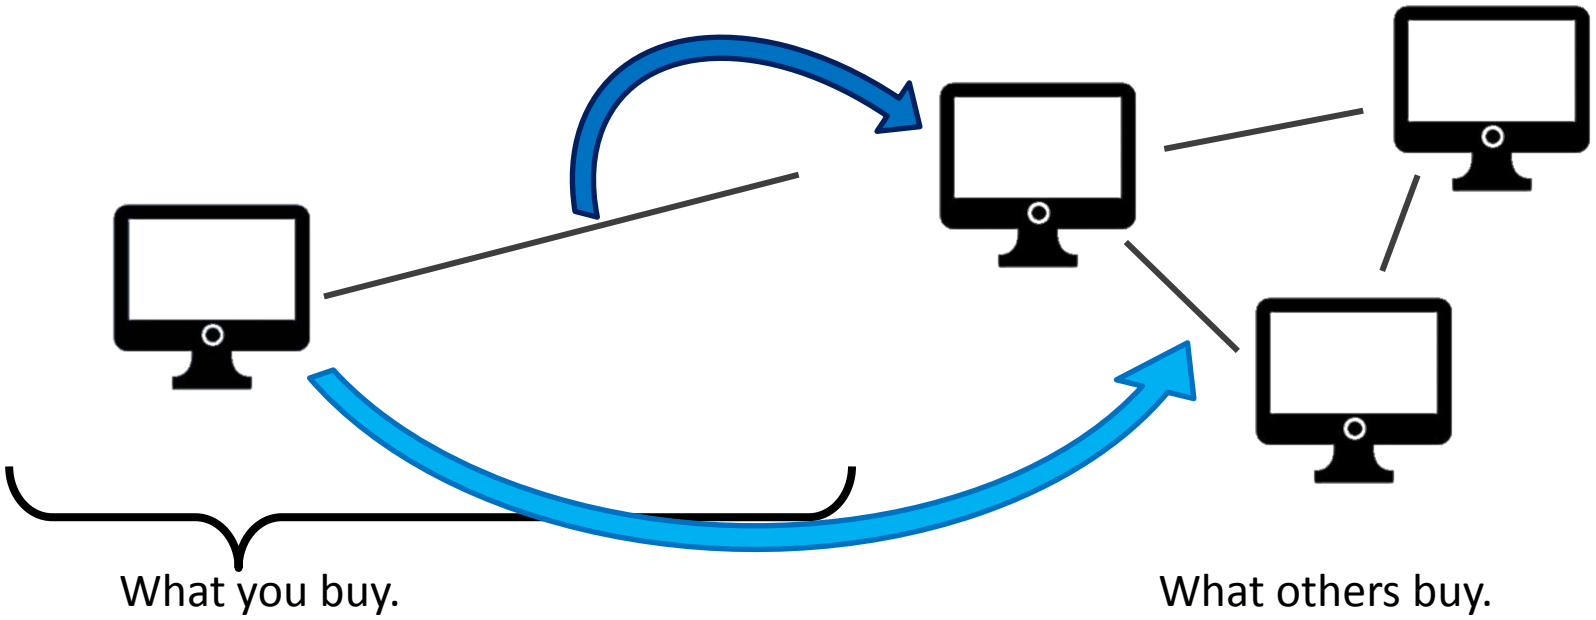

		Network	
		Bandwidth	CPU / RAM / Storage
Network	Bandwidth	?	*
	CPU / RAM / Storage	*	Economies of Scale

Miner Inter-Dependence

1. Bandwidth

   (#2 x #2)

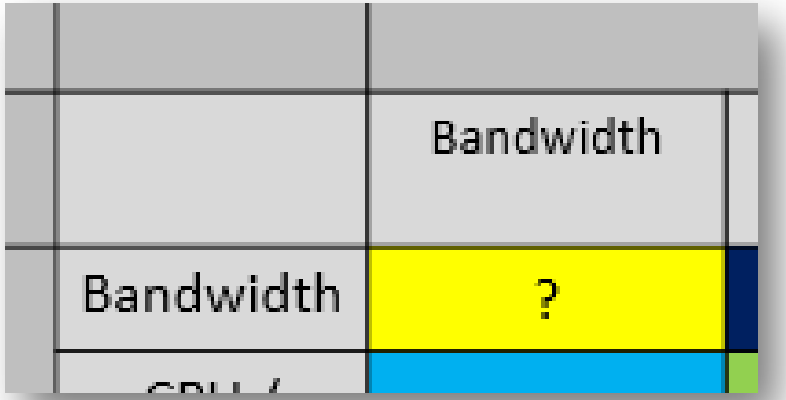

2. Fees

   (#4 x #4)

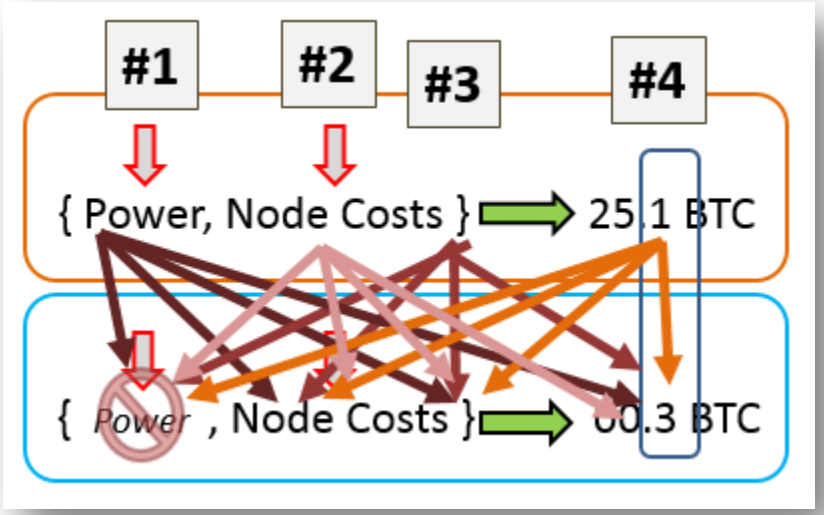

Conclusion: Sidechains convert miners.

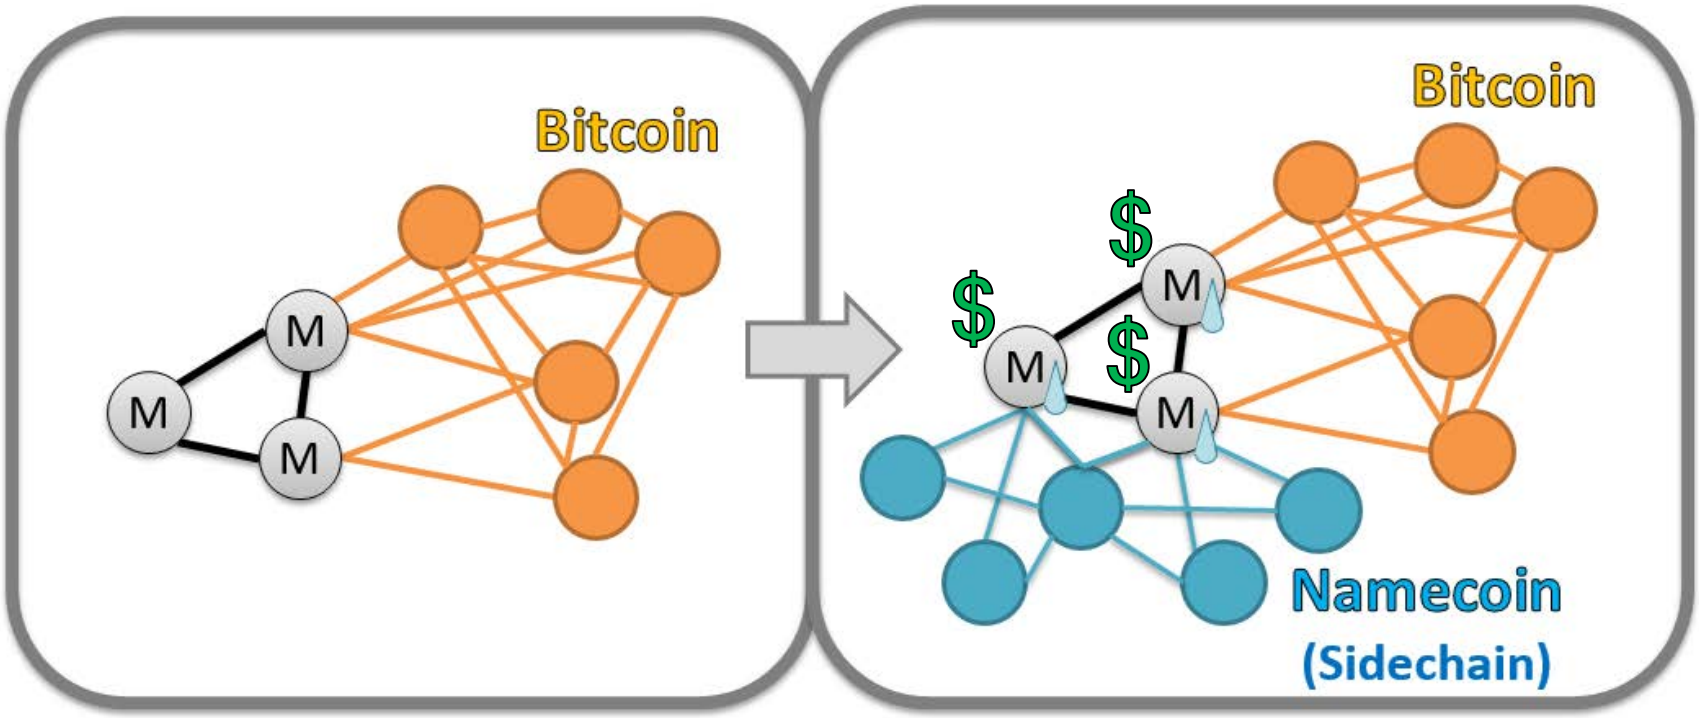

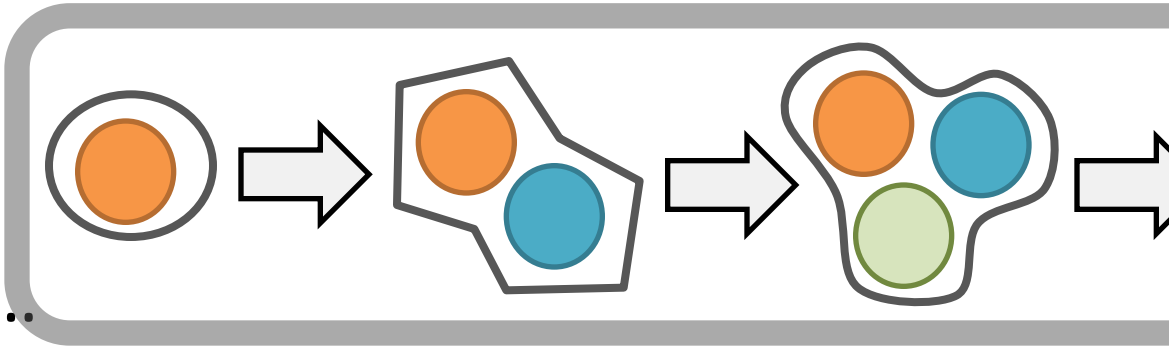

- Nodes are safe.
- Miners not harmed…
- But, miners are *changed*. They “must” absorb pSCs.

Conclusions

Result 1: Total effect limited to “mining conversion”.

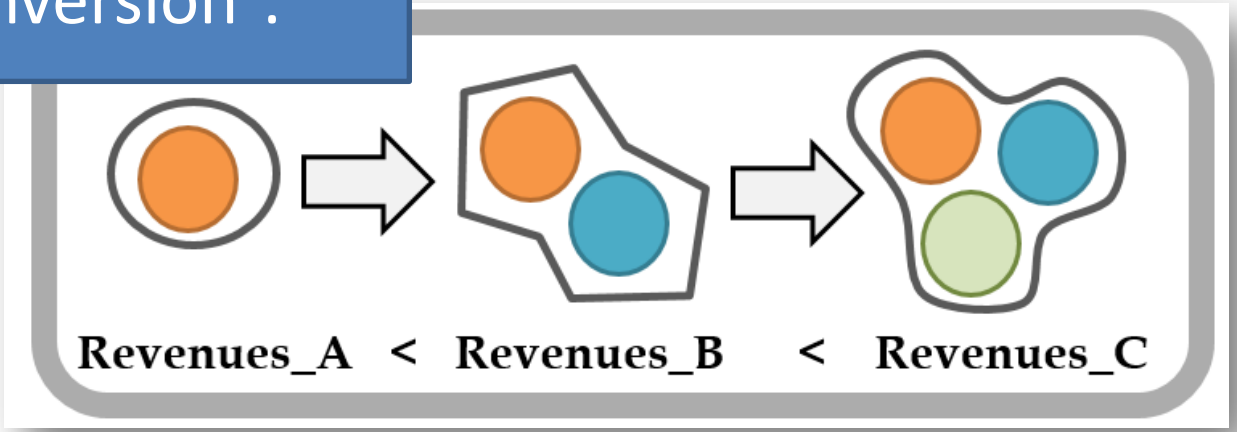

Result 2: Possible for Total_Effect = 0.

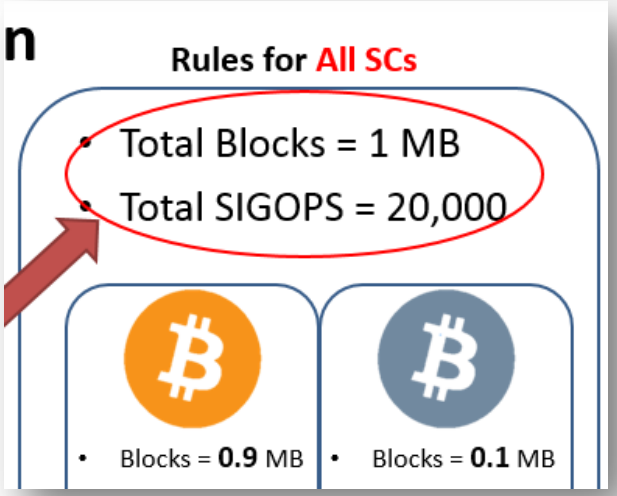

Conclusions

Result 3: Worst case, trespassing res. limits will cause miners to affect each other.

Bonus Points: Operating vs Startup

- Sidechains need [1] nodes…

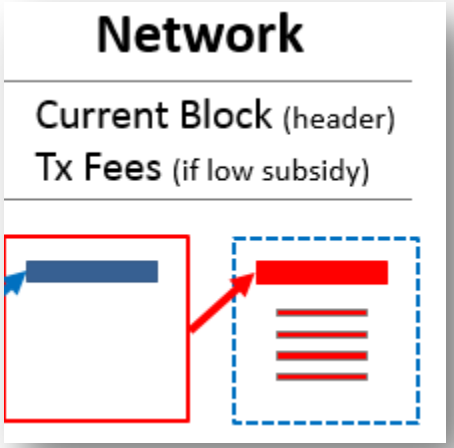

(Sidechains can use UTXO snapshots / checkpoints…and delete old history.)

- …but they *don't* necessarily need [2] an entire history.

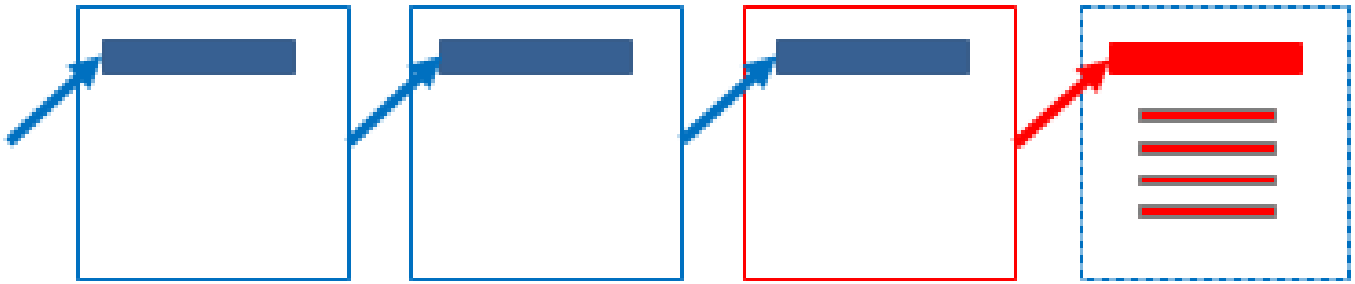

(While Bitcoin needs whole [growing] history, sidechains **may** only need a smaller, fixed amount [of space]).

Part 2 (Interlude): The Docile Miner

1. How ***relevant*** is Mining to Mainchain Bitcoin *investors* and *users*?
2. Mining is innovative and important. Miners (and their decisions) aren't.

Only talking about the Mainchain.

Often Heard

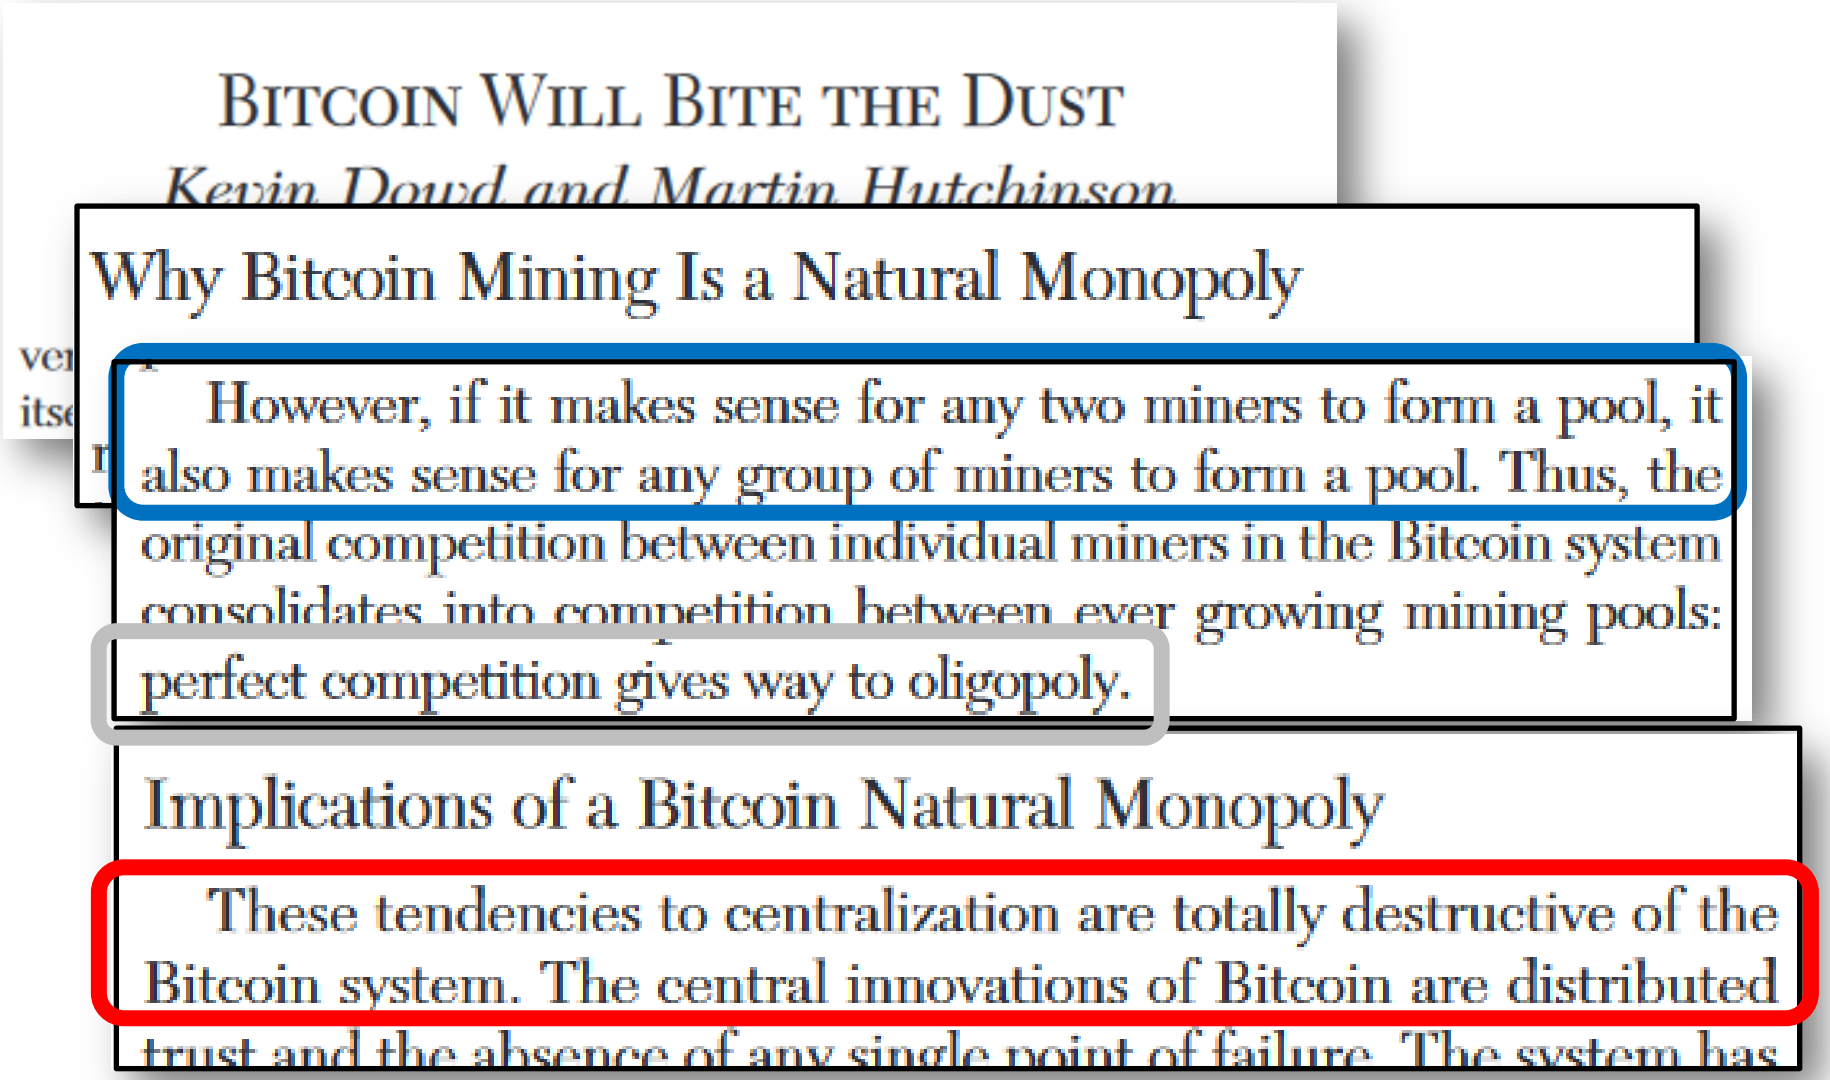

BITCOIN WILL BITE THE DUST

Kevin Dowd and Martin Hutchinson

Why Bitcoin Mining Is a Natural Monopoly

However, if it makes sense for any two miners to form a pool, it also makes sense for any group of miners to form a pool. Thus, the original competition between individual miners in the Bitcoin system consolidates into competition between ever growing mining pools: perfect competition gives way to oligopoly.

Implications of a Bitcoin Natural Monopoly

These tendencies to centralization are totally destructive of the Bitcoin system. The central innovations of Bitcoin are distributed trust and the absence of any single point of failure. The system has

Point about Miners → Point about Bitcoin

Often Heard

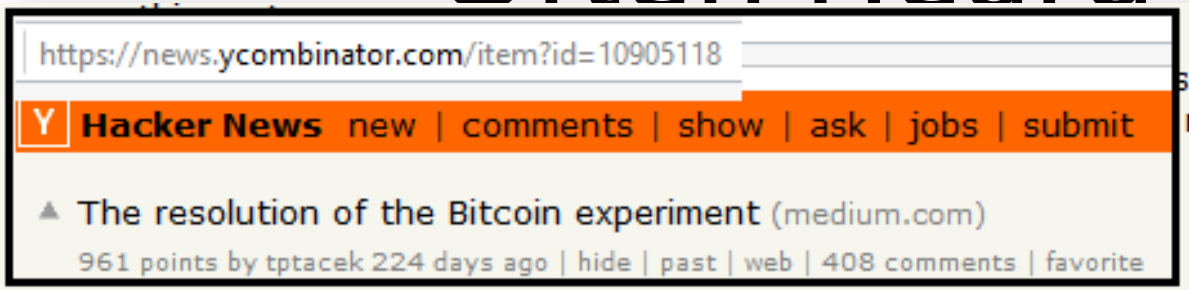

sues with Bitcoin core. At the
rationalization of his decision

▲ mike_hearn 224 days ago [-]

he article. I mentioned Classic briefly at the end. I did not dwell on it
epeating the same process as XT went through.

aring for XT, we also went and talked to the Chinese miners. They told us that the original 20mb limit Gavin proposed was too high, but that they could accept 8mb. So we compromised and went with 8 + a growth function. Then after XT was launched they changed their mind and said any growth after 8 at all was totally unacceptable. Now they're telling the Classic guys that 2 is the most they could handle. Did the Chinese internet border really get 4x worse in the span of 3 months? I doubt it.

Western miners aren't much better. One told me quite clearly they'd start voting for BIP101 back in November (though: voting in such a way that it wouldn't actually activate!). But they didn't. When I followed up, they again said it was on their todo list and they'd start really soon. But they didn't.

The miners have proven over and over again that what they say they will accept and what they actually do accept is not aligned. So right now I'm seeing some excitement (maybe more like desperate hope) that Bitcoin Classic will solve anything. Maybe now the "Scaling Bitcoin" conferences have come and gone and Core's reputation is much worse, they'll have better luck, but even then the *best case scenario* is that Bitcoin gets a 2mb limit. That isn't nearly enough and big backlogs will still occur.

More to the point, even in the best case scenario, the community will essentially accept that Bitcoin is controlled by the Chinese government and grows or shrinks at their whim.

Bitcoin != Mining

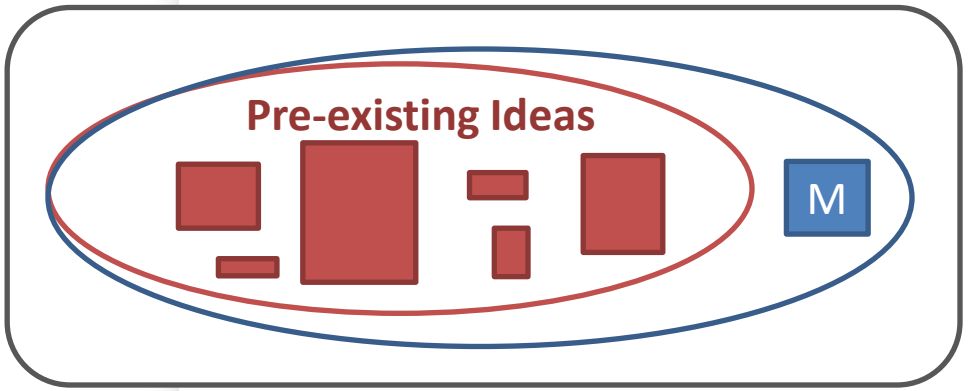

Bitcoin: A Peer-to-Peer Electronic Cash Syst

Satoshi Nakamoto
satoshin@gmx.com
www.bitcoin.org

Abstract. A purely peer-to-peer version of electronic cash would allow onlin
payments to be sent directly from one party to another without going through
financial institution. Digital signatures provide part of the solution, but the mai
benefits are lost if a trusted third party is still required to prevent double-spending
We propose a solution to the double-spending problem using a peer-to-peer networ
The network timestamps transactions by hashing them into an ongoing chain
hash-based proof-of-work, forming a record that cannot be changed without redoin

We're >50% Done Already

-- The Purpose(s) of Mining

By convention, the first transaction in a block is a special transaction that starts a new coin owned by the creator of the block. This adds an incentive for nodes to support the network, and provides a way to initially distribute coins into circulation, since there is no central authority to issue them. The steady addition of a constant of amount of new coins is analogous to gold miners expending resources to add gold to circulation. In our case, it is CPU time and electricity that is expended.

1. Support Network
2. Distribute coins slowly.
   - More important.
   - (*never* affected by any centralization)
   - Unaffected by sidechains.

“Mining” (Process) vs. “Miners” (Agents)

1. Support Network

- Mining is important, miners aren’t.

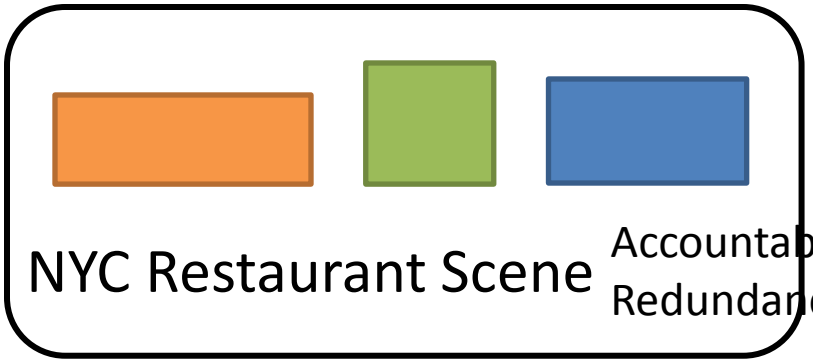

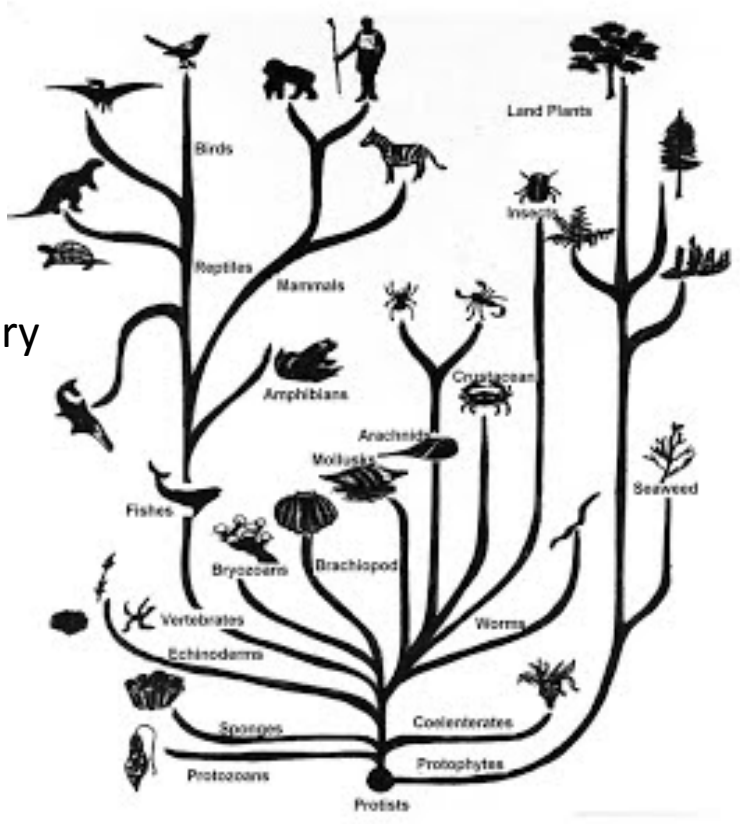

- Competition (Process) vs. Competitors (Agents)
- Indifference → Health
- Scale Fallacy, Anti-fragility.

"Mining" (Process) vs. "Miners" (Agents)

1. Support Network

- Mining is important, miners aren't.

A process, where **anyone *could*** join.

The *people* who did actually choose to join.

Amazing! Wonderful!

Boring. Pedestrian. Irrelevant.

- Com s.
  Com
- Indif
- Scale gility.

Respond to attack by:
- Wait for new miners (easy - ACJ).
- Require additional confirmations.
- (Change mining algorithm.)

Attacker wins, buys miner's hardware, or just coerces miner.

“Chinese” Co-location

- Efficiency improvements (hashes/$) are **good**.
- ***Specialization*** maximizes security, specifically vs. “rented” hashpower (Botnets, AWS).
- Miners **must** take all EIs, to remain competitive.

1. Support Network

Miners are like plants. (On AutoPilot.)

Complaints about **location** = complaints about **ASICs** = complaints about **efficiency**. All miner-choices are efficiency-maximizing.

Improvements in Mining Strategy

Majority is not Enough: Bitcoin Mining is Vulnerable

Ittay Eyal and Emin Gün Sirer

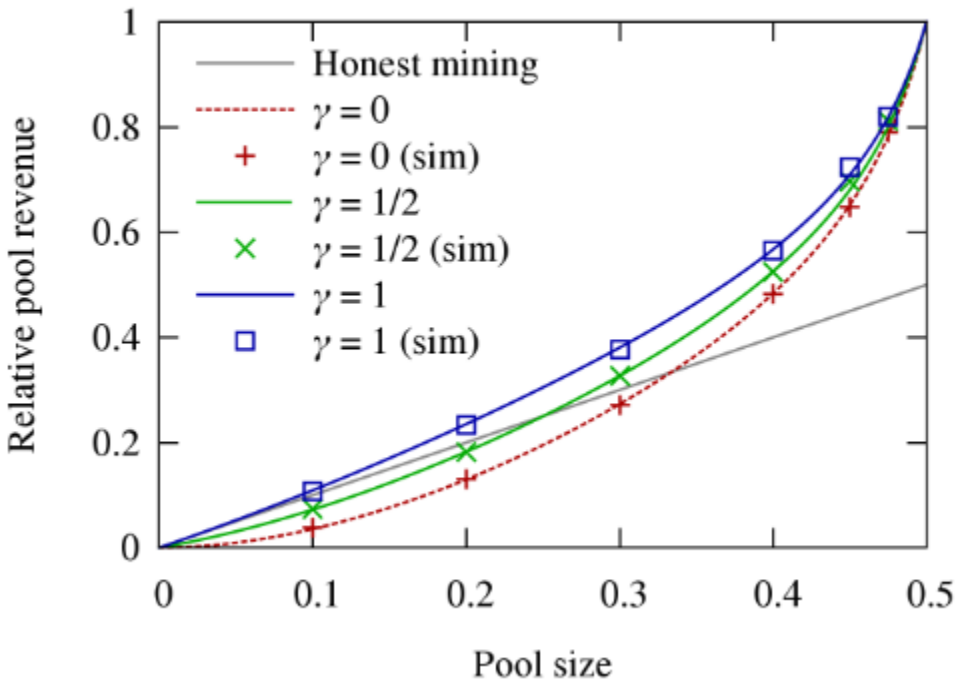

November 04, 2013

Bitcoin Is Broken

Ittay Eyal and Emin Gün Sirer

Bitcoin is broken. And not just superficially so, but fundamentally, at the core protocol level. We're not talking about a simple buffer overflow here, or even a badly designed API that can be easily patched; instead, the problem is intrinsic to the entire way Bitcoin works. All other

analysis, and both show that Selfish-Mine results in hig
the hc
a thr
pends

Complaints about ***strategy*** = complaints about **ASICs** = complaints about **efficiency**.

A Majority is Always Enough

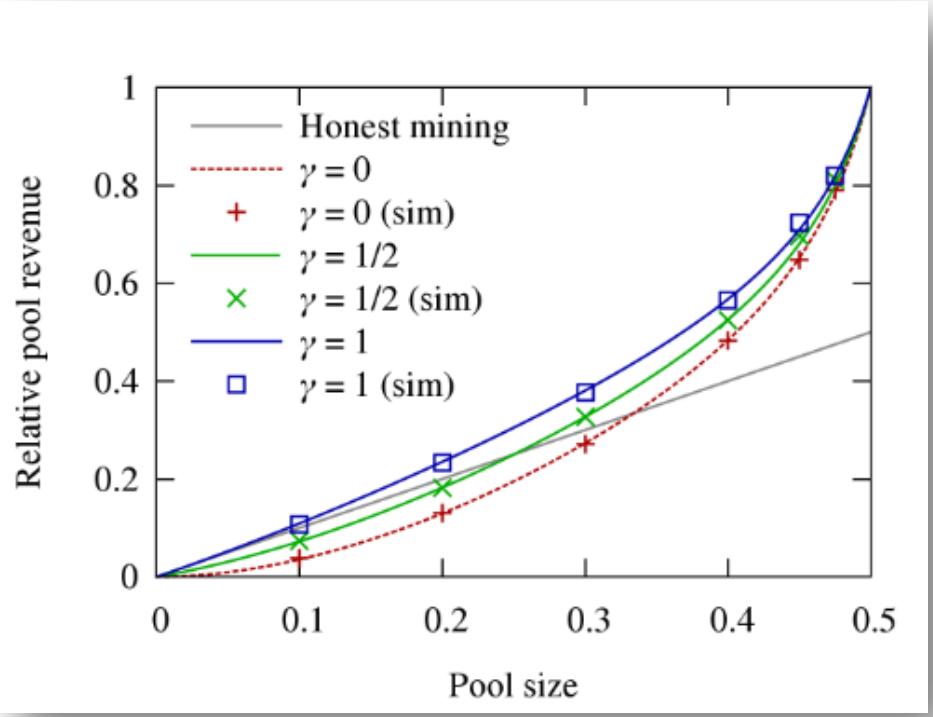

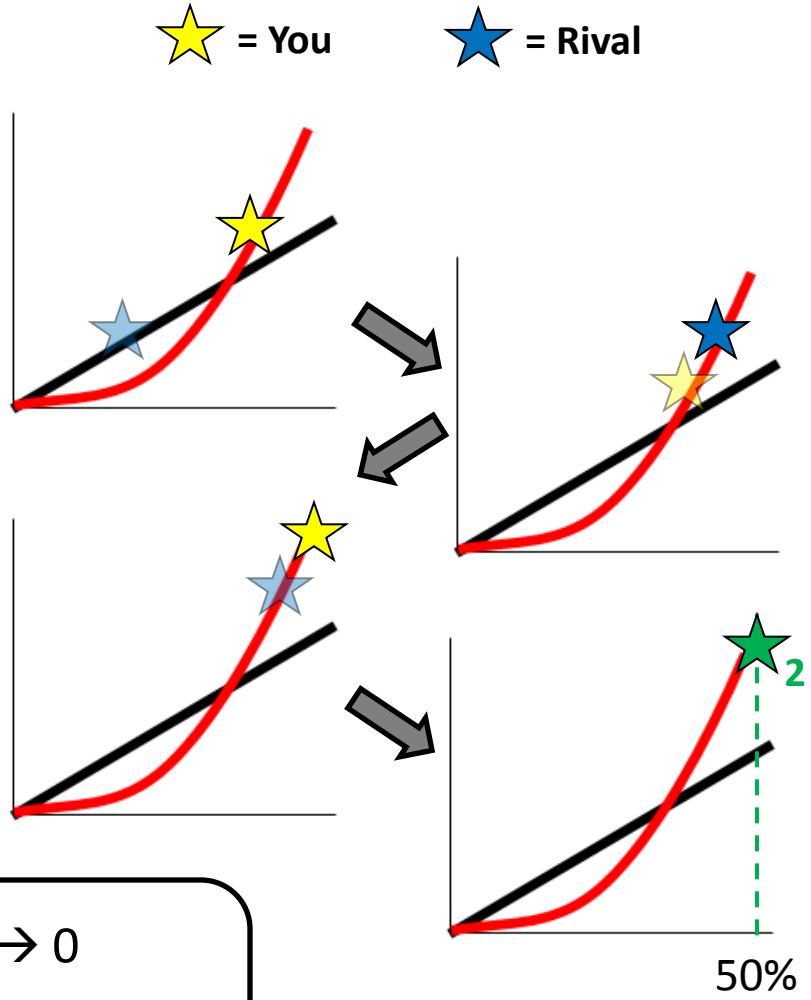

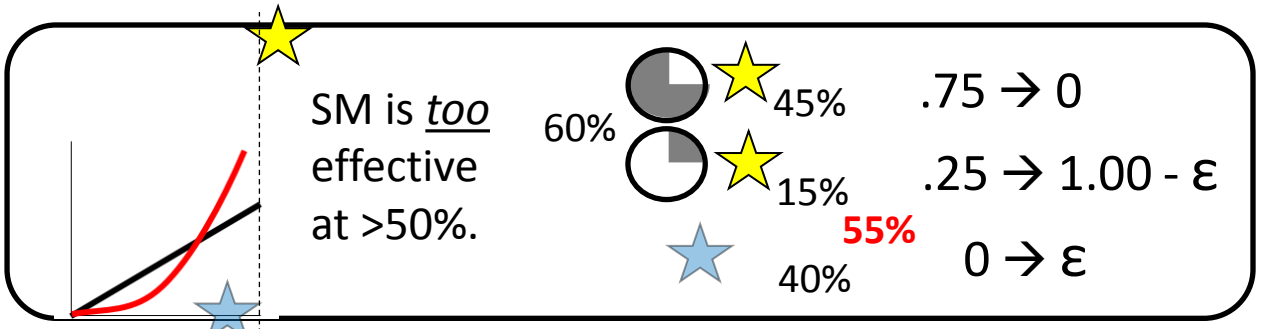

* New = Before – Headache
* Rel Efficiency: 0% or 100%
* Folk Theorem

Miners Can't Do Very Much

- Abilities
  - Individual miner can **only** filter transactions…for *their* block.
  - 51% Miner-Group can **only** order txns.
- Decision Criteria
  - Txns have no identities nor context. As a result, there's no basis for arbitrary censorship (only for economic rationing). This is *ideal!*
  - With 51%, miners can (try to) reorder / re-filter txns in the immediate past. *At tremendous cost / risk.*
- Inabilities. Miners can **never**:
  - "Steal" (Move money without a private key).
  - "Print" (Mint BTC, in excess of the pre-determined schedule).
- Profitable Attack
  - Miner must **double-spend** with *their* funds.
  - …yet, attack affects *all* funds. (Attack must pay off, big!)
  - How many confirmations? ***Subjective*** answer = strategic ***response***.

Miners Can't Do Very Much

- Abilities
  - Individual miner can **only** filter transactions…for *their* block.
  - 51% Miner-Group can **only** order txns.
- Decision Criteria
  - Txns have no identities nor context. As a result, there's no basis for arbit… *deal!*
  - With… imm…
- Inabilities: Miners can **never**:
  - "Steal" (Move money without a private key).
  - "Print" (Mint BTC, in excess of the pre-determined schedule).
- Profitable Attack
  - Miner must **double-spend** with *their* funds.
  - …yet, attack affects *all* funds. (Attack must pay off, big!)
  - How many confirmations? ***Subjective*** answer = strategic ***response***.

With Lightning Network, miners cede even more of their influence (to users).

The Short Leash

- Competition erases profits.
  - Best practices are copied, by rivals.
  - Rivals compete, benefits pass to consumers.
  - Un-copy-able resources become “rents”.

The incentive may help encourage nodes to stay honest. If a greedy attacker is able to assemble more CPU power than all the honest nodes, he would have to choose between using it to defraud people by stealing back his payments, or using it to generate new coins. He ought to find it more profitable to play by the rules, such rules that favour him with more new coins than everyone else combined, than to undermine the system and the validity of his own wealth.

- Mining is *extremely competitive.* “Contestable Market”
  - Anyone can join (ie, anyone can provide hashes).
  - Every hash has an equal chance of winning.
  - Profits will constantly be erased, by difficulty increases.
- Miners more resemble **Subsistence Farmers**.

The Short Leash

Changing the PoW hash function.

- C
  - sumers.
  - Un

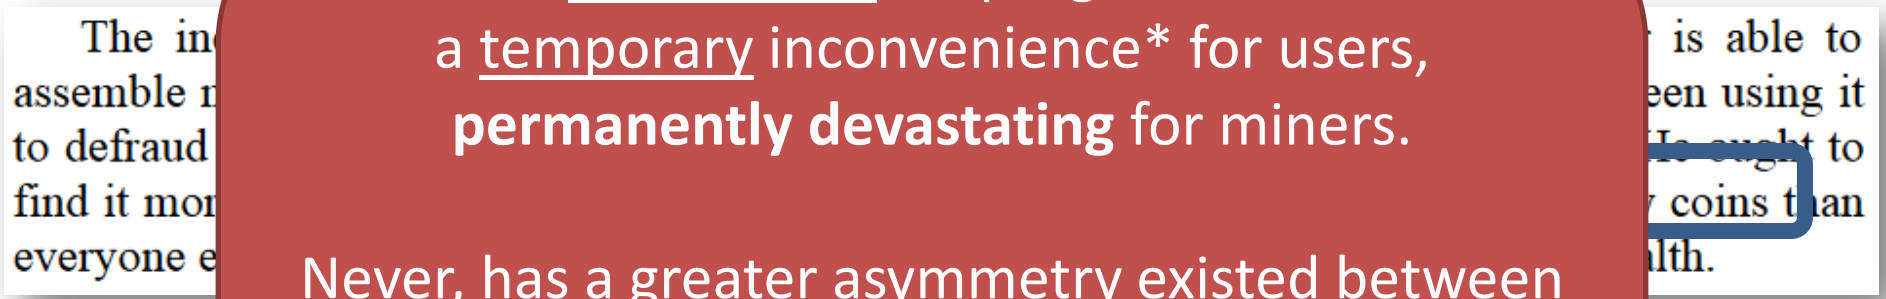
The in
assemble
to defraud
find it mo
everyone e
is able to
een using it
to
coins than
lth.

- Mi
  - A
  - E
  - F
  - Market"
  - ses.
- Mi

A trivial detail for programmers,
a temporary inconvenience* for users,
permanently devastating for miners.

Never, has a greater asymmetry existed between producer and consumer. The miner's ASIC equipment is in perfect competition…with a near-infinite family of hash function combinations.

*We can improve.

The Flow of Influence

https://medium.com/the-coinbase-blog/scaling-bitcoin-the-great-block-siz

Brian Armstrong Follow
Co-Founder and CEO at @Coinbase.
Jan 2 · 10 min read

Luckily, **bitcoin has a built in upgrade mechanism with an elegant design.** If a majority of bitcoin miners "vote" for a particular upgrade then by definition this is the new version of bitcoin. The number of votes each miner

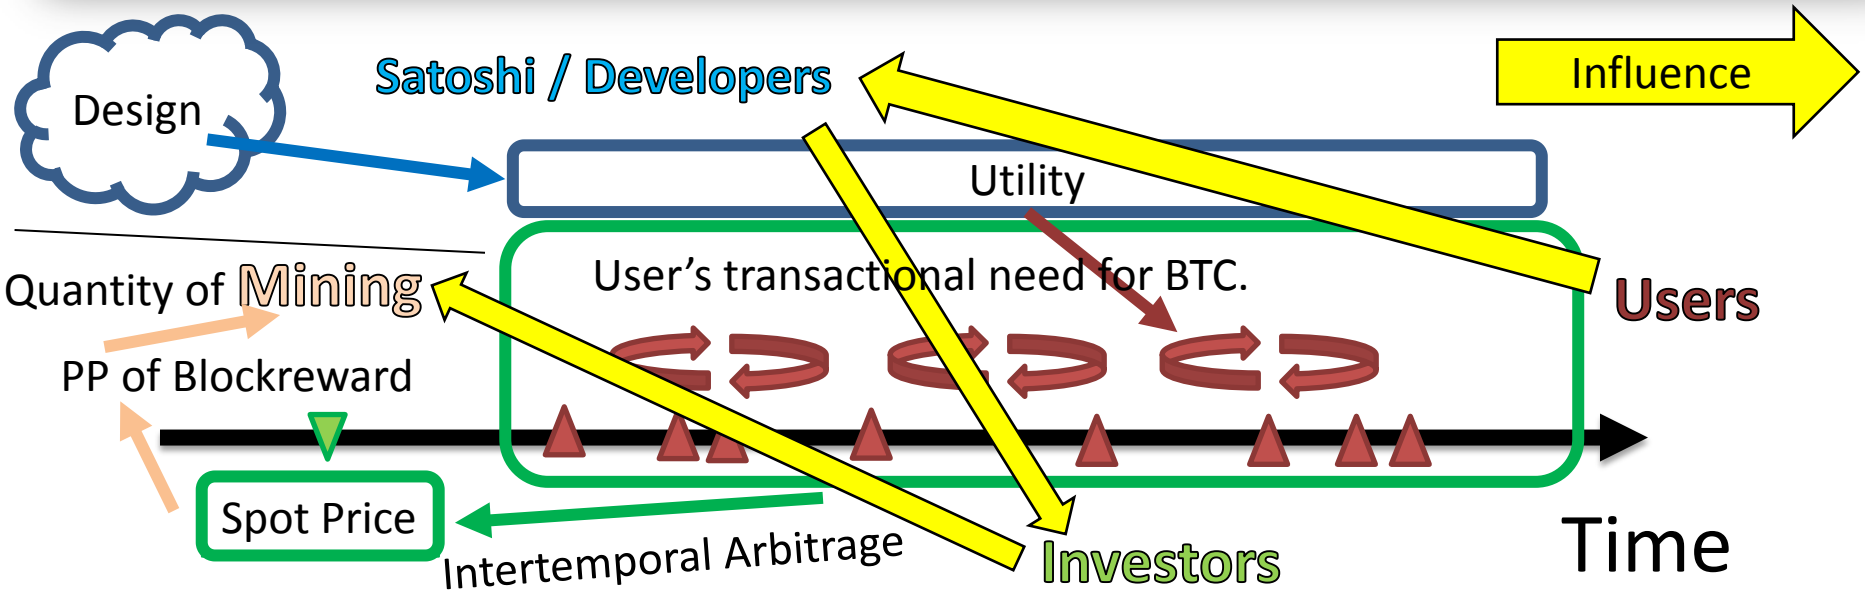

- Users give away "real" goods and services…for digital tokens.
- Expectations of future usage drive "buying of BTC".
- Miners are *lowest* on the "hierarchy of influence".

Miners As Waves

Stimulus = { Spot Price of BTC, Mining Tech & Best Practices }

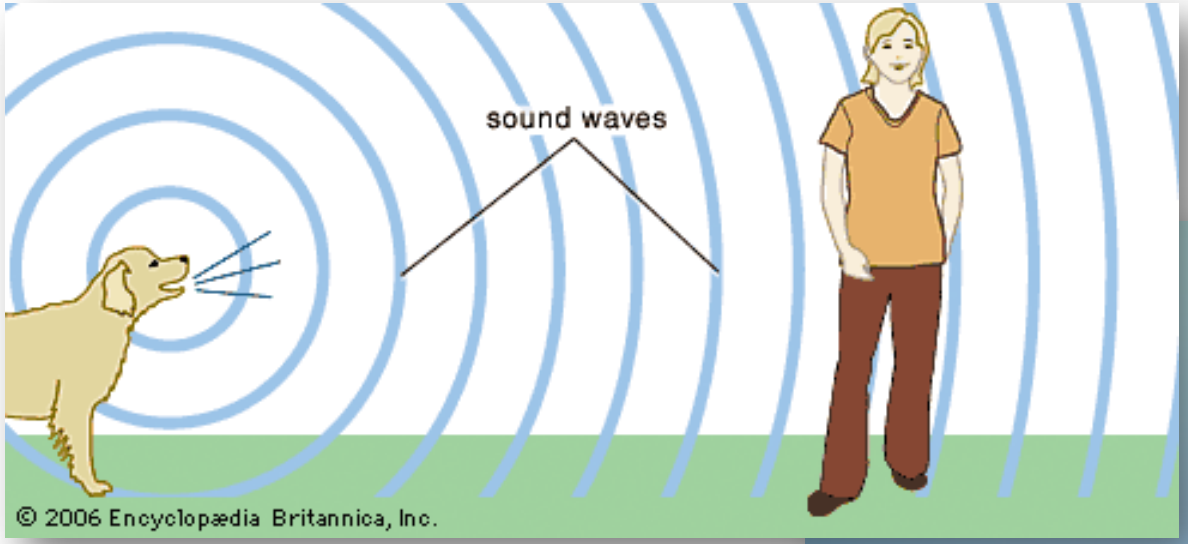

1. Customer is a single numeric function – no substitutes, no excuses, no differentiation, no creativity, .

Conclusion

- Users have a safe relationship with miners.
- Miners don't have many decisions to make:
  - Supply hashes? Include tx? Attempt double-spend?
  - Miners ***must*** walk "the path of efficiency", improving their hardware, software, location, strategy, (etc) .
- Miners are on a short leash.
- Mining is **caused**, it is not "a cause".

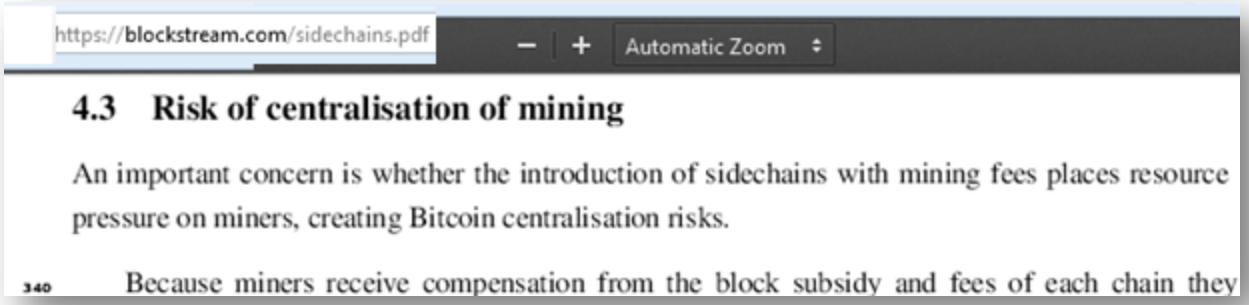
https://blockstream.com/sidechains.pdf

Automatic Zoom

4.3 Risk of centralisation of mining

An important concern is whether the introduction of sidechains with mining fees places resource pressure on miners, creating Bitcoin centralisation risks.

340 Because miners receive compensation from the block subsidy and fees of each chain they

Anyone can mine.
Anyone can become a peer.
Progress sans identity.

Node Network

Mining

Part 3 - Bandwidth

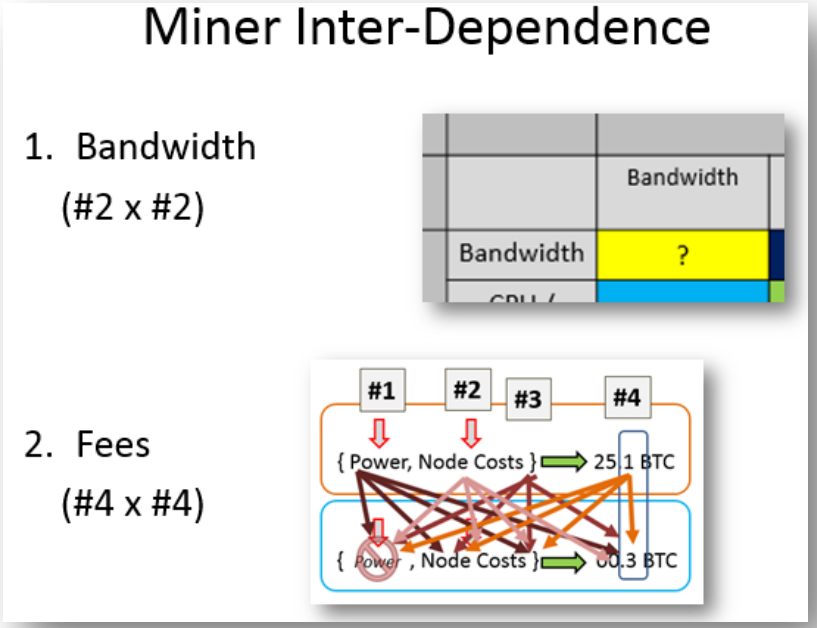

1. Isn't bandwidth just another full node cost?
2. If not, *why* not?

Bandwidth – Beyond the Limits

- Node Costs
  - Opt-In
  - Internalized
  - Anti-Fragile
- Mining
  - "There's only
  - How can diffe
    common netv

"Tame" Sidechains

- Add new features to Bitcoin.
- Contracts are firewalled – opt in, and don't affect each other.
- Contracts are managed, to maximize BTC value.

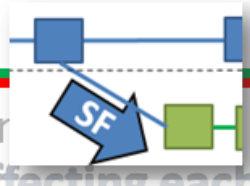

"Aggressive" Sidechains

- All the benefits of "Tame"...and:
- ***Permissionless Innovation***.
- No debate over precise "split" of validation resources (MB, SigOps).

What happens to mining, on mainchain Bitcoin?

Bandwidth – Beyond the Limits

- Node Costs
  - Opt-In
  - Internalized
  - Anti-Fragile
- Mining
  - "There's only
  - How can diffe common netv

"Tame" Sidechains

- Add new features to Bitcoin.
- Contracts are firewalled – opt in, and don't affect each other.
- Contracts are managed, to maximize BTC value.

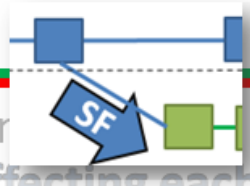

"Aggressive" Sidechains

- All the benefits of "Tame"...and:
- ***Permissionless Innovation***.
- No debate over precise "split" of validation resources (MB, SigOps).

What happens to mining, on mainchain Bitcoin?

There's something **special** about "propagation".

- Or, let's call it "connectivity" or "bandwidth".

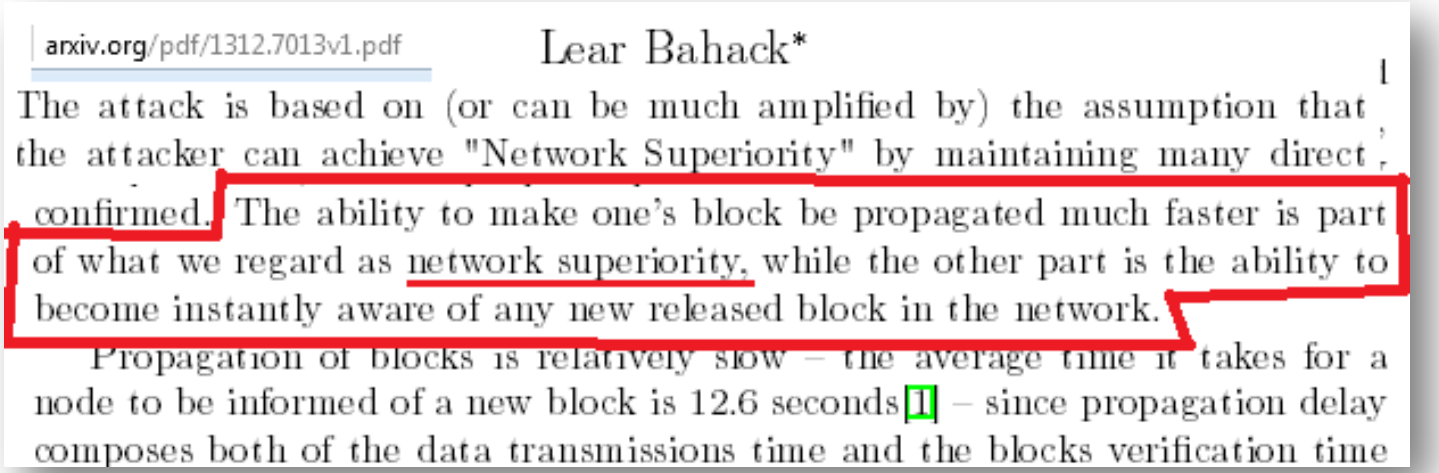
arxiv.org/pdf/1312.7013v1.pdf

Lear Bahack*

The attack is based on (or can be much amplified by) the assumption that the attacker can achieve "Network Superiority" by maintaining many direct confirmed. The ability to make one's block be propagated much faster is part of what we regard as network superiority, while the other part is the ability to become instantly aware of any new released block in the network.

Propagation of blocks is relatively slow – the average time it takes for a node to be informed of a new block is 12.6 seconds [1] – since propagation delay composes both of the data transmissions time and the blocks verification time

- Compare the **non-special**: labor, power, hardware, cooling tech, land.
- These are **internalized**: Improvement = higher profits.

Section Agenda

- Problem (9 Slides)
- Other Research (4)
- Solution (7)
- Safety (4)
- Improvements (2)

Inherently Interpersonal

Just as a chain is only as strong as its weakest link, a broadcast network is only as fast as its slowest bilateral connection.

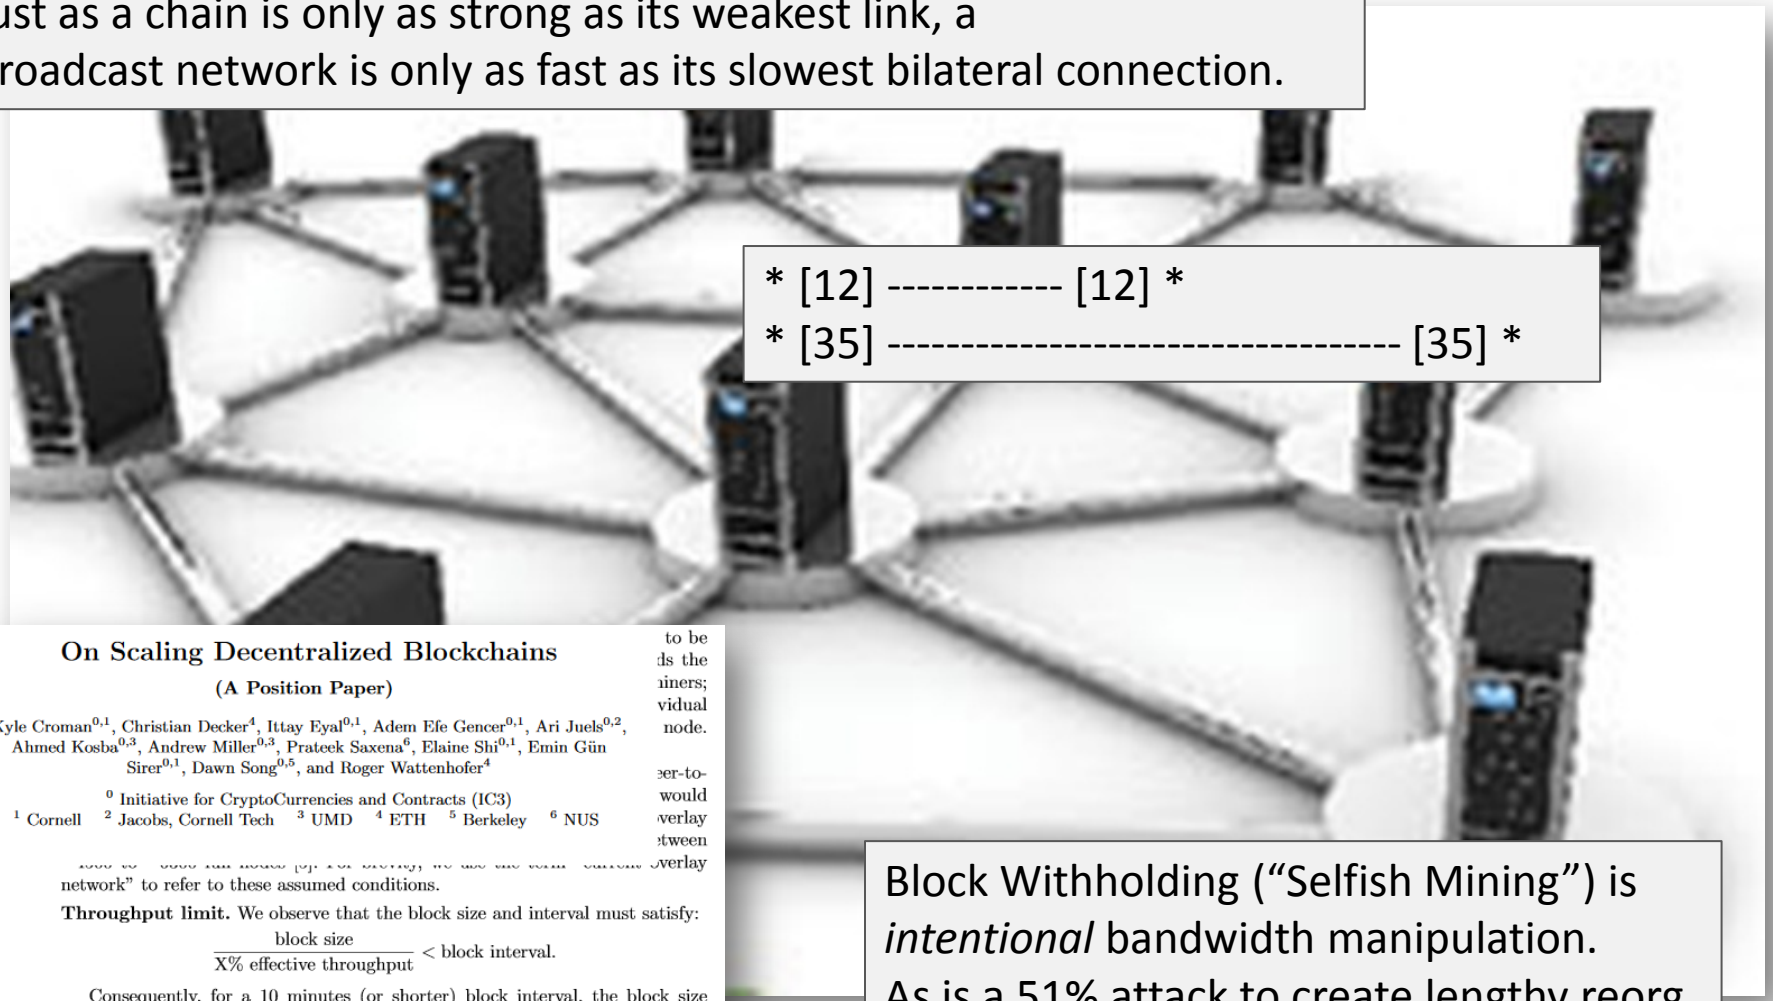

* [12] ---------- [12] *
* [35] ---------- [35] *

On Scaling Decentralized Blockchains

(A Position Paper)

Kyle Croman[0,1], Christian Decker[4], Ittay Eyal[0,1], Adem Efe Gencer[0,1], Ari Juels[0,2], Ahmed Kosba[0,3], Andrew Miller[0,3], Prateek Saxena[6], Elaine Shi[0,1], Emin Gün Sirer[0,1], Dawn Song[0,5], and Roger Wattenhofer[4]

[0] Initiative for CryptoCurrencies and Contracts (IC3)
[1] Cornell [2] Jacobs, Cornell Tech [3] UMD [4] ETH [5] Berkeley [6] NUS

to be
ds the
iners;
vidual
node.
eer-to-
would
verlay
etween
verlay

network" to refer to these assumed conditions.

Throughput limit. We observe that the block size and interval must satisfy:

$$\frac{\text{block size}}{\text{X\% effective throughput}} < \text{block interval}.$$

Consequently, for a 10 minutes (or shorter) block interval, the block size should not exceed 4MB for X=90%; and 38MB for X=50%.

Block Withholding ("Selfish Mining") is *intentional* bandwidth manipulation.
As is a 51% attack to create lengthy reorg.

Orphaning “Costs” - Bifurcation

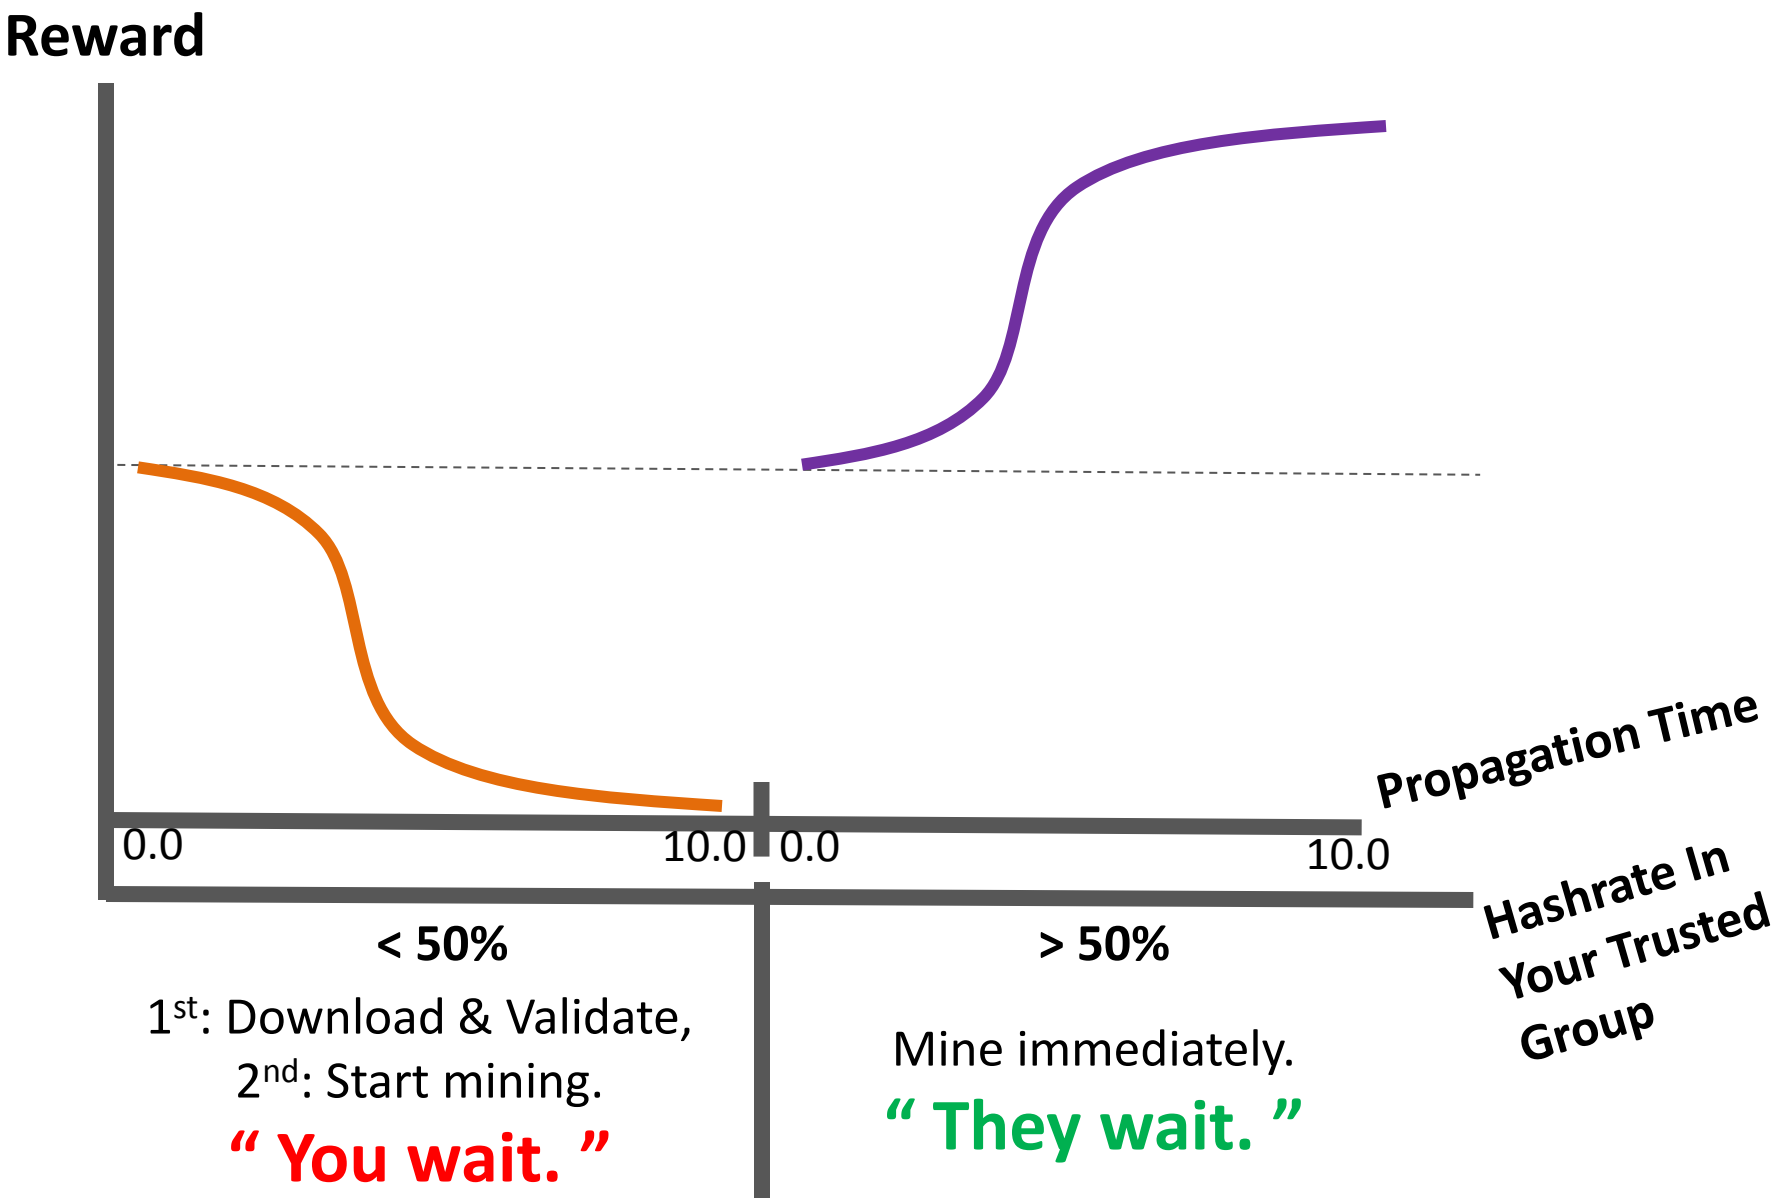

The Propagation Paradox: Connectivity Down, Profits…Up?!

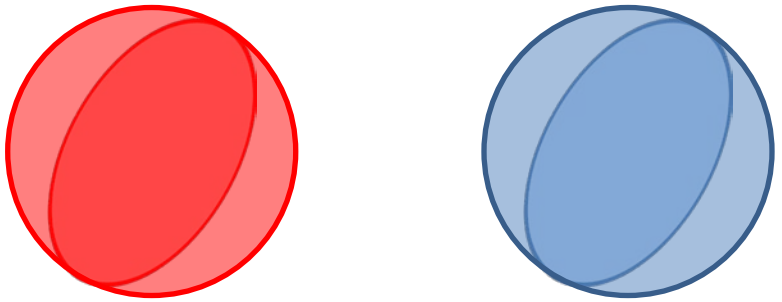

Planet:	Mars	Earth
Hashrate:	75%	25%
Txn Volume:	10%	90%

- 75% of hashrate teleports to Mars.
- It takes 1 hour for messages to pass between planets.
- (Mars miners are not necessarily coordinating with each other).

Paradox: Bandwidth Bad, Profits Up

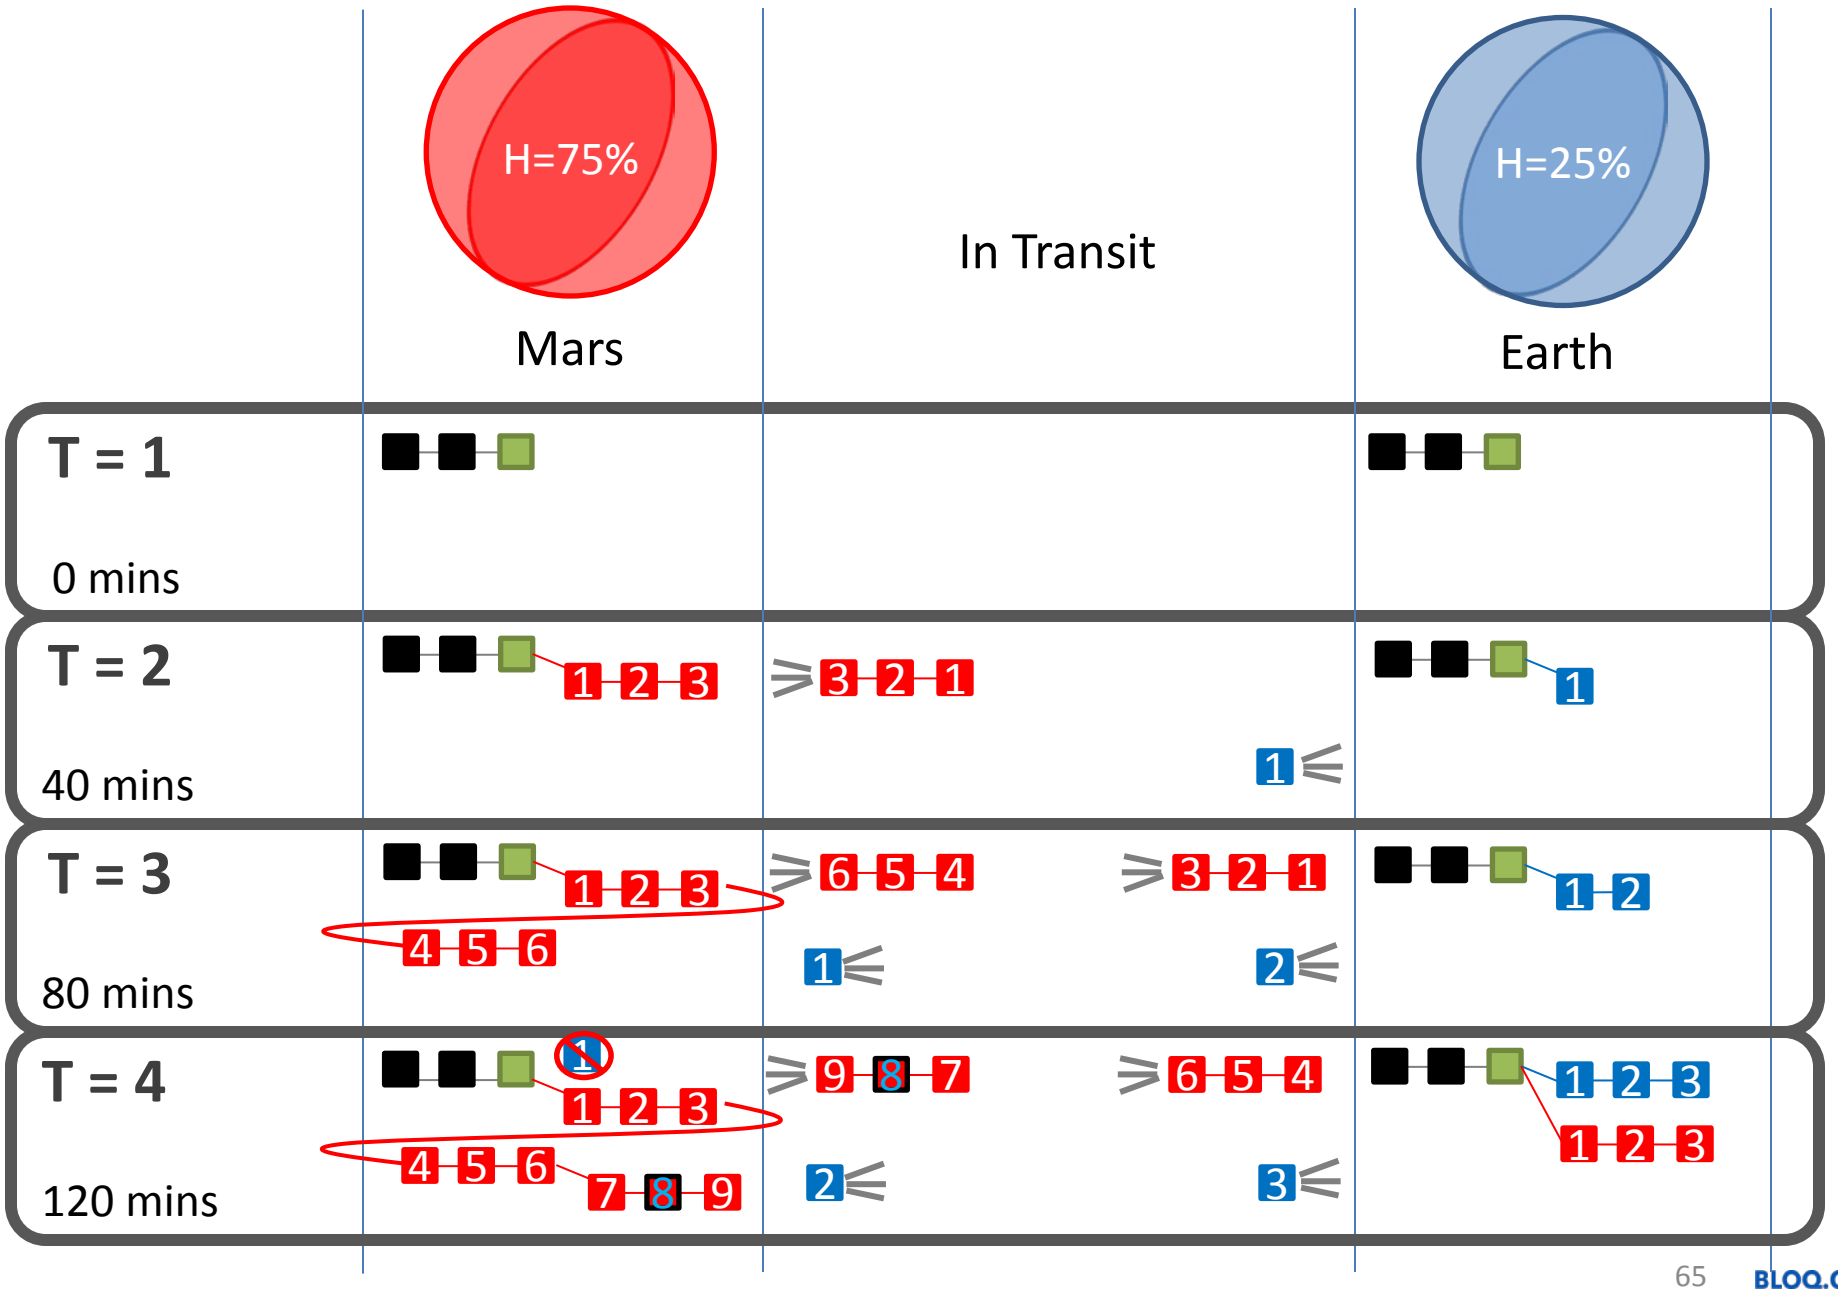

Paradox: Bandwidth Bad, Profits Up

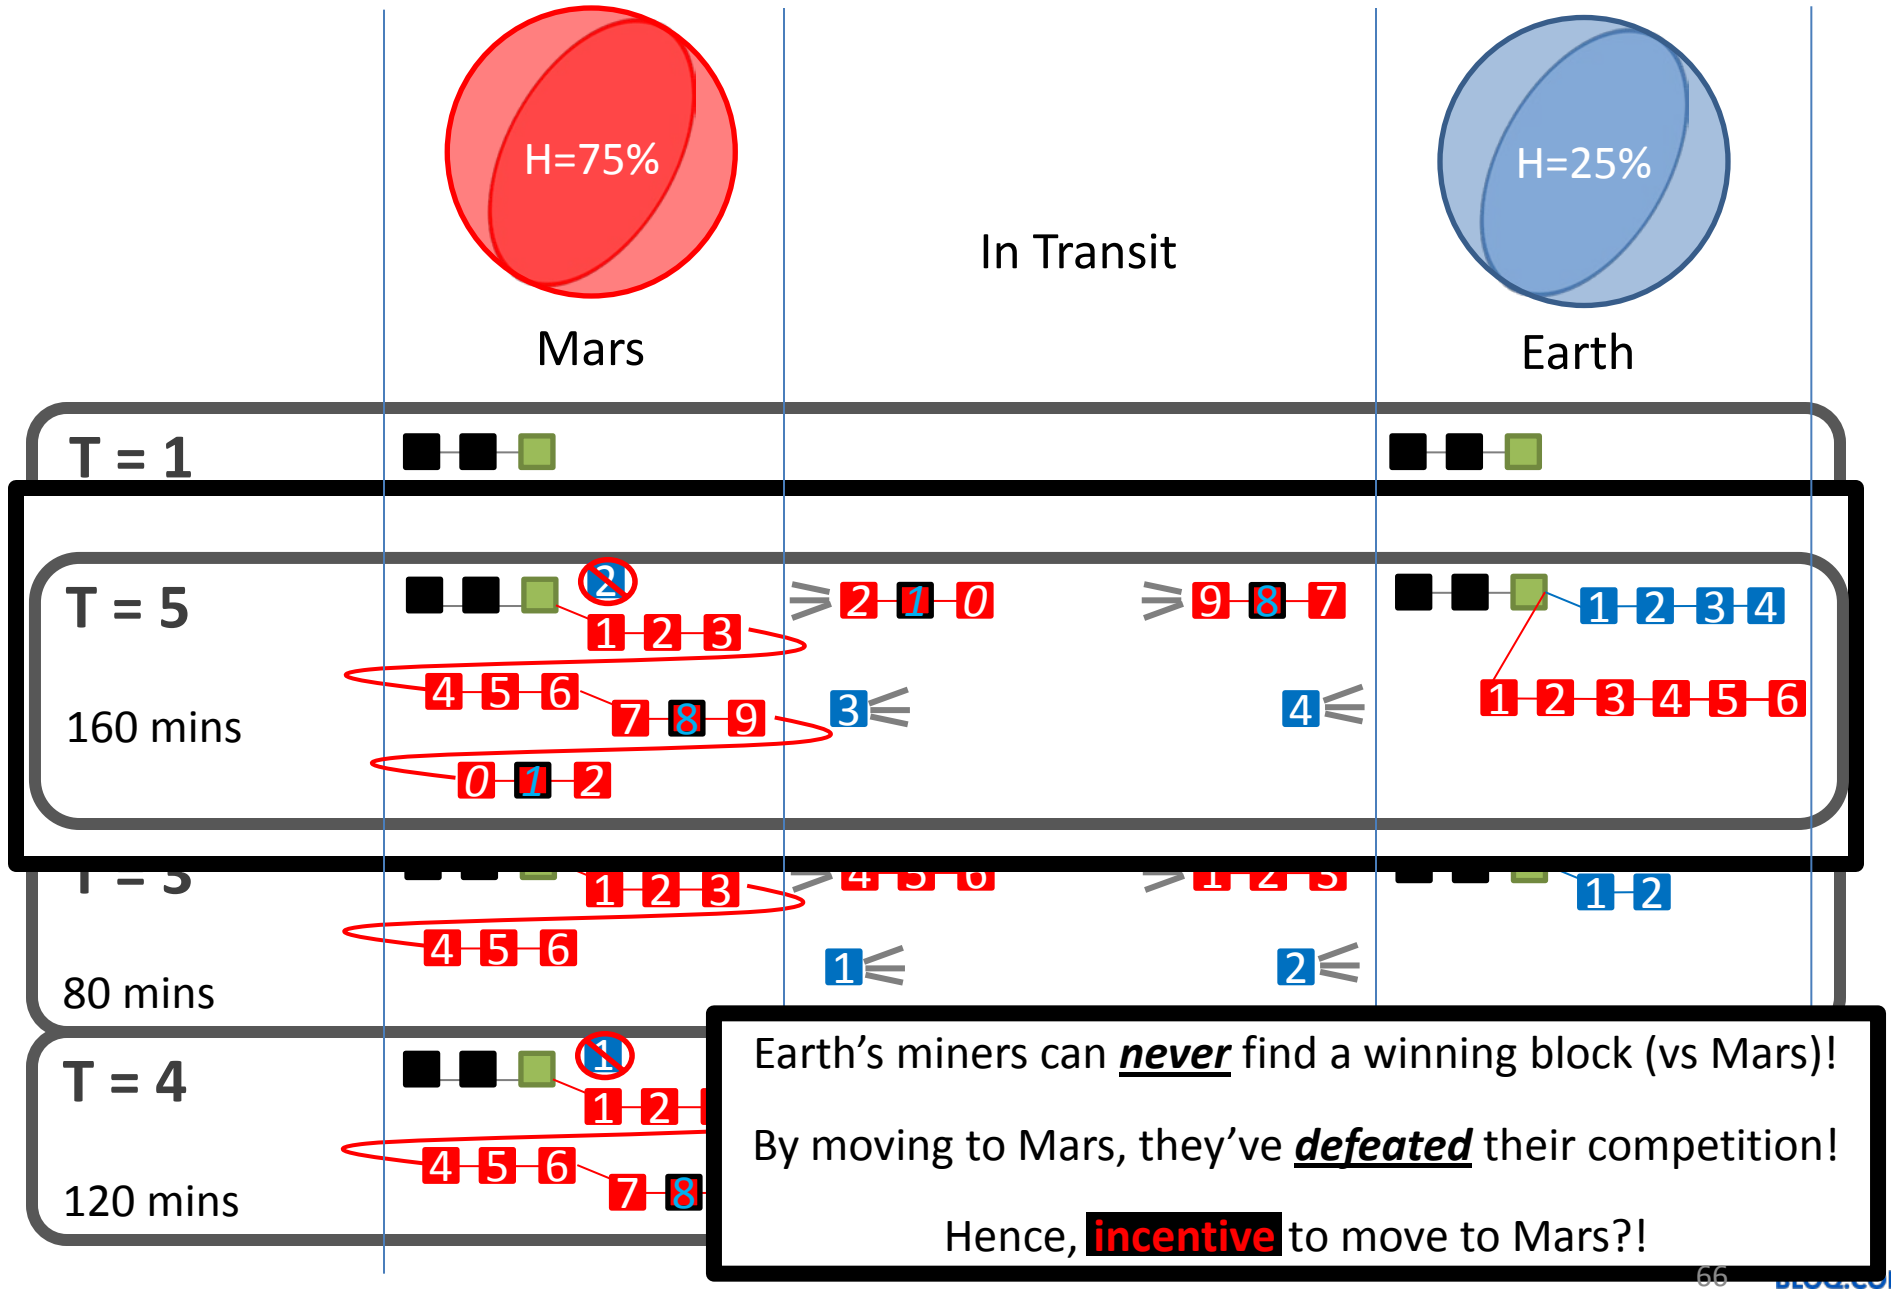

Incentive Glitch

- Miners have an incentive to make their connectivity *worse*.
- This “glitch” is principle behind all selfish mining variants (esp. where miners fill their own blocks), as well as the 51% attack.
- Also, heart of the Byzantine Generals problem. “I didn’t get that message.” (sincerely?).
- CoreDev complaints about “bandwidth”.
- **What’s going on? Can we fix this?**

Miners: The Dual-Role

- Miners **“sell”** blocks to the network...
- ...but who is “the network”? ...who’s *buying*?
- Miners, also!

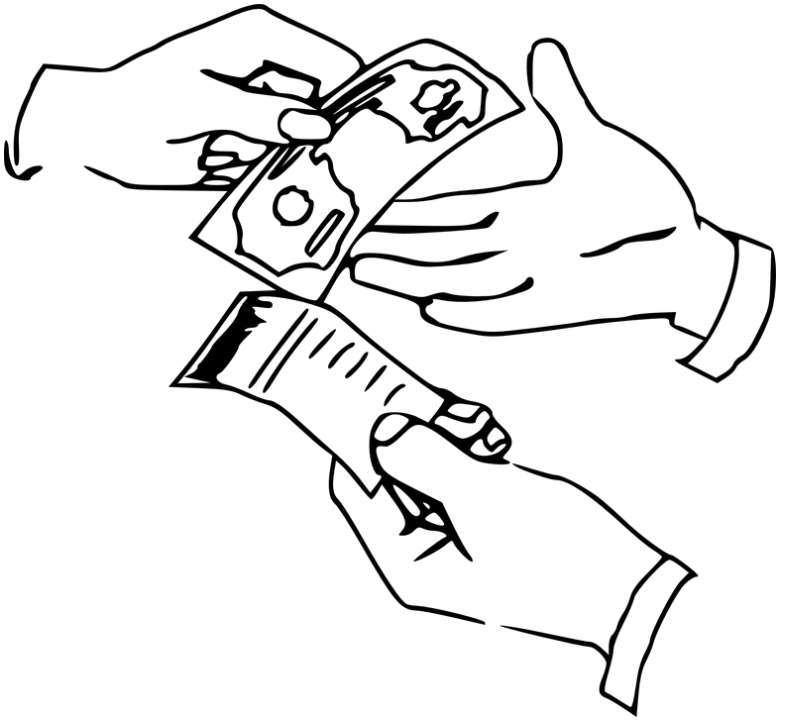

A block is **Bitcoin** if it is part of the heaviest chain, so ...
“Buyers” = Network = 51% of *the future hashrate*.

Orphaning “Costs” - Bifurcation

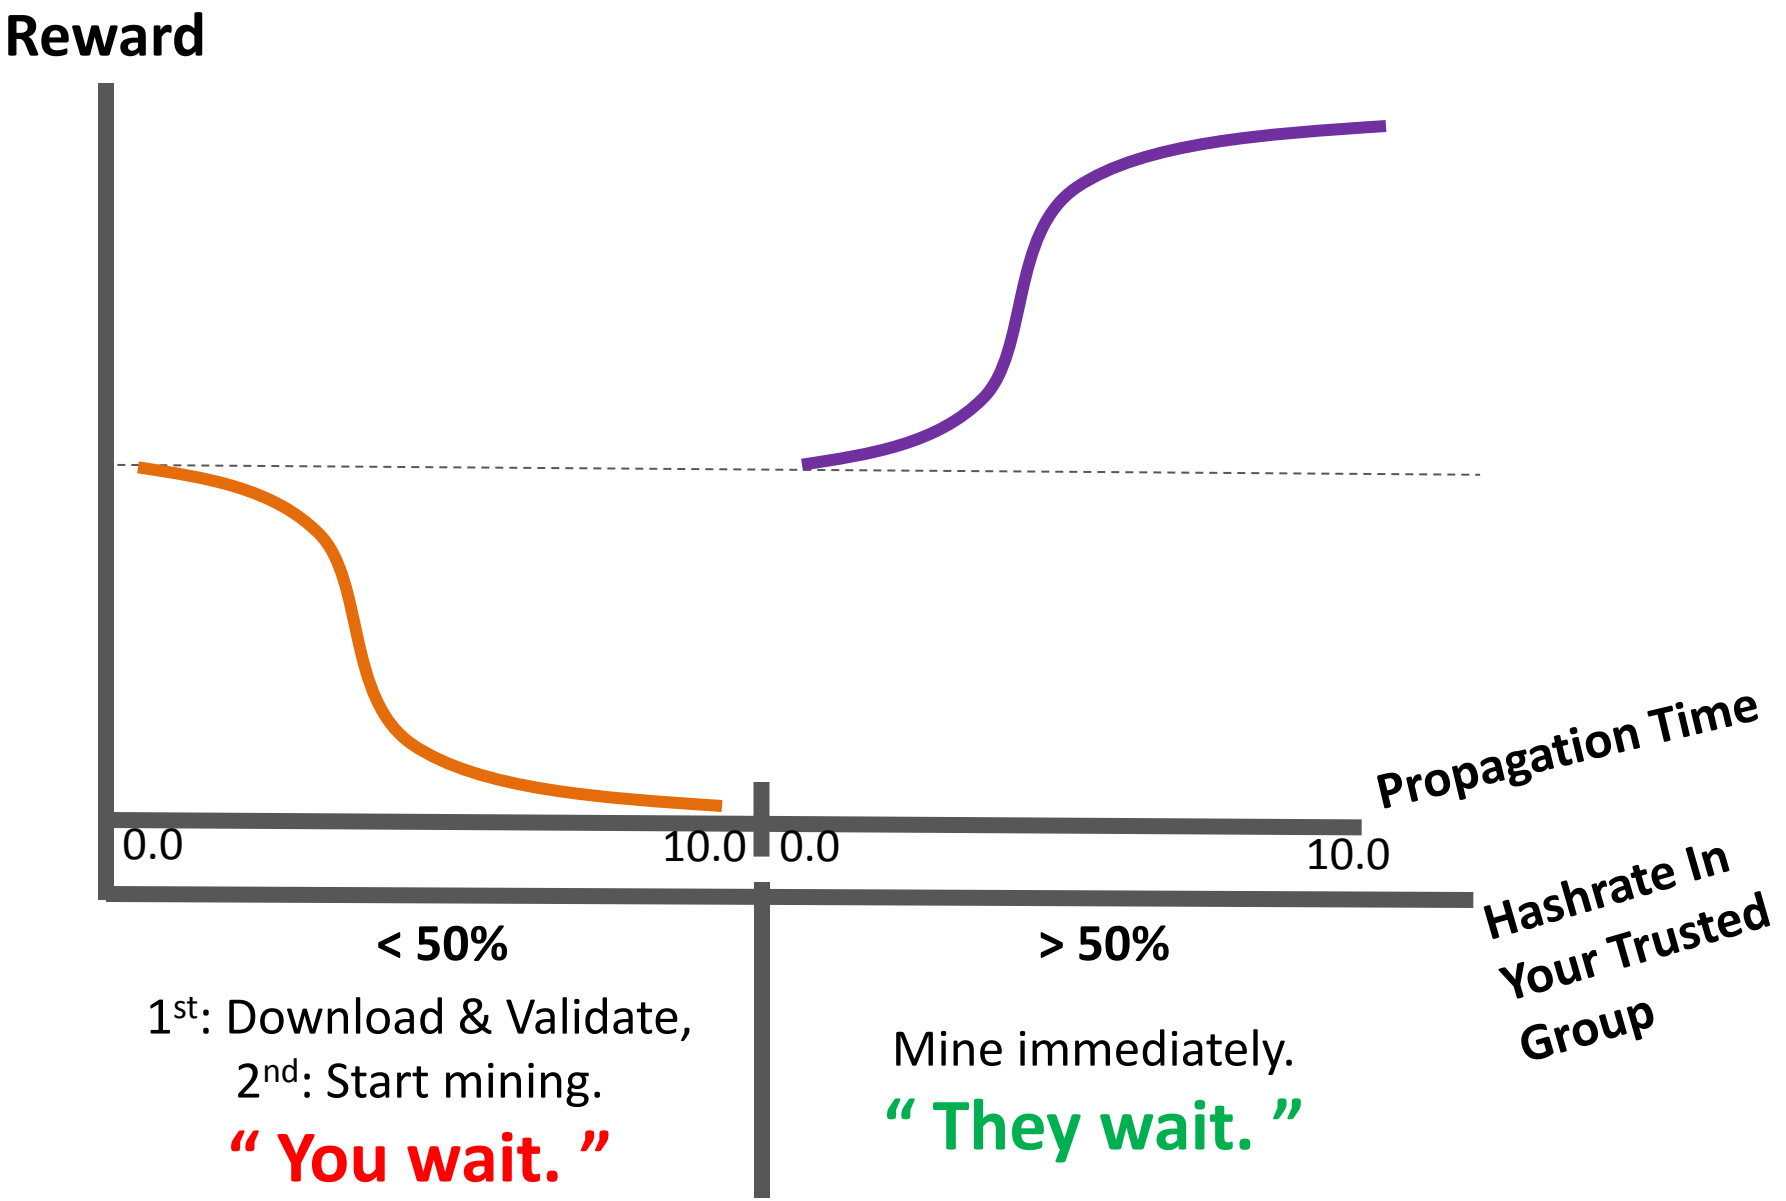

Other Research

A Transaction Fee Market Exists Without a Block Size Limit

Peter R†

August 4, 2015

The price *per byte*, ρ,[15] for the miner to produce a given quantity of block space follows by differentiating M_{supply} with respect to Q:

$$\rho_{\text{supply}}(Q) \equiv \frac{d}{dQ} M_{\text{supply}} = \frac{R}{T} \frac{d\tau}{dQ} e^{\frac{\tau(Q)}{T}}. \quad (8)$$

$$\rho_{\text{supply}}(Q) \approx \underbrace{\frac{1}{\gamma C} \frac{R}{T}}_{\text{Static approximation}} \overbrace{e^{\frac{Q}{\gamma C T}}}^{\text{Correction}}.$$

$\frac{1}{\gamma C}$: time it takes per uncompressed megabyte to propagate block solutions to the other miners,

Consistency

A Transaction Fee Market Exists Without a Block Size Limit

Peter R†
August 4, 2015

The price *per byte*, ρ,[15] for the miner to produce a given quantity of block space follows by differentiating M_{supply} with respect to Q:

$$\rho_{\text{supply}}(Q) \equiv \frac{d}{dQ} M_{\text{supply}} = \frac{R}{T} \frac{d\tau}{dQ} e^{\frac{\tau(Q)}{T}}. \quad (8)$$

$$\rho_{\text{supply}}(Q) \approx \underbrace{\frac{1}{\gamma C} \frac{R}{T}}_{\text{Static approximation}} \underbrace{e^{\frac{Q}{\gamma C T}}}_{\text{Correction}}.$$

Independent of Q

$\frac{1}{\gamma C}$ time it takes per uncompressed megabyte to propagate block solutions to the other miners,

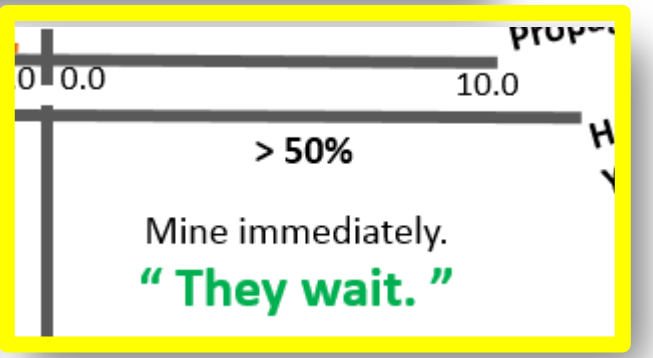

We made three important simplifying assumptions in this paper:

(2) in Sections 7 to 9, we assumed that this time parameter had a lower bound, in part, due to the capacity of the channels used to communicate the solutions and by the coding gain with which they could be compressed, as described by the Shannon-Hartley theorem,

(1) The time it takes to propagate information to the other miners is not in general constant across the network,[22] while the mempool is largely homogenous. This suggests that, assuming equal hashing costs, miners who can propagate their block solutions faster will earn a larger surplus. Relatedly, recent evidence also suggests that miners may begin mining prior to fully receiving and validating new blocks.[23] How do these phenomena affect the current analysis?
(2) Imagine the existence of a mining cartel, interconnected with high-capacity relay channels and committed to standardized mempool policies (to facilitate dense compression of block solutions). Such a cartel could greatly reduce the time required to propagate solutions to its other members. Do we expect such cartels to form and what might be their effect?
(3) When a miner accepts a transaction that increases the set of unspent outputs (UTXO), he takes on a liability equal to the present value of the cost of storing those new outputs indefinitely far into the future. Is a healthy fee market expected to emerge that charges users the true cost of expanding Bitcoin's UTXO set?

the analysis presented in this paper breaks down when the block reward falls to zero.

We made three important simplifying assumptions in this paper:

(2) in Sections 7 to 9, we assumed that this time parameter had a lower bound, in part, due to the capacity of the channels used to communicate the solutions and by the codi[No SPV]h which they could be compressed, as described by the Shannon-Hartley theorem,

(1) The time it takes to propagate information to the other miners is not in general constant across the network [22] while the mempool is largely homogenous. This [Pre-Propagation, "Scheduled Blocks"] osts, miners who can propagate their block solutions faster will earn a larger surplus. Relatedly, recent evidence also suggests that miners may begin mining prior to fully receiving and validating [SPV mining] How do these phenomena affect the current analysis?

(2) Imagine the existence of a mining cartel, interconnected with high-capacity relay channels and committed to standardized mempool policies (to facilitate dense co[Equilibrium Behavior]ons). d greatly reduce the time required to propagate solutions to its othe e expect such cartels to form and what might be their effect?

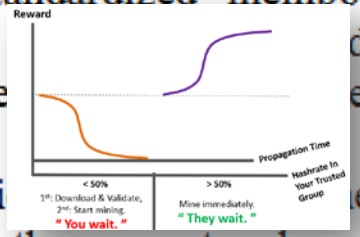

(3) When a miner accepts a transacti e set of unspent outputs (UTXO), he takes on a liability equal to the present value of the cost of storing those new [Externality Problem] itely far into the future. Is a healthy fee market expected to emerge that he true cost of expanding Bitcoin's UTXO set?

the analysis presented in this paper breaks down when the block reward falls to zero.

Fixing the Incentive Glitch

https://bitslog.wordpress.com/2016/01/08/spv-mining-is-the-solution-not-the-problem/

Words on Bitcoin Design, Privacy, Security and Crypto by Sergio Demian Lerner

BITSLOG

"SPV mining" is the solu

Is was Wednesday, March 26, 2014. It was a cold w
university campus. The room, full of people, was wa
Bitcoin: Andresen, Bonneau, Eyal, Maxwell, Miller,
talk was about "SPV Mining", but at that time that misleading name had not been coined.

A block header contains these fields:

Field	Purpose
Version	Block version number
hashPrevBlock	256-bit hash of the previous block header
hashMerkleRoot	256-bit hash based on all of the transactions in the block
Time	Current timestamp as seconds since 1970-01-01T00:00 UT
Bits	Current target in compact format
Nonce	32-bit number (starts at 0)

My point was simple, I had to convince everyone in the room that SPV Mining was not only inevitable (if not prevented by a softfork or hardfork) but was the solution to many of the Bitcoin protocol problems, including a future hard fork to scale Bitcoin. One of the problems with Bitcoin at that time was the monetary bias towards bigger miners, because bigger miners have a lower rate of orphan blocks (but not bigger mining pools, as the block withholding attack later proved).

Mining in the dark

Orphaning “Costs” - Bifurcation

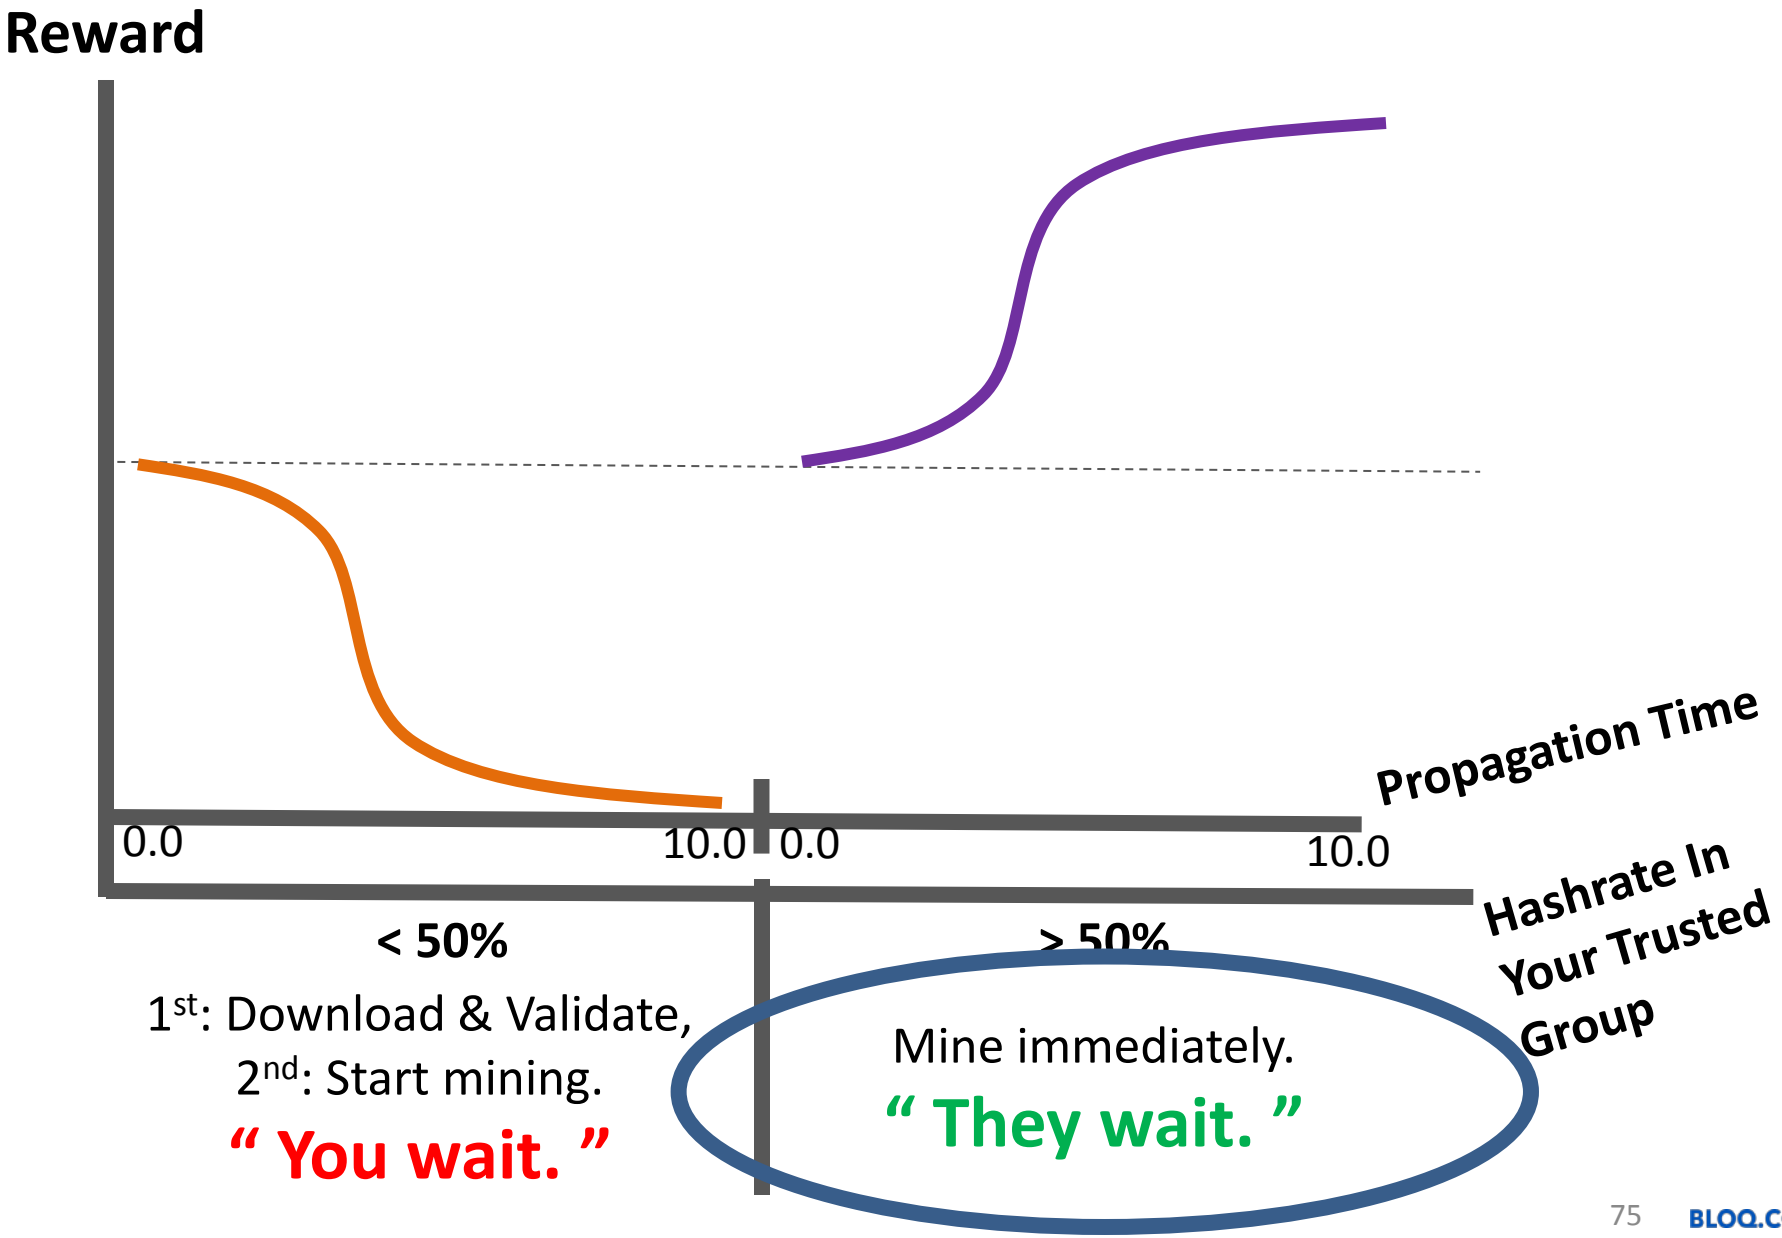

“SPV Mining”

- Proof-of-work is done on headers (not blocks)!
  - Headers are 80 bytes. Travel time near-instant.
  - Work(Invalid Header) = Work (Mine Valid BTC Block)
- Headers: expensive to fake + teleport instantly.

A block header contains these fields:

Field	Purpose
Version	Block version number
hashPrevBlock	256-bit hash of the previous block header
hashMerkleRoot	256-bit hash based on all of the transactions in the block
Time	Current timestamp as seconds since 1970-01-01T00:00 UT
Bits	Current target in compact format
Nonce	32-bit number (starts at 0)

- Strat: [1] Notice new header. [2] Mine on it.
  [3] Meanwhile, download and validate its block.
  [4] Post-validation, insert all (still valid) txns.

Space Ex. w Teleporting Headers (& Fast Coinbases)

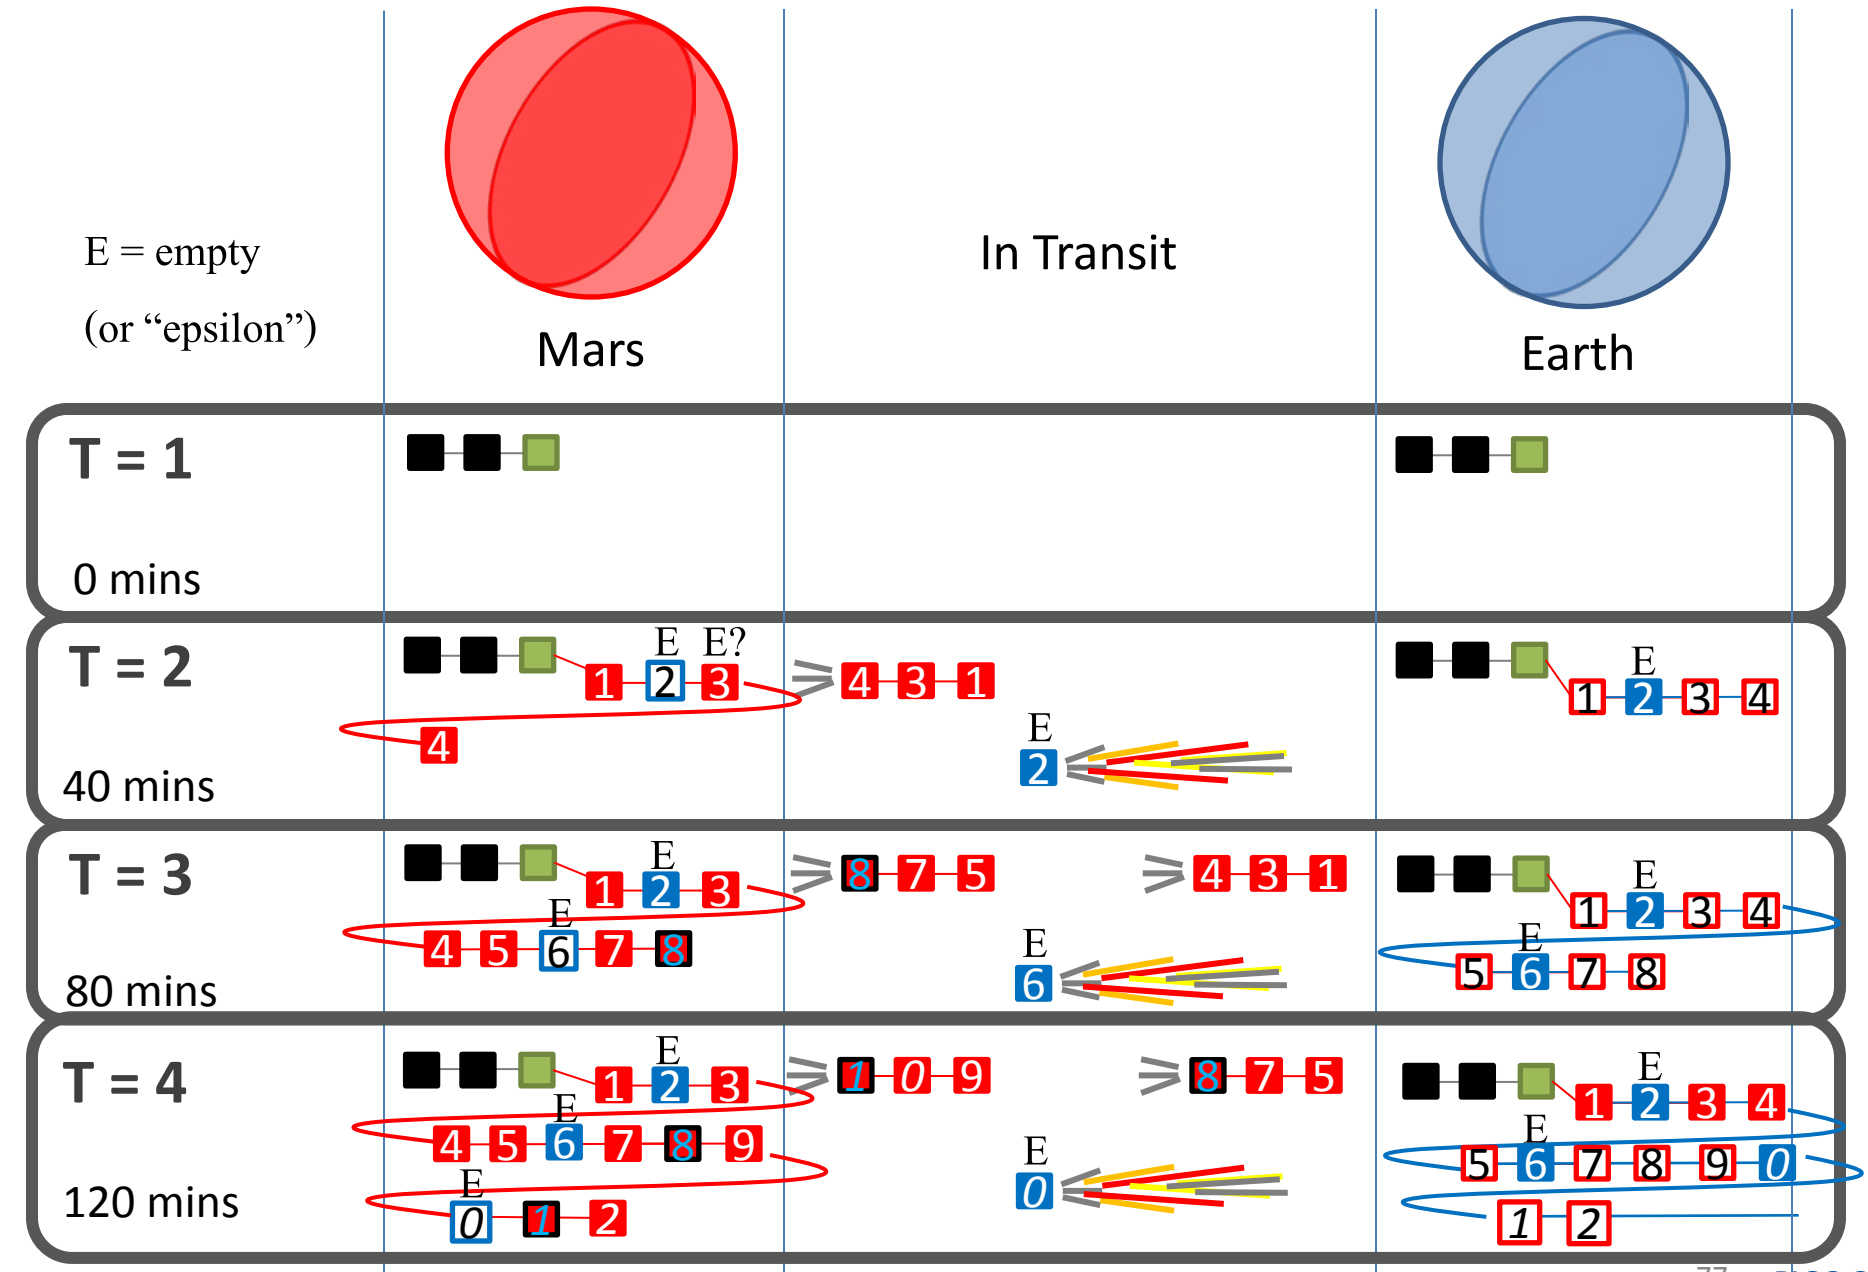

The Propagation Paradox: Connectivity Up, Profits…Down?!

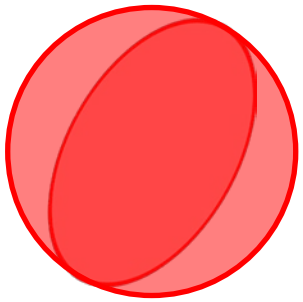

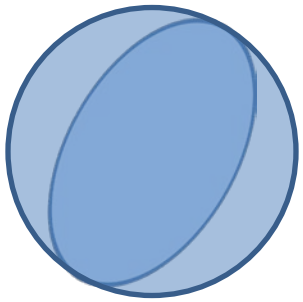

Planet:	Mars	Earth
Hashrate:	75%	25%
Txn Volume:	10%	90%

- 75% of hashrate teleported to Mars (where bandwidth is poor).
- It takes 1 hour for messages to pass between planets.
- (Mars miners are not necessarily coordinating).

For Subsidy (50, 25, 12.5) …

- …problem solved.
- Because:
  - Miners can now “mine a block”, at any time…
  - …but they can’t start *safely* including new txns, until they have latest txn data.

Incentive Problem: Fixed

- Should a **group of 26%** leave Mars to go back to Earth?

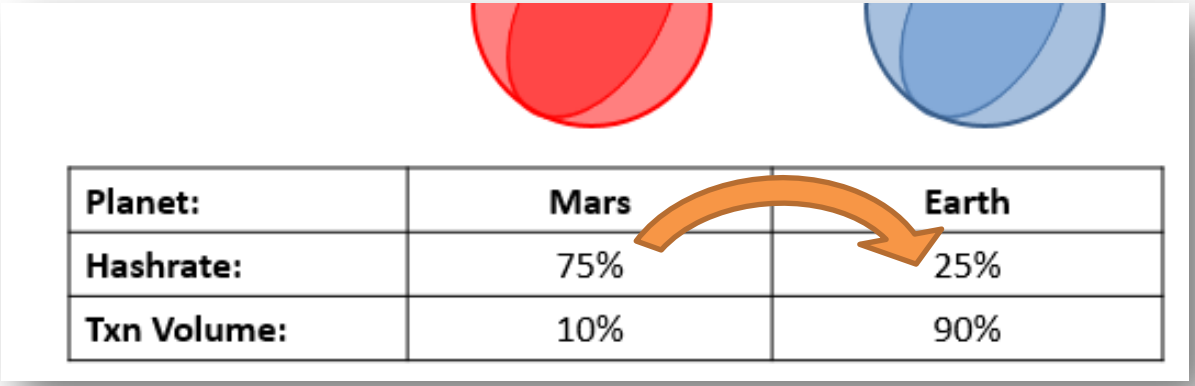

Planet:	Mars	Earth
Hashrate:	75%	25%
Txn Volume:	10%	90%

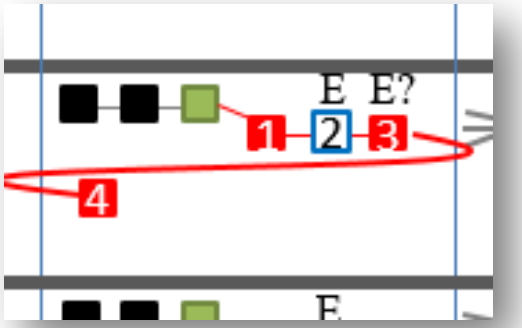

- [1] fewer "epsilon txns", [2] less e-tx competition.
- Stay: **10%** * (26/**75**) vs.
- Return: **90%** * (26/**51**)
- Incentive to…
  - …move to Earth.
  - …invest in connectivity.
  - …co-locate with tx-fees.

Fees drive miner incentives!

- Node Costs
  - Opt-In
  - Internalized
  - Anti-Fragile
- Mining
  - "There's only
  - How can diffe common net

Is SPV Mining “Safe”? (1 of 2)

- Shouldn’t we force miners to validate?

33 “SPV mining” is the solution, not the problem (bitslog.wordpress.com)
submitted 6 months ago by sangzou
9 comments share save hide give gold report

all 9 comments - sorted by: best

[–] luke-jr 5 points 6 months ago

Looks to me like this would destroy the *purpose* of mining in the process...

The safe way to "SPV mine" is:

1. Be able to switch back to the old block if the new one is determined to be invalid (*not possible* with current miners)
2. Never *publish* a block found SPV mining, until you have completed verification of the block it is built on top of (and it is determined to be valid).

BTCC's pool did not significantly harm Bitcoin during the BIP66 adoption despite their "SPV mining" *because* they implemented policy 2. Without policy 1, they continued to *waste hashpower* until the next block, but that mostly only affected their own income/relevance, and would be no different than if they had simply shut off their miners for that time period. With both, the Bitcoin network would actually benefit for the reasons laid out by Sergio, but without the problems of "SPV mining".

- (Don’t we, already?)

Is SPV Mining “Safe”? (2 of 2)

- Or, is validation the *node’s job*?
- If we aren’t running our own node, we are trusting someone else to validate for us.
- Isn’t that anathema to Bitcoin?

‘Bandwidth’ Will Continue to Improve

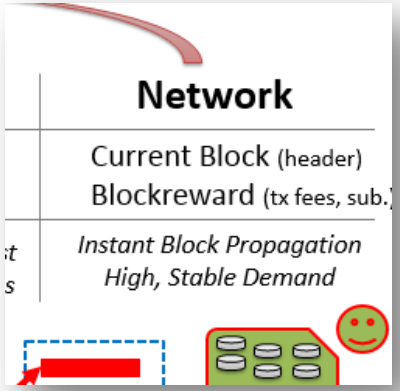

A block header contains these fields:

Field	Purpose
Version	Block version number
hashPrevBlock	256-bit hash of the previous block header
hashMerkleRoot	256-bit hash based on all of the transactions in the b
Time	Current timestamp as seconds since 1970-01-01T00
Bits	Current target in compact format
Nonce	32-bit number (starts at 0)

- **SPV Mining**
- Bloom Filters
- FIBRE
- UTXO Commitments
- **Block Scheduling***
- Clever Transaction Strategies

Clever SPV Optimization

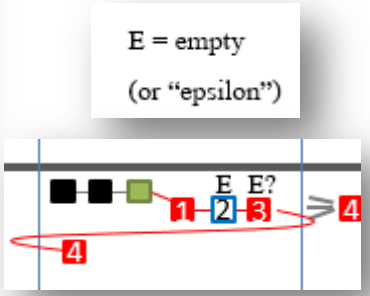

"Hipster Mining"

- Low Quality
- 'Unpopular' / Passed over by rival miners.

Canonical TX Priority

- It helps if miners agree on the definition of “transaction priority”. (Semi-scheduled blocks).

Priority transactions

Historically it was not required to include a fee for every transaction. A large portion of miners would mine transactions with no fee given that they had enough "priority". Today, low priority is mostly used as an indicator for spam transactions and almost all miners expect every transaction to include a fee.

Transaction priority is calculated as a value-weighted sum of input age, divided by transaction size in bytes:

```
priority = sum(input_value_in_base_units * input_age)/size_in_bytes
```

- Even crazier idea: pre-defining each upcoming block (across-block priority).

Conclusion

- We want to push bandwidth into the ‘full node costs’ category, but this is hard. Bandwidth is a distance *between* two people.
- Ideally, Miners would always have up-to-date information.

SPV Mining

More Work Needed, But Promising Solutions Exist

A block header contains these fields:

Field	Purpose
Version	Block version number
hashPrevBlock	256-bit hash of the previous block header
hashMerkleRoot	256-bit hash based on all of the transactions in the block
Time	Current timestamp as seconds since 1970-01-01T00:00 UT
Bits	Current target in compact format
Nonce	32-bit number (starts at 0)

Conclusion

- We … e ‘full noc … d. Ban … o people.
- Idea … p-to-date info …

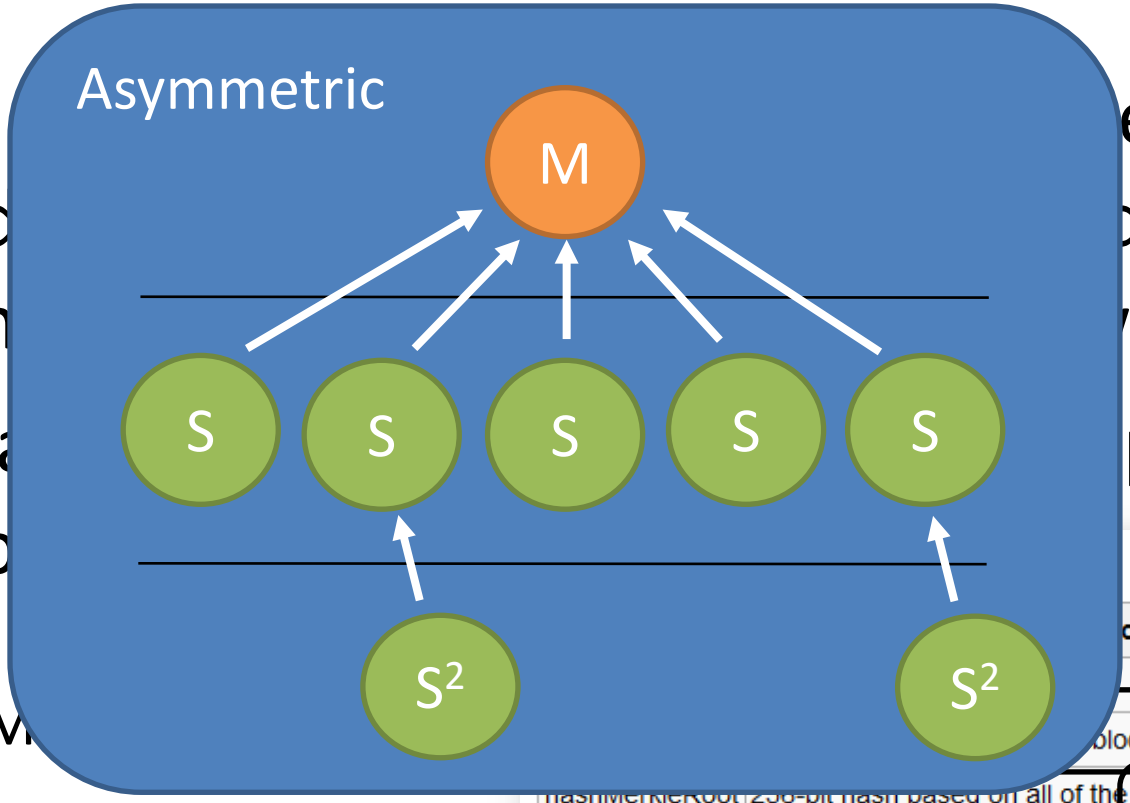

SPV M…

	…ose
	…block header
hashMerkleRoot	256-bit hash based on all of the transactions in the block
Time	Current timestamp as seconds since 1970-01-01T00:00 UT
Bits	Current target in compact format
Nonce	32-bit number (starts at 0)

More Work Needed, But

If bandwidth is to be increased, it *might* be safest to do so via sidechains.

Part 4 - Fees

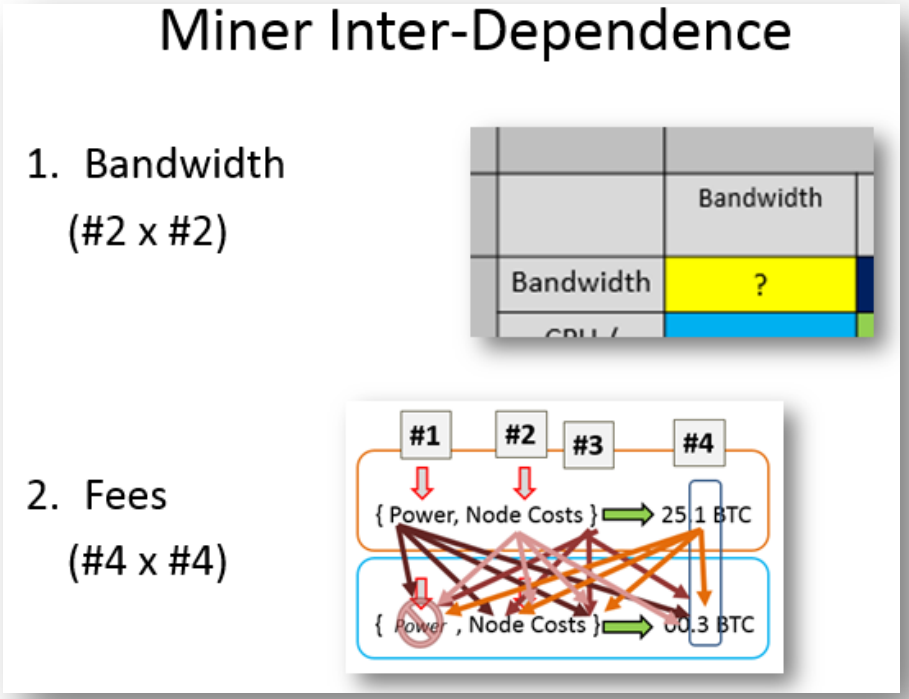

1. Under what conditions will revenues fall?
2. How can we prevent this?

Tx Fees – Beyond the Limits

- Node Costs
  - Opt-In
  - Internalized
  - Anti-Fragile
- Mining
  - "There's only
  - How can diffe
    common netv

"Tame" Sidechains

- Add new features to Bitcoin.
- Contracts are firewalled – opt in, and don't affect each other.
- Contracts are managed, to maximize BTC value.

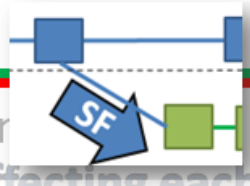

"Aggressive" Sidechains

- All the benefits of "Tame"...and:
- ***Permissionless Innovation***.
- No debate over precise "split" of validation resources (MB, SigOps).

What happens to mining, on mainchain Bitcoin?

Section Agenda

- Review / Assumptions (4)

1. Three Different Perspectives (5)
2. Bitcoin's Transition (to Equilibrium) (3)
3. The "Market" for Block-Access (15)
4. Coase vs Folk -- Miner Coordination (9)
5. Demand Curve Calculus (8)

An ***OLD*** Topic

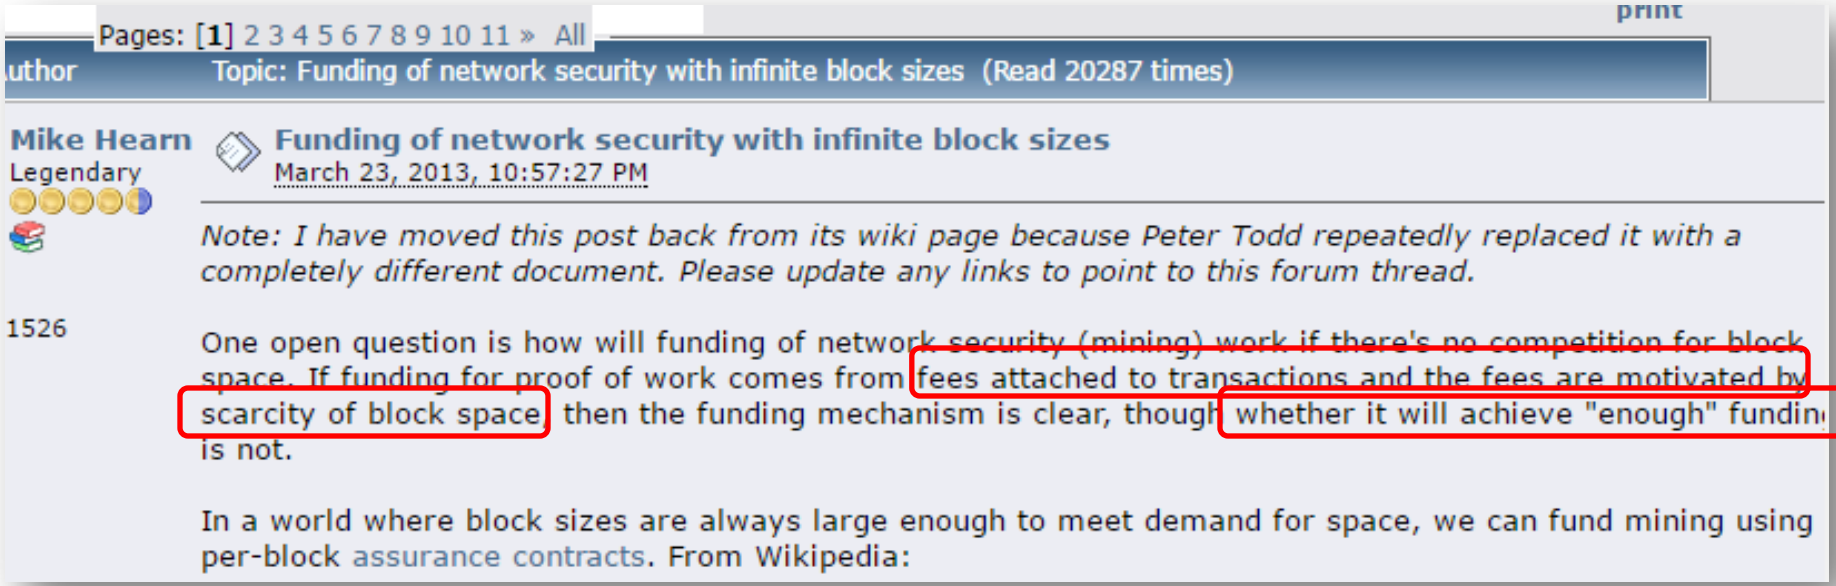

Pages: [1] 2 3 4 5 6 7 8 9 10 11 » All

print

uthor Topic: Funding of network security with infinite block sizes (Read 20287 times)

Mike Hearn
Legendary

Funding of network security with infinite block sizes
March 23, 2013, 10:57:27 PM

Note: I have moved this post back from its wiki page because Peter Todd repeatedly replaced it with a completely different document. Please update any links to point to this forum thread.

1526

One open question is how will funding of network security (mining) work if there's no competition for block space. If funding for proof of work comes from fees attached to transactions and the fees are motivated by scarcity of block space then the funding mechanism is clear, though whether it will achieve "enough" fundin is not.

In a world where block sizes are always large enough to meet demand for space, we can fund mining using per-block assurance contracts. From Wikipedia:

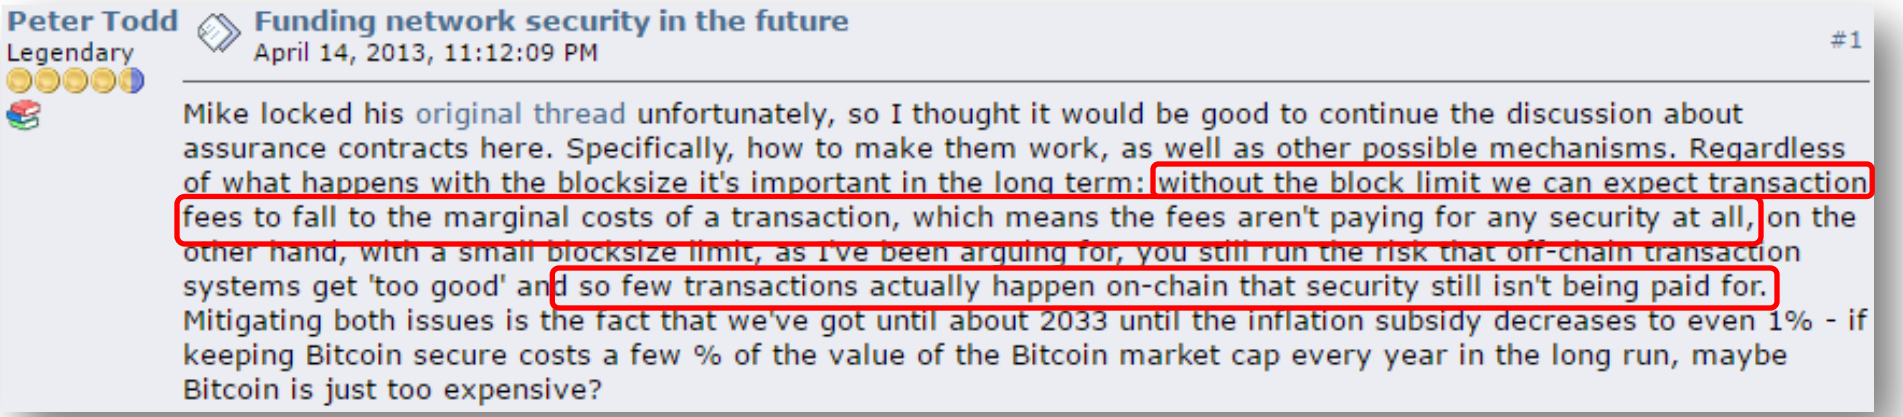

Peter Todd
Legendary

Funding network security in the future
April 14, 2013, 11:12:09 PM

#1

Mike locked his original thread unfortunately, so I thought it would be good to continue the discussion about assurance contracts here. Specifically, how to make them work, as well as other possible mechanisms. Regardless of what happens with the blocksize it's important in the long term: without the block limit we can expect transaction fees to fall to the marginal costs of a transaction, which means the fees aren't paying for any security at all, on the other hand, with a small blocksize limit, as I've been arguing for, you still run the risk that off-chain transaction systems get 'too good' and so few transactions actually happen on-chain that security still isn't being paid for. Mitigating both issues is the fact that we've got until about 2033 until the inflation subsidy decreases to even 1% - if keeping Bitcoin secure costs a few % of the value of the Bitcoin market cap every year in the long run, maybe Bitcoin is just too expensive?

(Worst Case) Assumptions

- That someone creates a sidechain of Bitcoin, with all economically-relevant limits removed.
- The sidechain txns are perceived as **perfect substitutes** for Bitcoin txns.

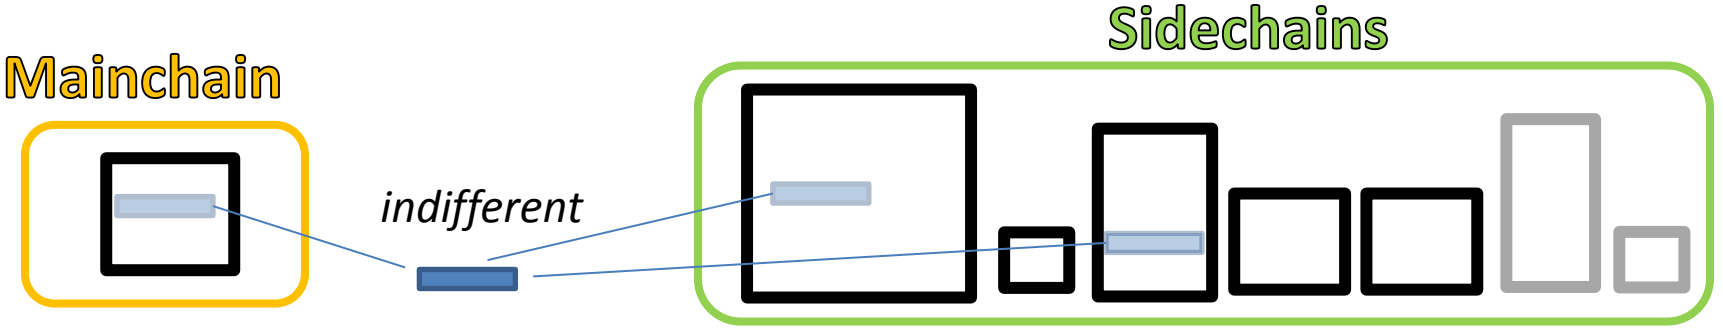

- This effectively removes Bitcoin's limits.
- Reminder: Bitcoin Core not *directly* affected.

Fees = Security

https://bitcoin.org/bitcoin.pdf

6. Incentive

By convention, the first transaction in a block is a special transaction that starts a new coin owned by the creator of the block. This adds an incentive for nodes to support the network, and provides a way to initially distribute coins into circulation, since there is no central authority to issue them. The steady addition of a constant of amount of new coins is analogous to gold miners expending resources to add gold to circulation. In our case, it is CPU time and electricity that is expended.

The incentive can also be funded with transaction fees. If the output value of a transaction is less than its input value, the difference is a transaction fee that is added to the incentive value of the block containing the transaction. Once a predetermined number of coins have entered circulation, the incentive can transition entirely to transaction fees and be completely inflation free.

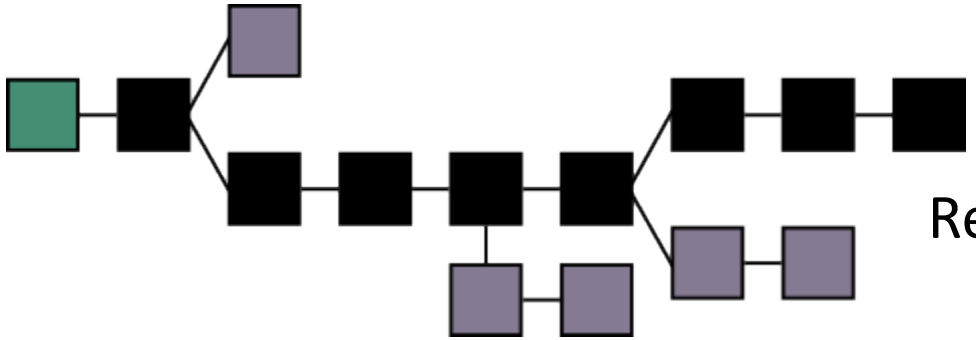

Rewrite Cost = $\text{Subsidy} + \sum Fees$

DoS Cost = $\sum Fees$

(Alternative) Assumptions

Characteristic	SC Demand Affects Mainchain...	Total Ecosystem Effect
Substitutes	...by **removing** need for a BTC tx.	?
Independent	...(not at all).	**Good** – Increases Total Fees.
Compliments	...by **inducing** need for a BTC tx.	**Great !** – Increases Fees 2x.

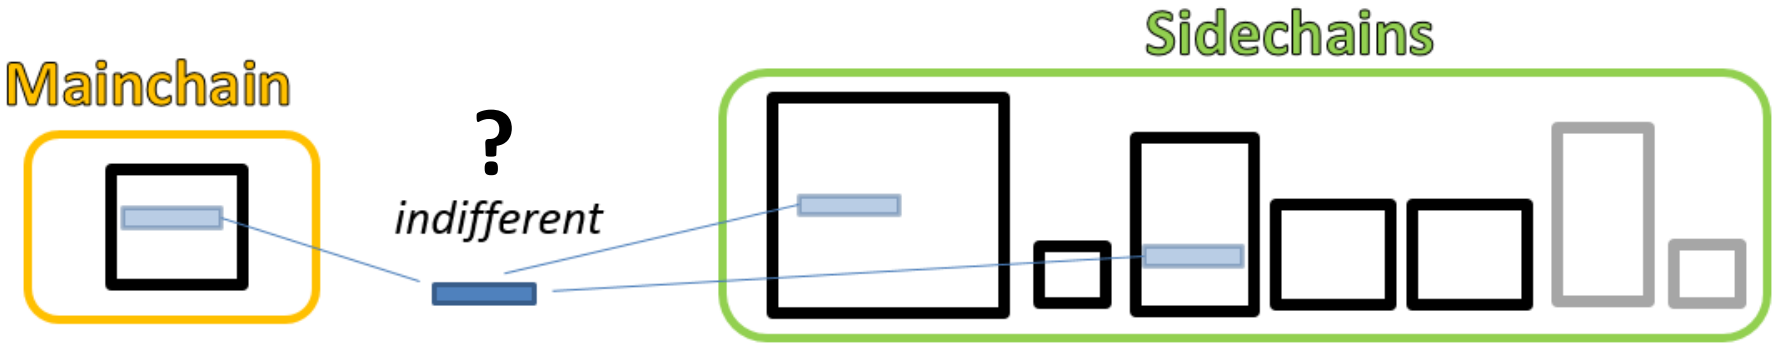

Section Agenda

- Review / Assumptions (Completed)

1. Three Different Perspectives (5)
2. Bitcoin's Transition (to Equilibrium) (3)
3. The "Market" for Block-Access (15)
4. Coase vs Folk -- Miner Coordination (9)
5. Demand Curve Calculus (8)

Economic Limits

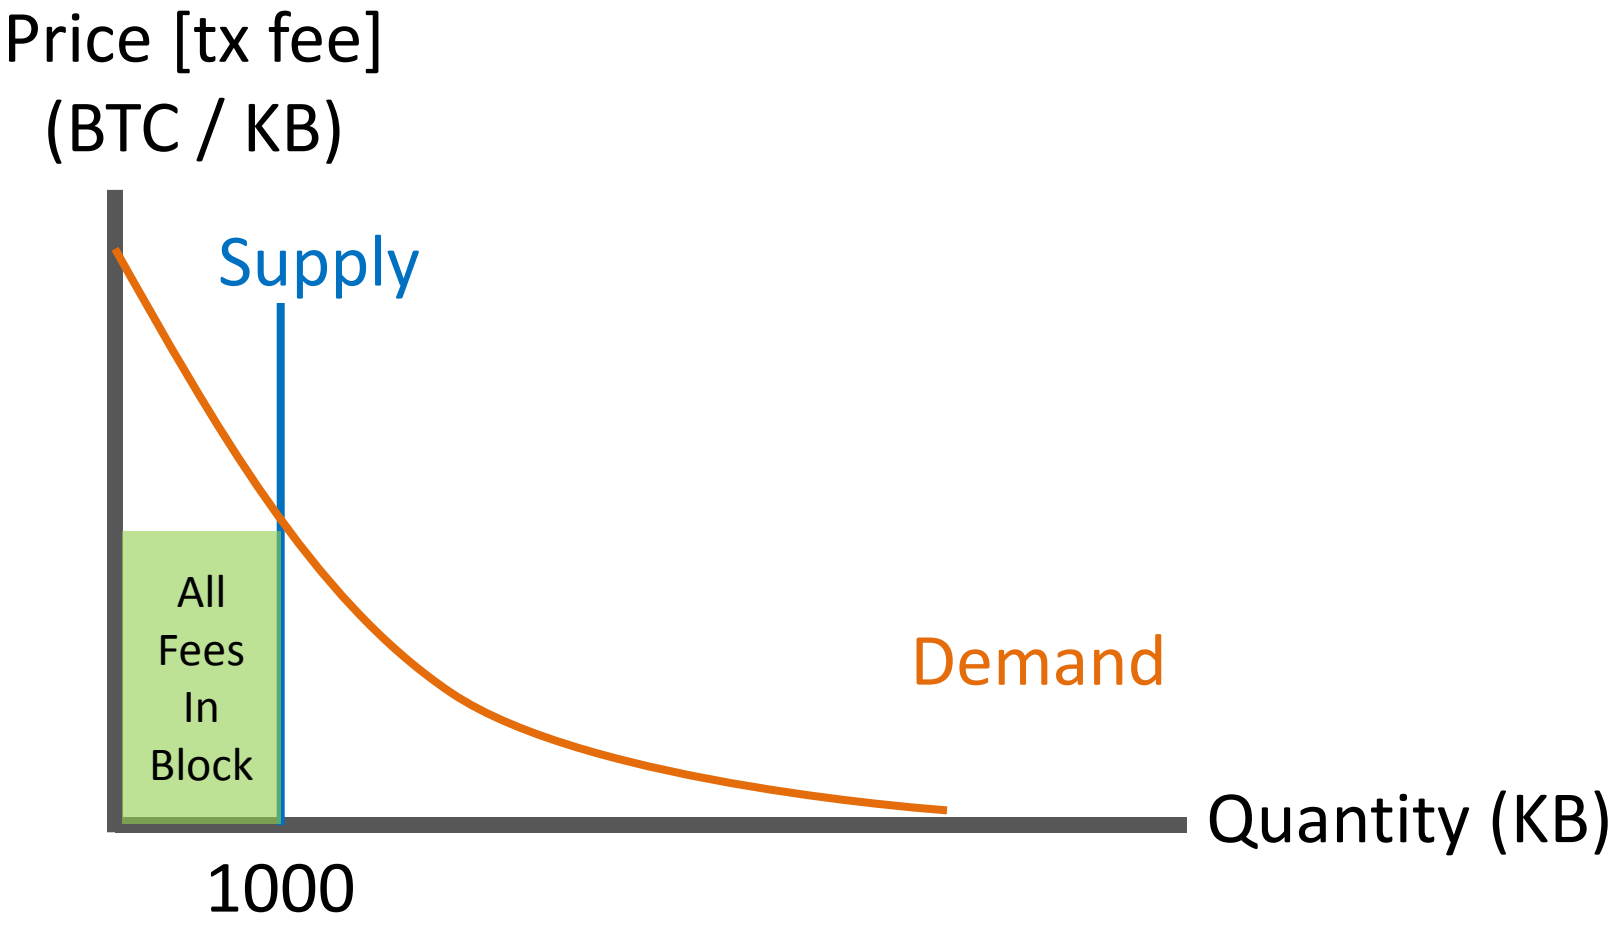

Economic Limits

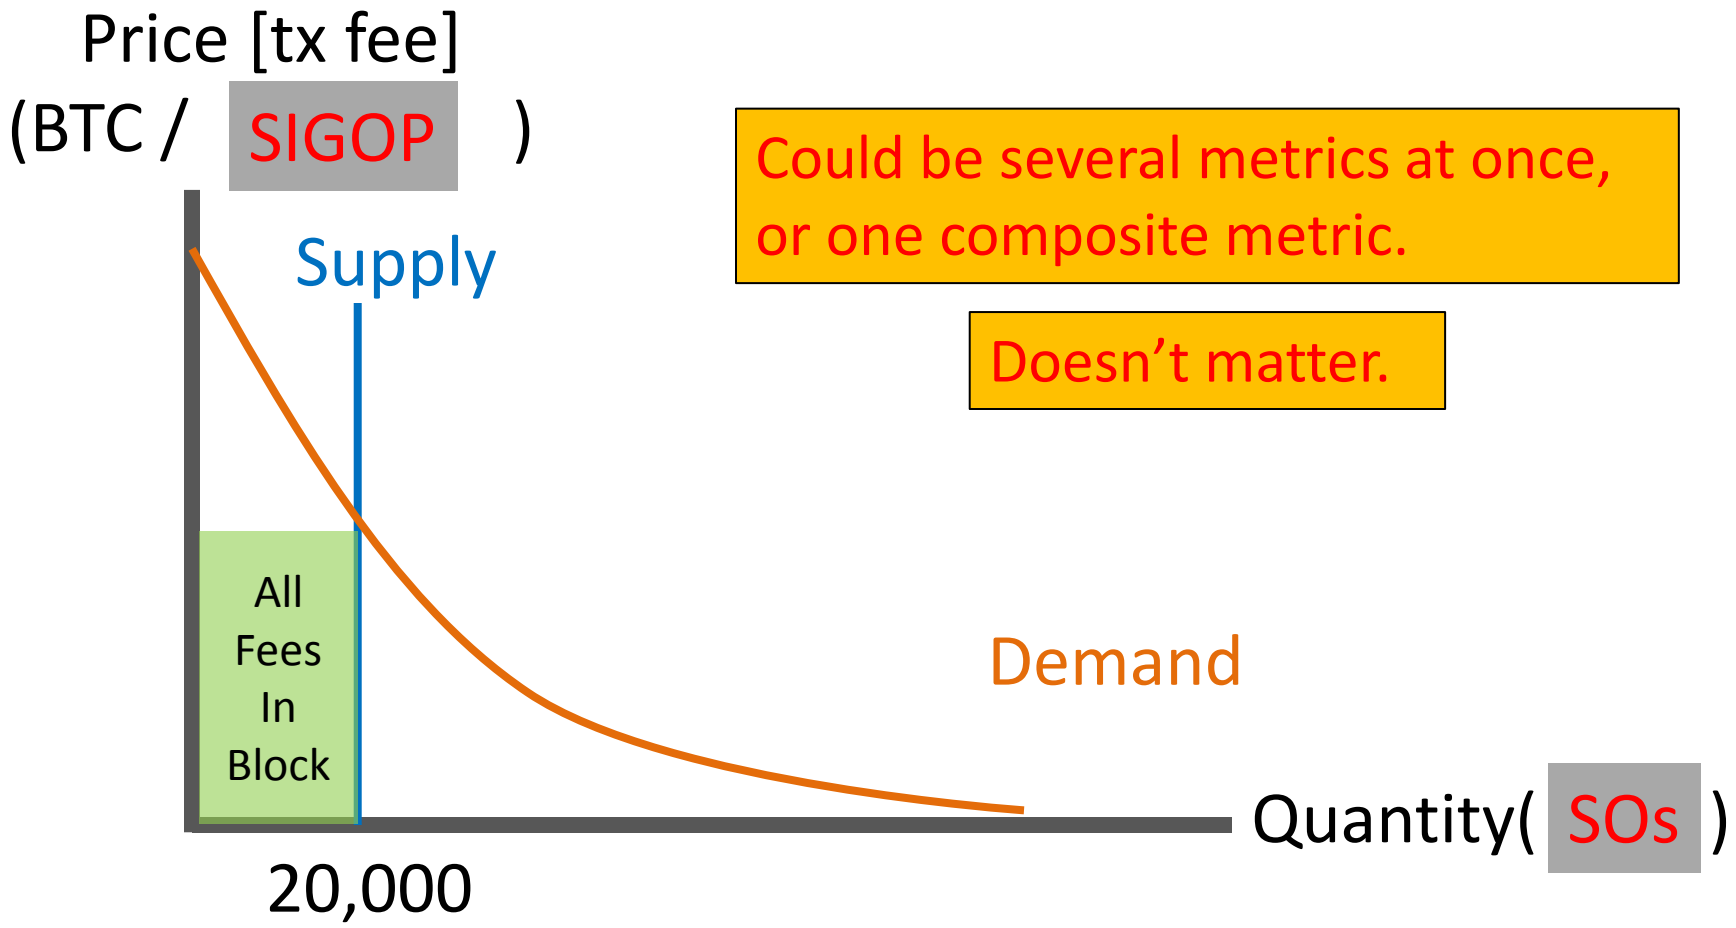

Could be several metrics at once, or one composite metric.

Doesn't matter.

3 Perspectives – Life Without Economic Limits

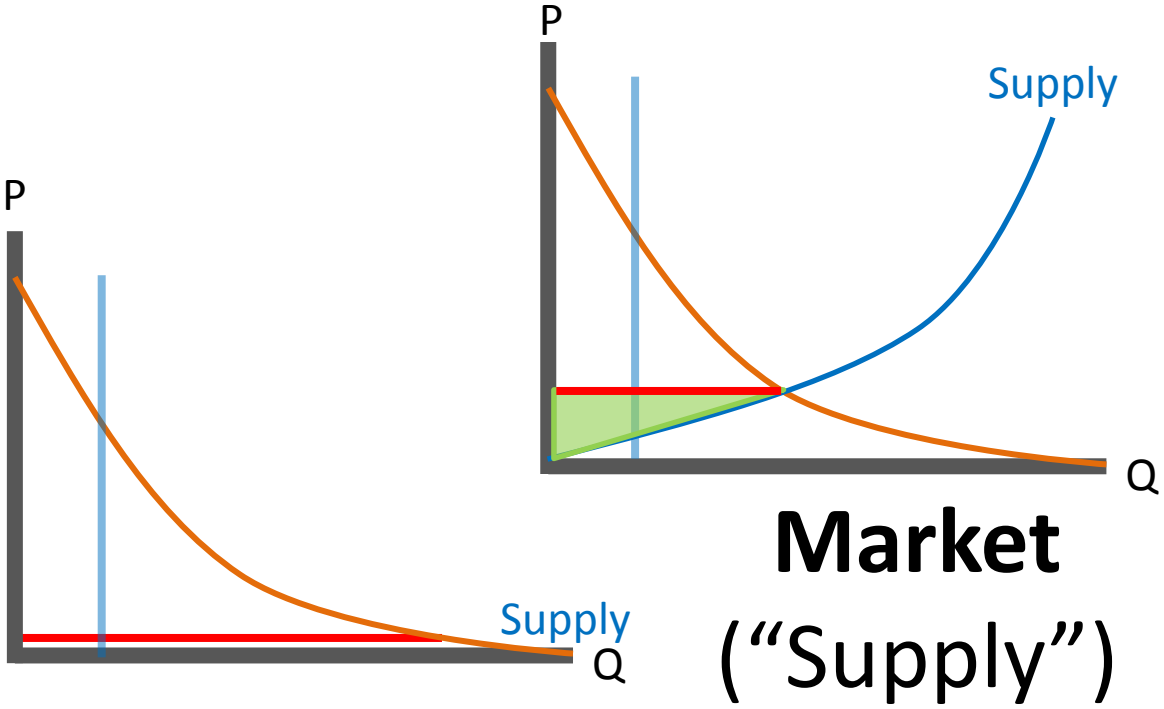

Market
("Supply")

Tragic
("Commons", "Anarchist")

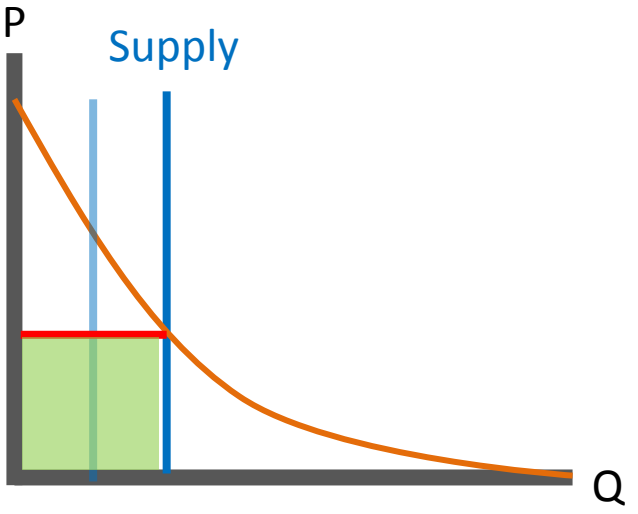

Cartel
("Managed", "Gouging")

3 Perspectives – Life Without Economic Limits

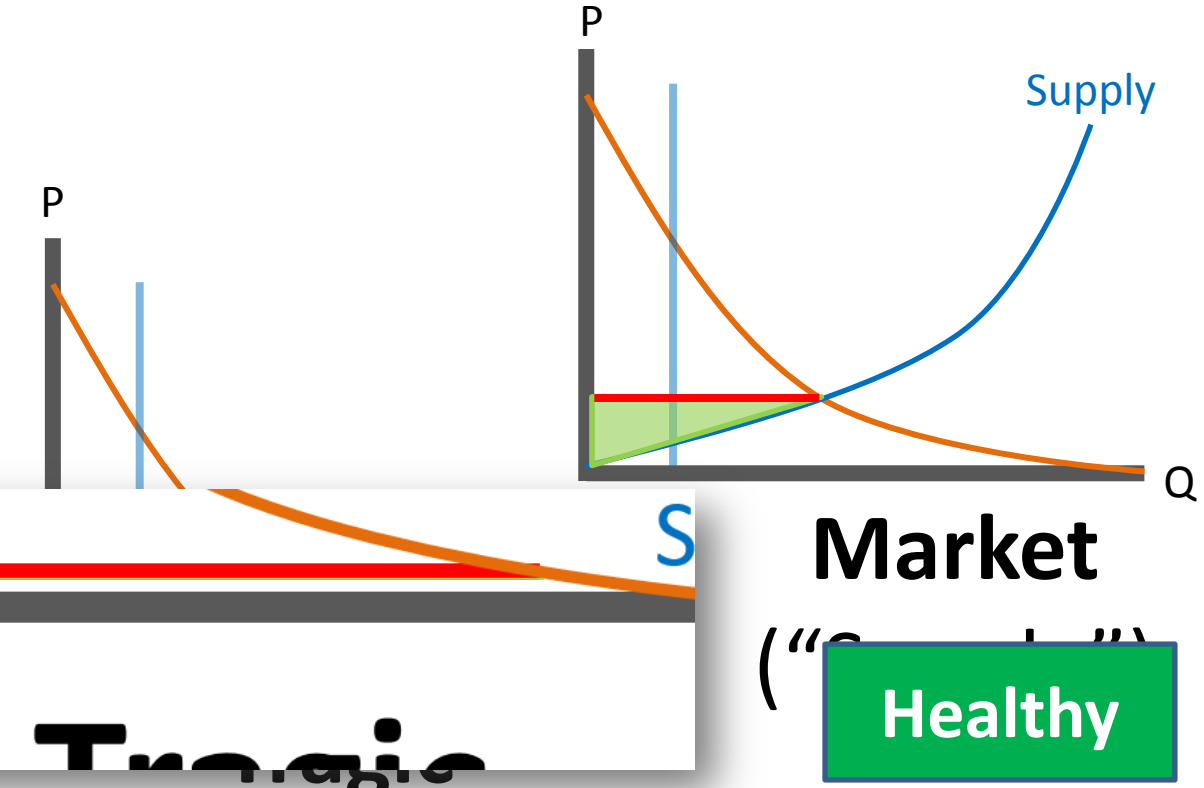

3 Perspectives – Life Without Economic Limits

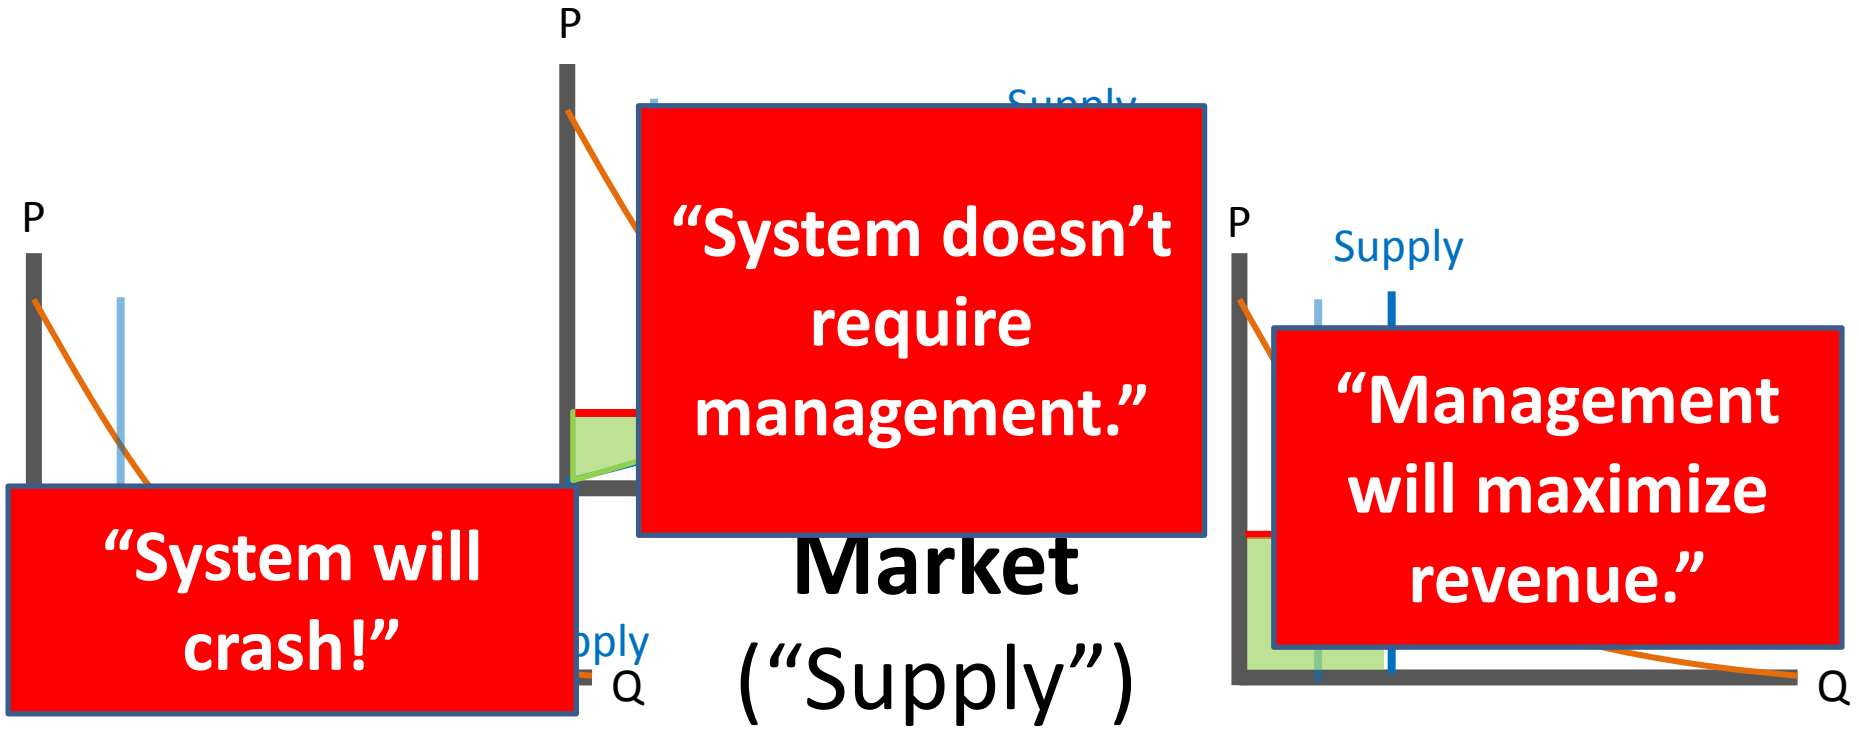

Tragic
("Commons", "Anarchist")

Market
("Supply")

Cartel
("Managed", "Gouging")

[2] Transitioning To Equilibrium

- Early Days are Not Representative

6. Incentive

By convention, the first transaction in a block is a special transaction that starts a new coin owned by the creator of the block. This adds an incentive for nodes to support the network, and provides

The incentive can also be funded with transaction fees. If the output value of a transaction is less than its input value, the difference is a transaction fee that is added to the incentive value of the block containing the transaction. Once a predetermined number of coins have entered circulation, the incentive can transition entirely to transaction fees and be completely inflation free.

– Fees ↑ , Subsidy ↓ . Implies: fee/sub ↑↑ .

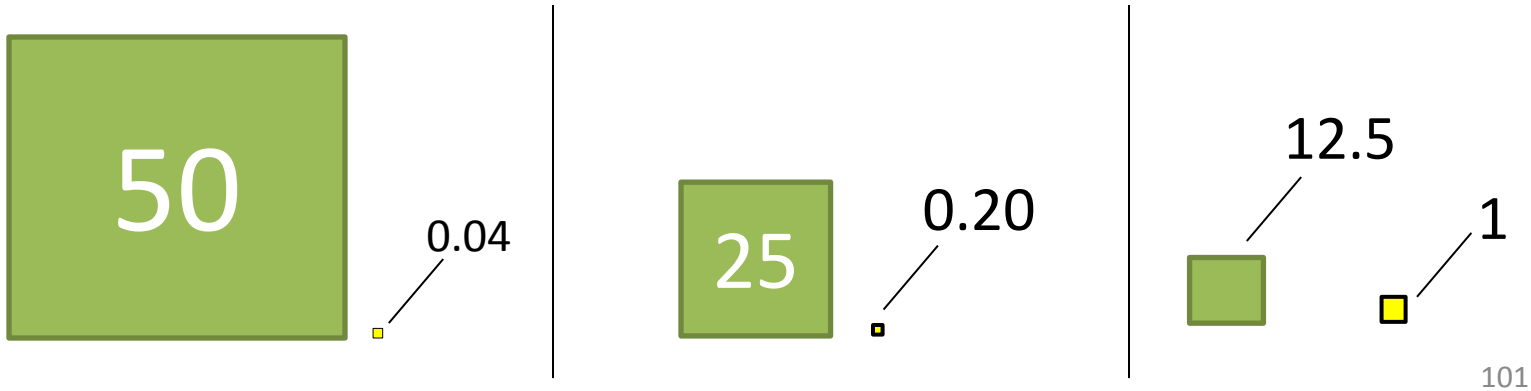

[2] Transition to Equilibrium

- What will change?
  - Miners will treasure blocks…for their Fees.
  - Growth Attitude will diminish, in favor of a (new) Adversarial Attitude.
  - Demand for BTC will increase and change (Hobby → Black Market → Home Network)
- How Different will it be?

[2] Transition to Equilibrium

- Conclusion: This will lead ***miners*** to engage in **“revenue (fee) maximization”**…
- …which will lead ***users*** to engage in **“fee minimization”**.

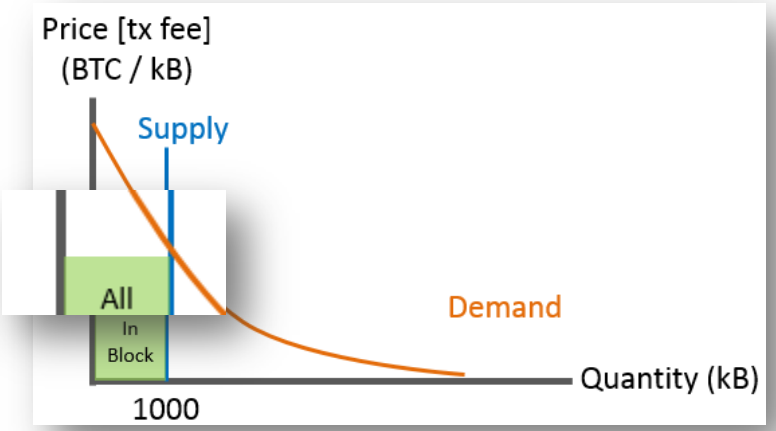

- Question: to *what extent* can miners fee-maximize?

[3] The “Market” for Block-Access

- [3] Sub-Agenda
  i. A “Normal” Market (6)
  ii. Abnormalities in Our Case (2)
  iii. Implications of these Abnormalities (4)
  iv. The Block Tree (3)

Preview: Per *Single* Block

(Across *Many* Blocks = Complicated)

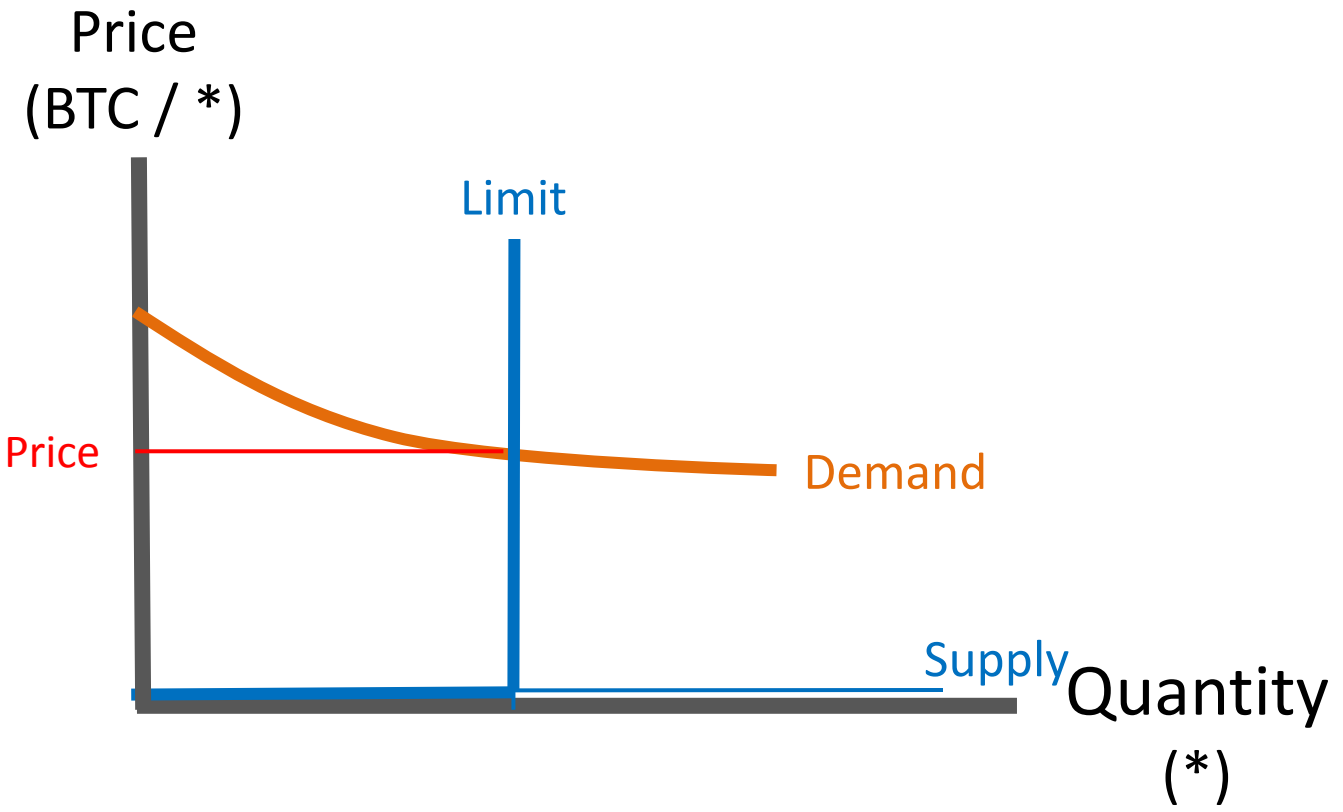

[i] Normal Market – Demand for Oil

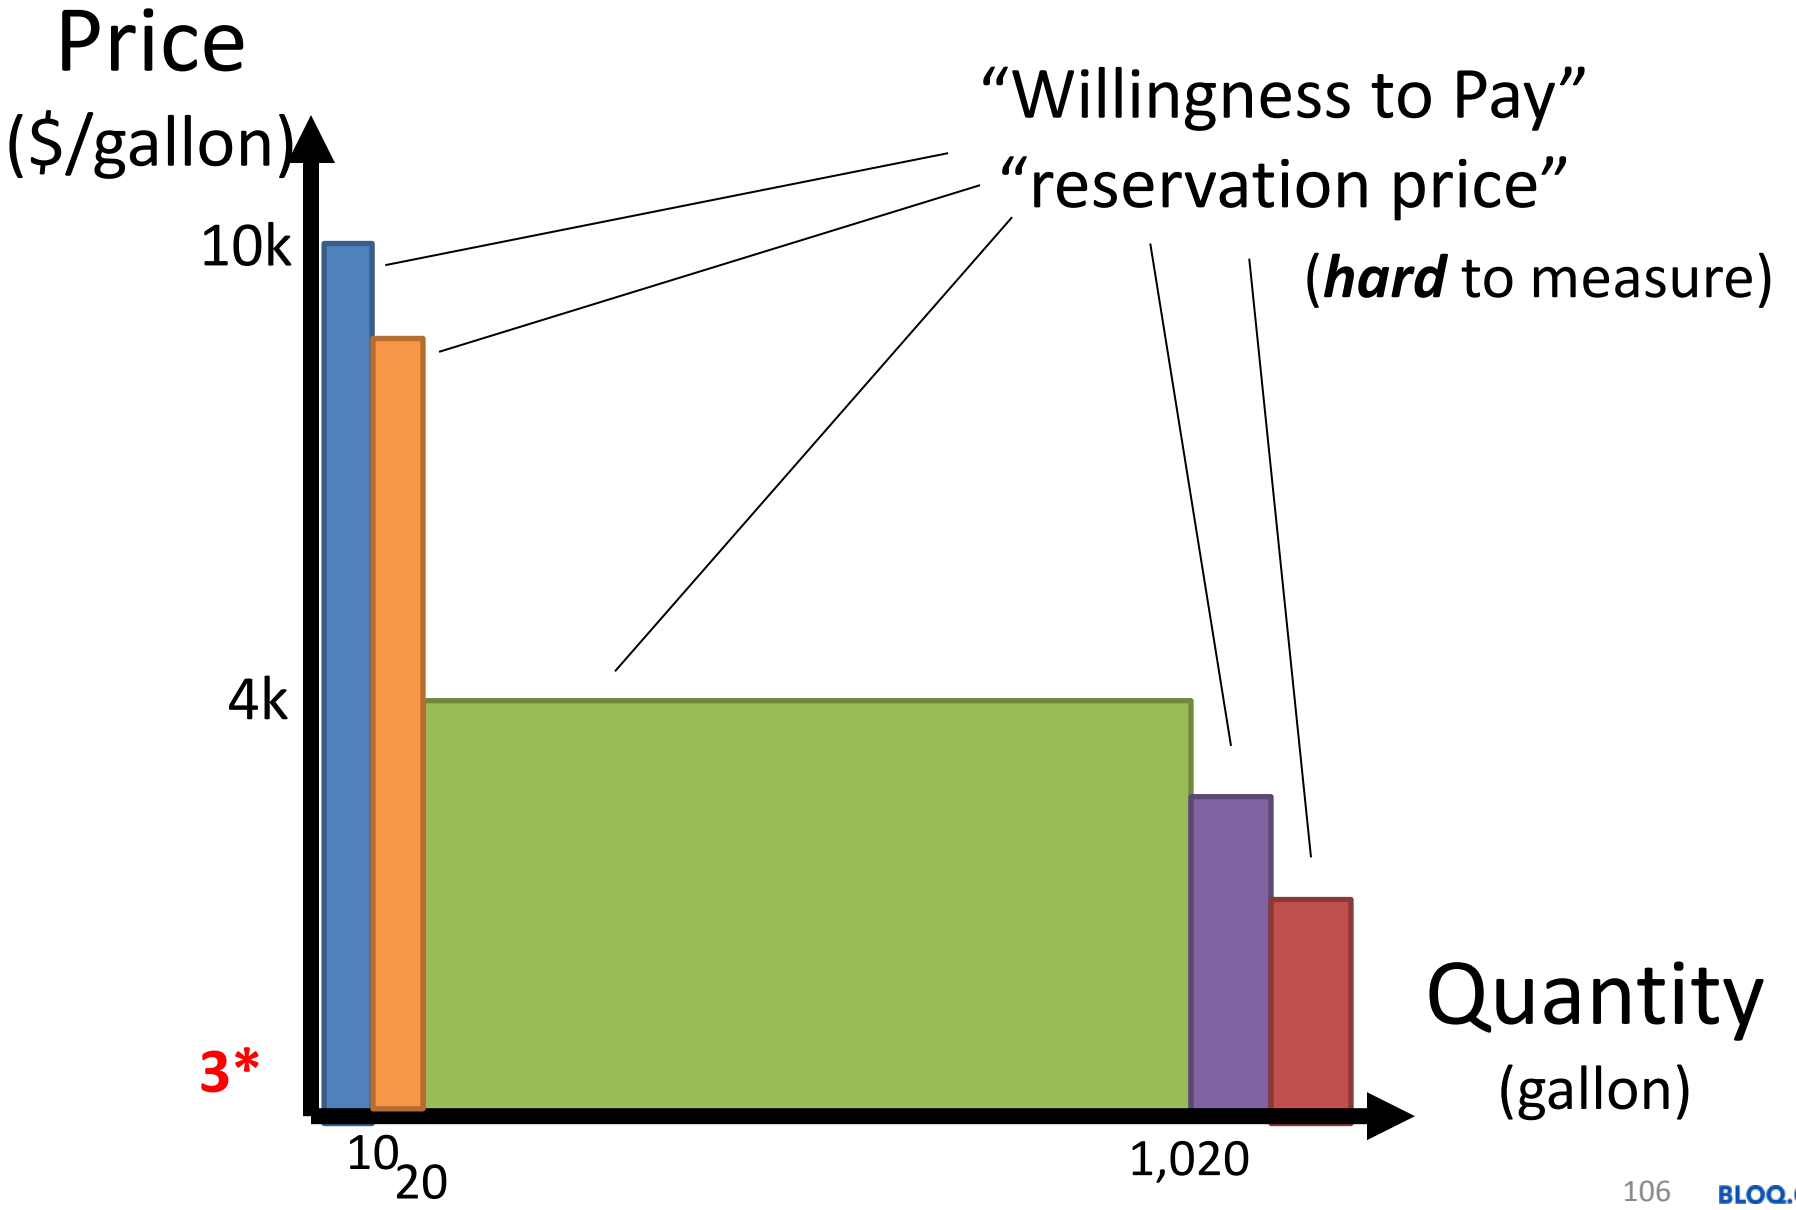

Normal Market – Demand for Oil

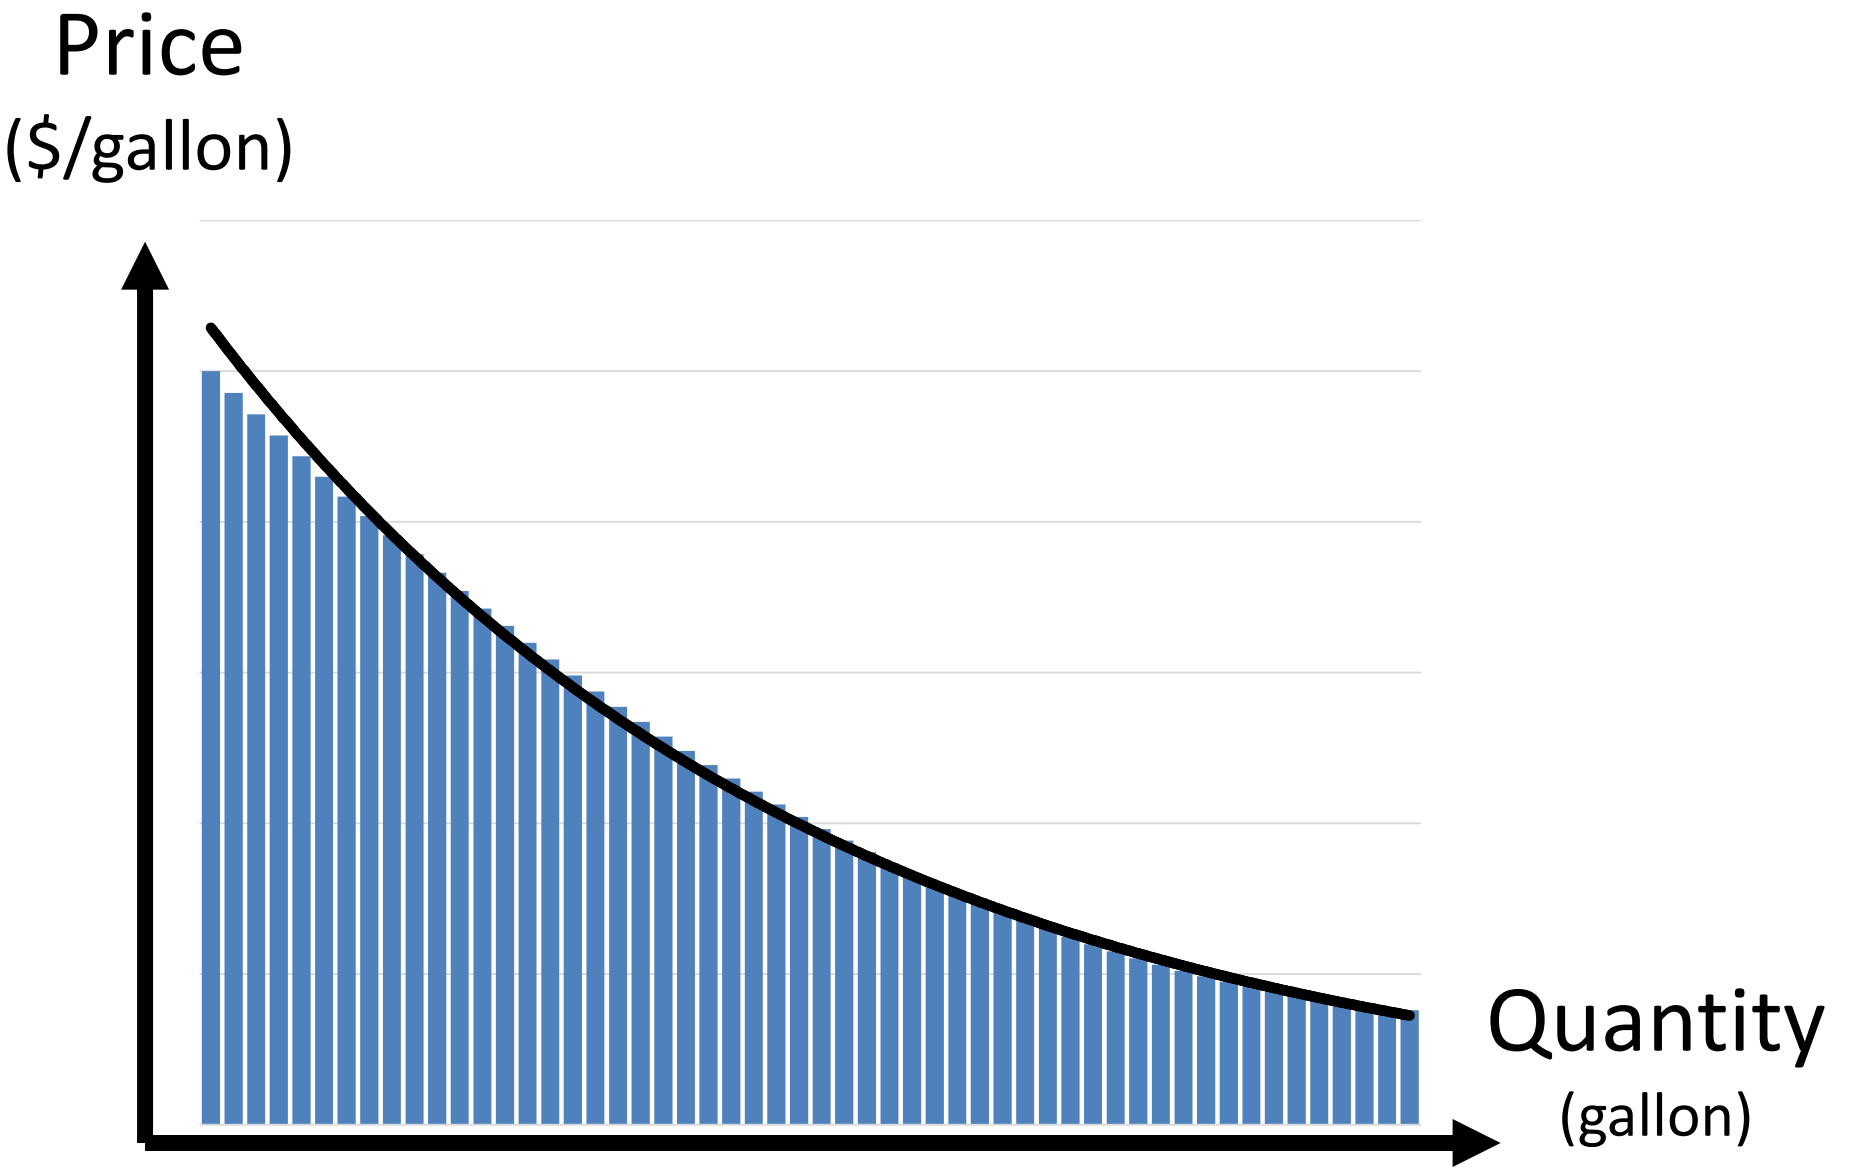

Normal Market – **Supply** of Oil

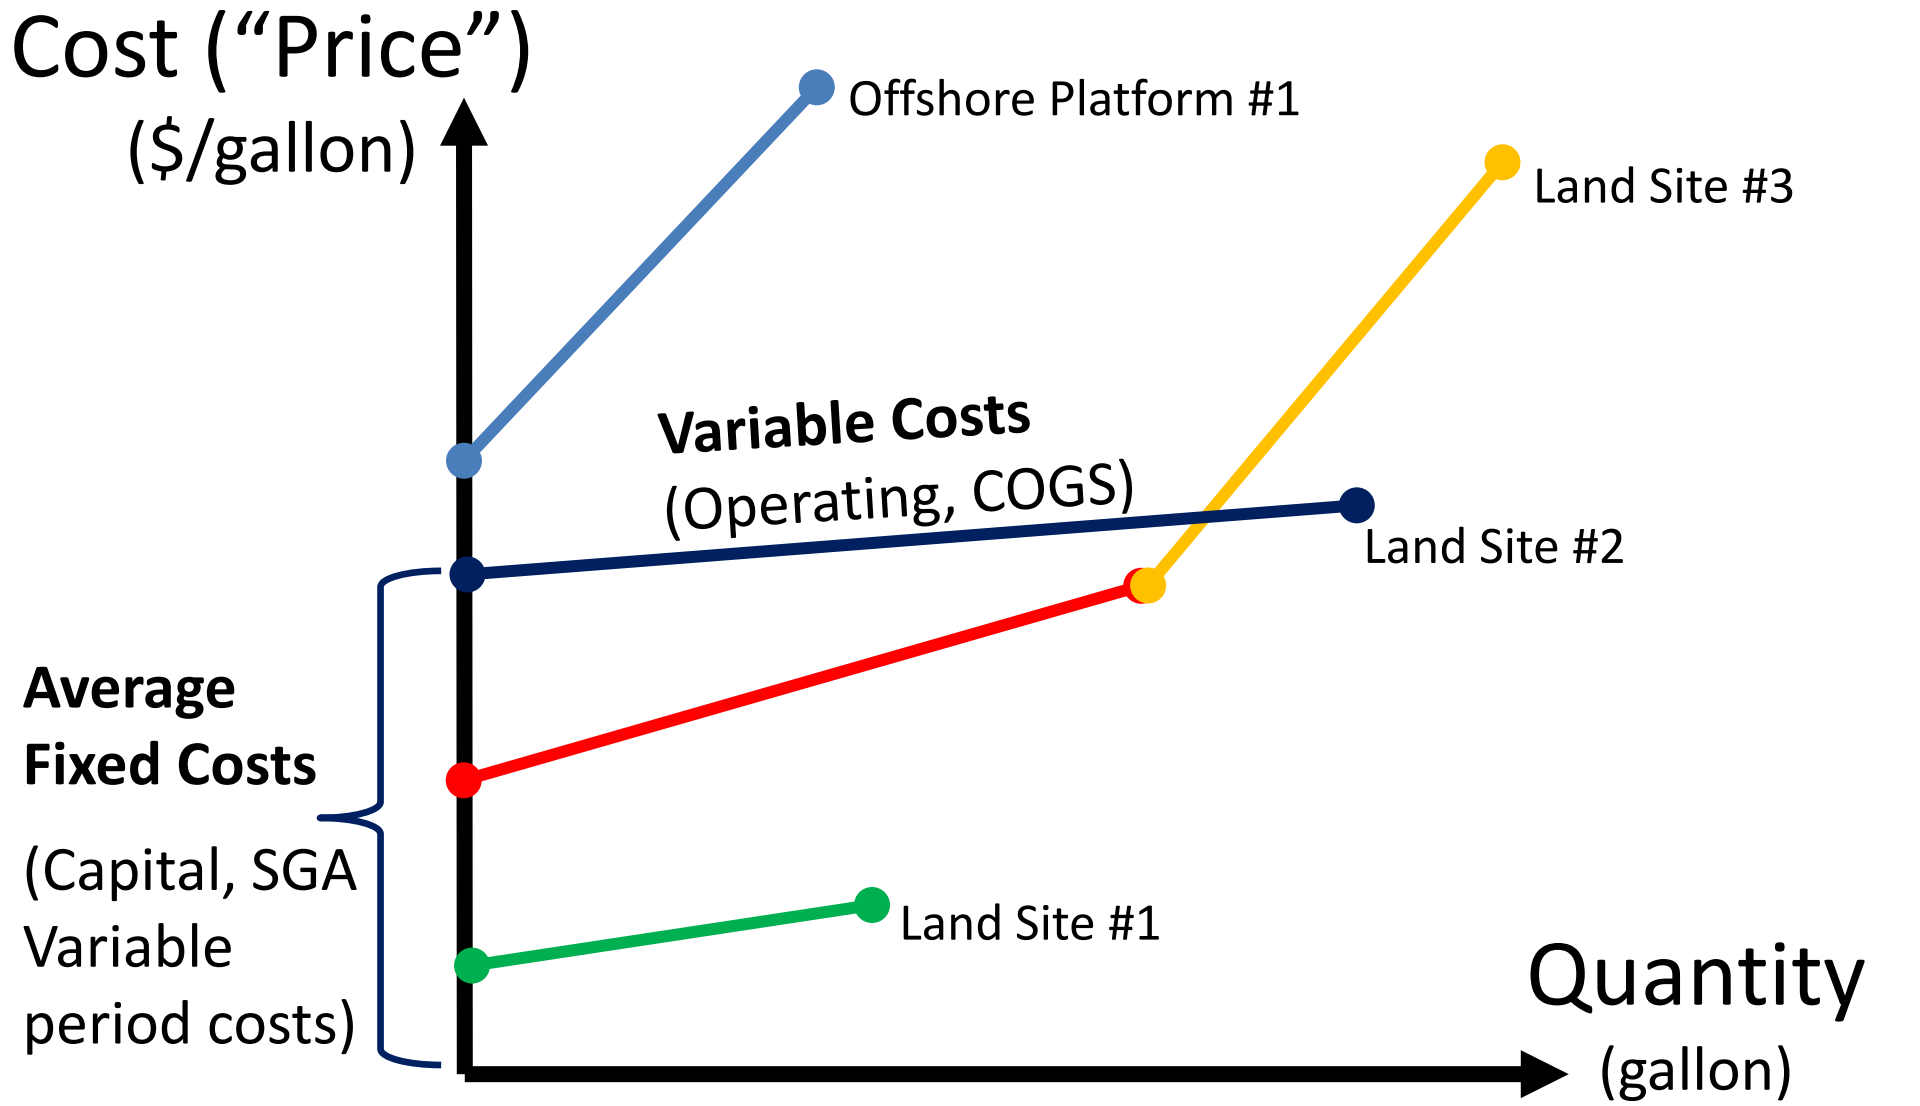

Normal Market – **Supply** of Oil

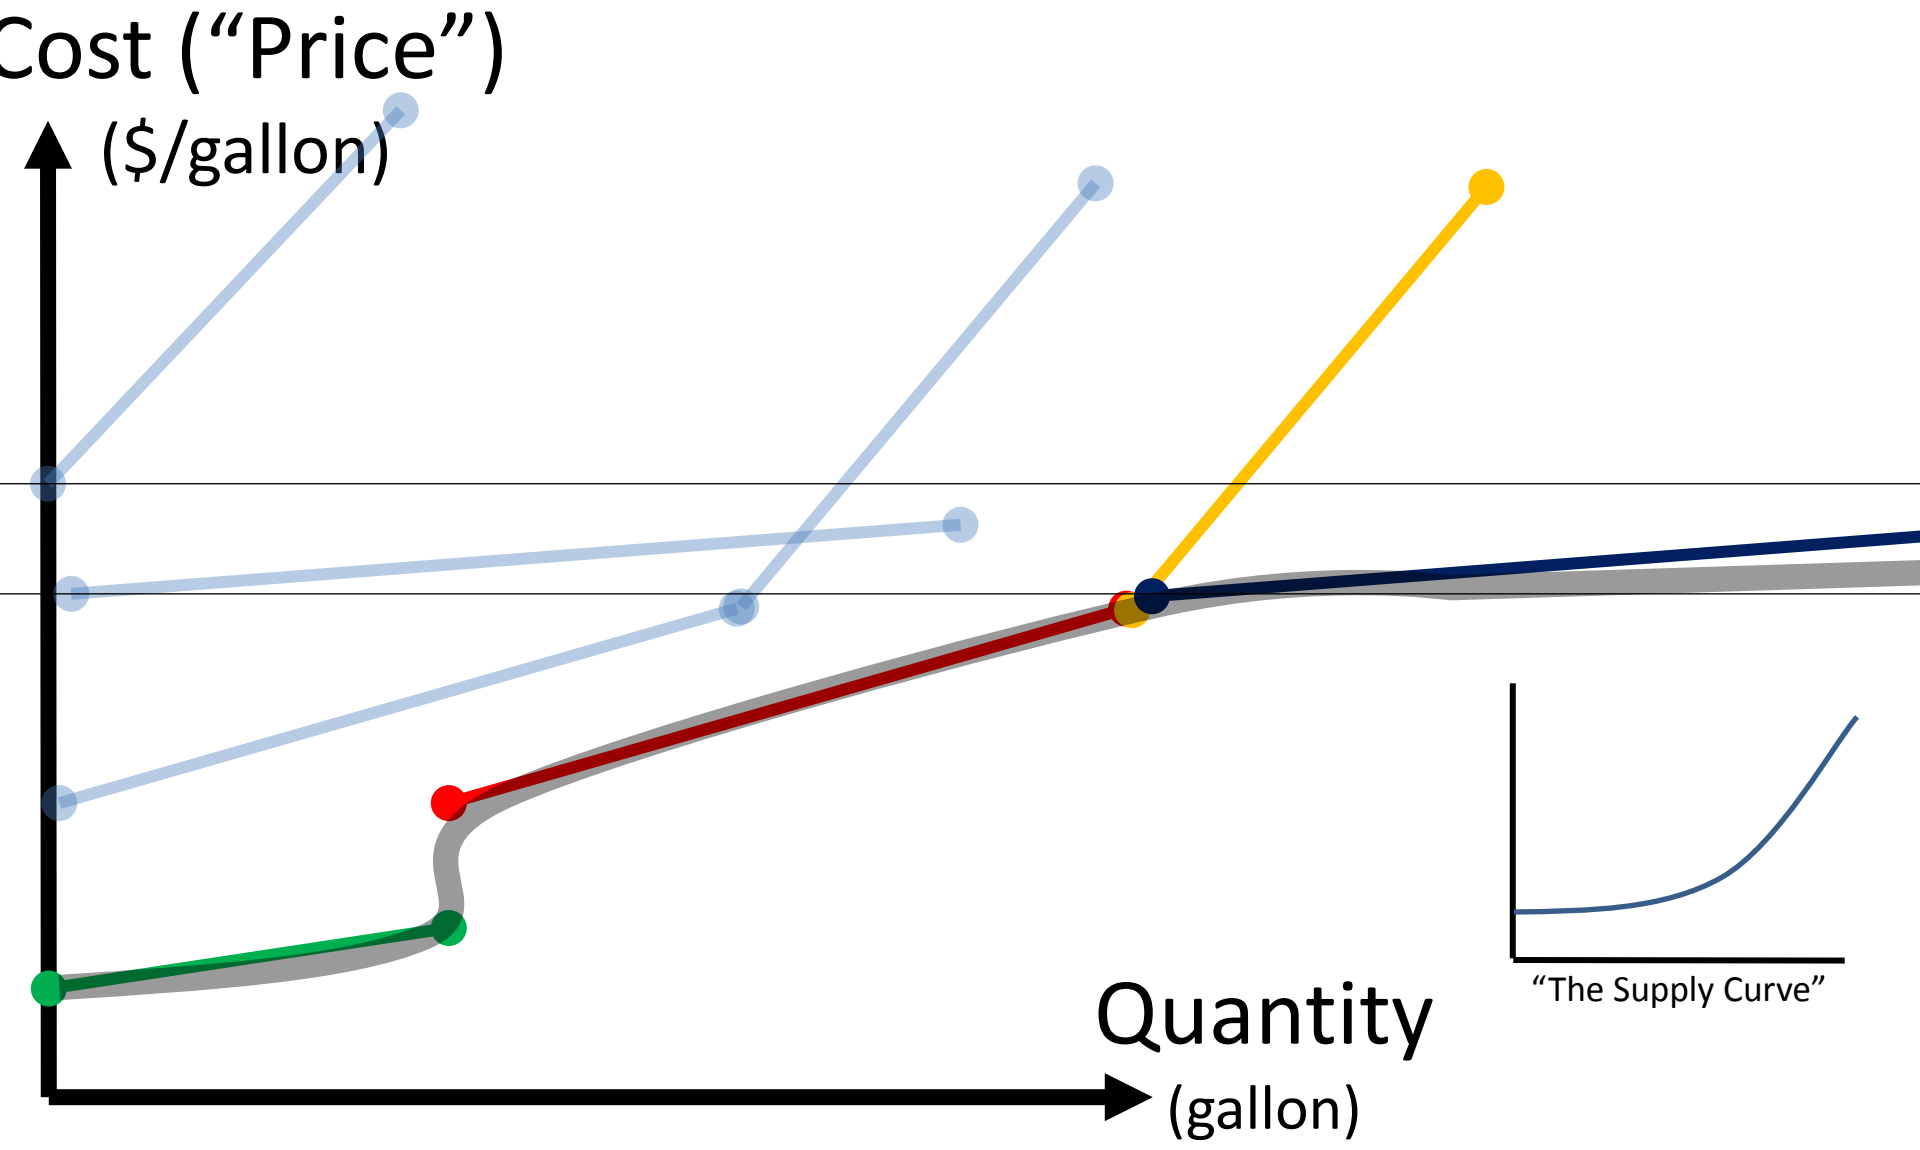

Features of Market Exchange

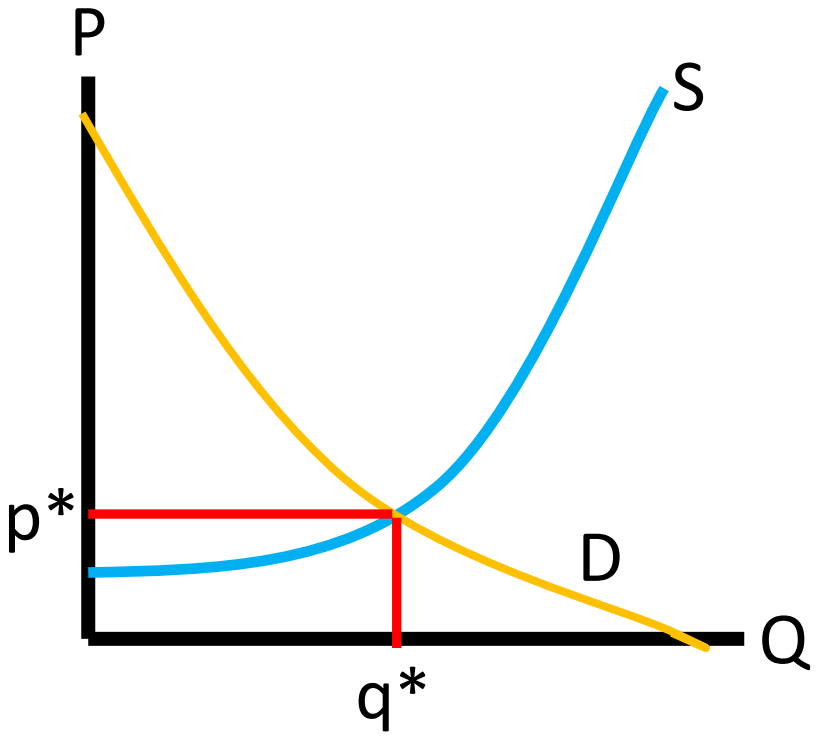

Competition

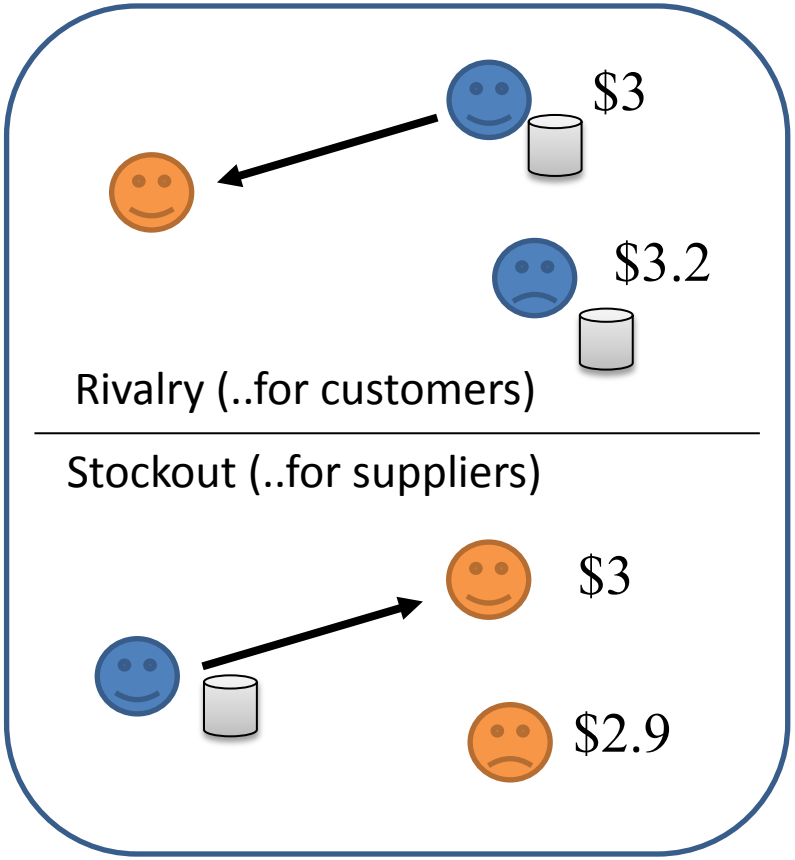

[ii] Block-Access Abnormalities

1. No choice. Buyers don't choose their miners.
2. No concept of "individual".
    1. Network is Public
    2. Pseudonymous
    3. Agent = User != Account

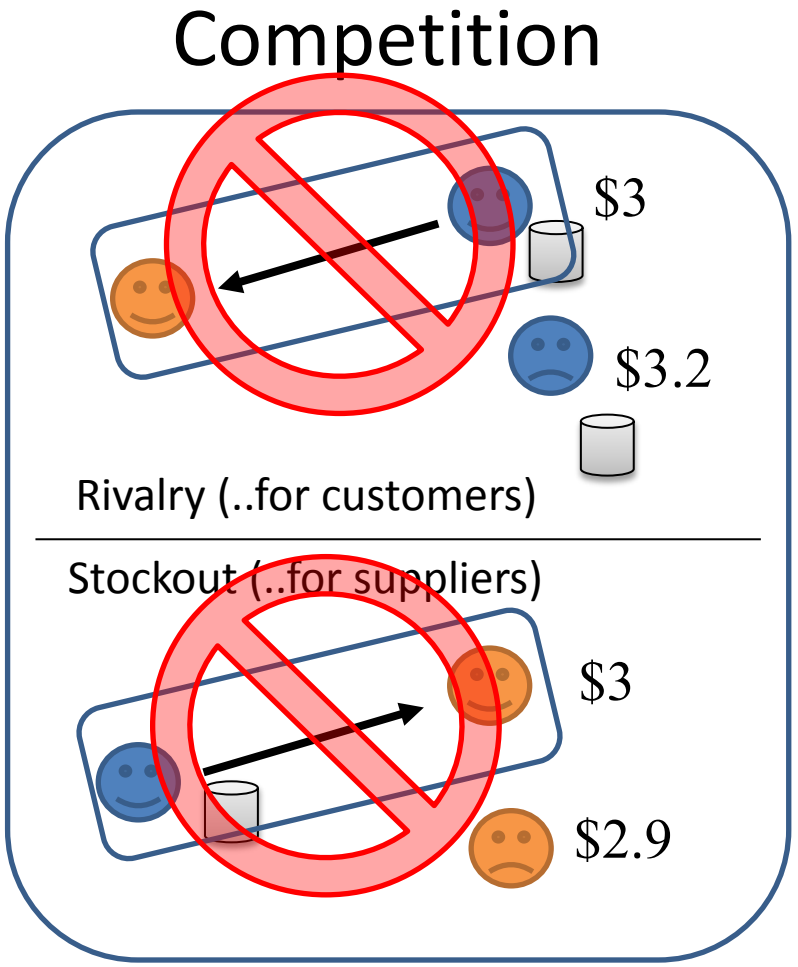

[ii] Abnormalities, cont.

3. Sellers (miners) don't control *block* production.
   - Can't "choose" to make more blocks.
   - Each block can, theoretically, hold ~+INF txns.

A block header contains these fields:

Field	Purpose	Updated when...	Size (Bytes)
Version	Block version number	You upgrade the software and it specifies a new version	4
hashPrevBlock	256-bit hash of the previous block header	A new block comes in	32
hashMerkleRoot	256-bit hash based on all of the transactions in the block	A transaction is accepted	32
Time	Current timestamp as seconds since 1970-01-01T00:00 UTC	Every few seconds	4
Bits	Current target in compact format	The difficulty is adjusted	4
Nonce	32-bit number (starts at 0)	A hash is tried (increments)	4

4. Sellers have nothing to sell!
   - Pre-Block, can't guarantee that they *will* find one.
   - Post-Block, include/exclude policy can't be changed!

[iii] Implications

1. No choice. Buyers don't choose their miners.
2. No concept of "individual".
   1. Public
   2. Pseudonymous
   3. Agent = User != Account

- Miners don't compete on price.
- In fact, there's a "shared customer pool" and "shared production schedule".

[iii] Implications

3. Sellers (miners) don't control *block* production.
   - Can't "choose" to make more blocks.
   - Each block can, theoretically, hold +INF txns.

- "Blocks" are not the relevant "thing supplied".

4. Sellers have nothing to sell!
   - Pre-Block, can't guarantee that they *will* find one.
   - Post-Block, include/exclude policy can't be changed!

- ..?

[iii] Implications

3. Sellers (miners) don't control *block* production.
   - Can't "choose" to make more blocks.
   - Each block can, theoretically, hold +INF txns.

- "Blocks" are not the relevant "thing supplied".

4. Sellers have nothing to sell to Bitcoin users.
   - Pre-Block, can't guarantee that they *will* find one.
   - Post-Block, include/exclude policy can't be changed!

- Miners aren't *transacting* with BTC-users.

[iii] The "Market" for Block Access

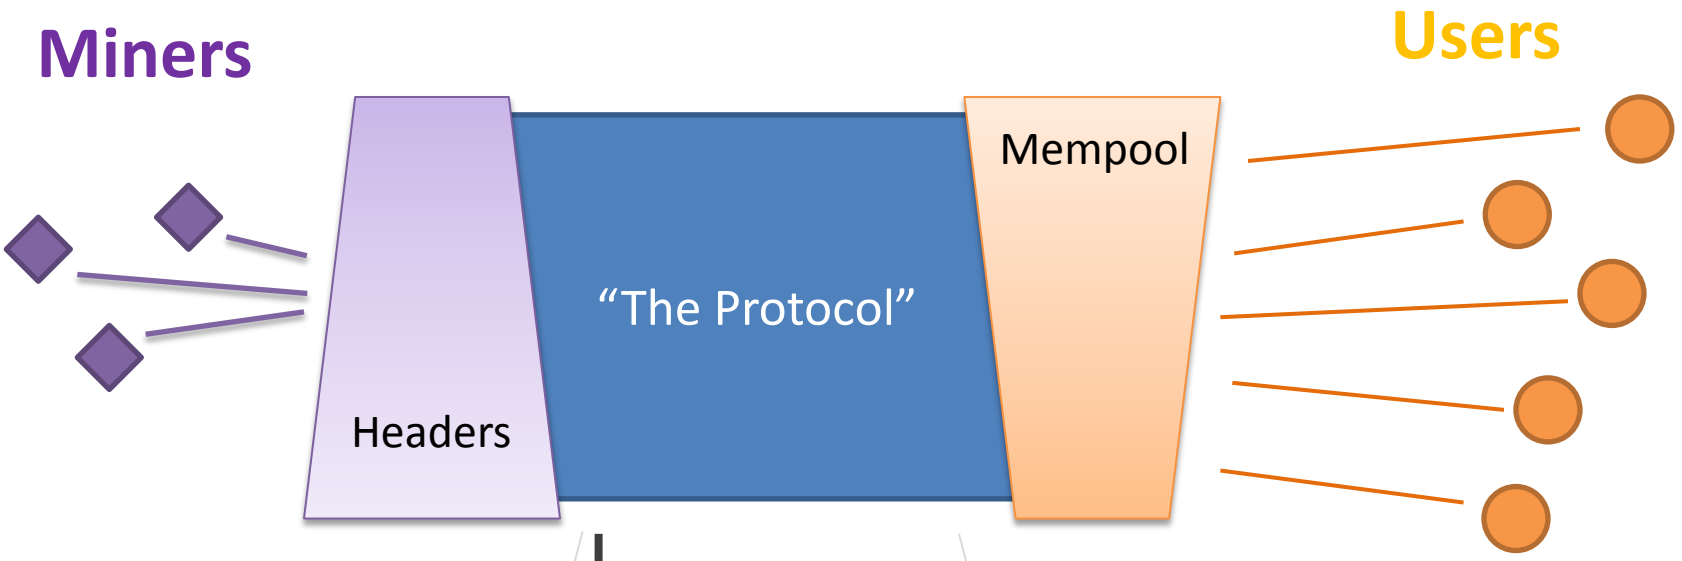

- **Contractors**, "work for BTC"
- {H + AP → BTC}
- H = "hashing activity"
- Brutal competition.

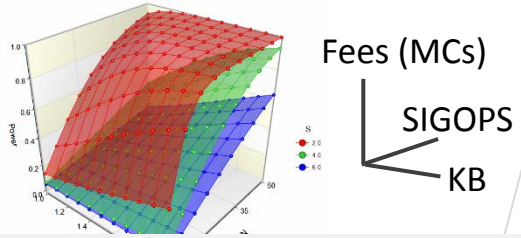

"Access policy" -- maps PTs to In/Out.

Every ~10 mins…

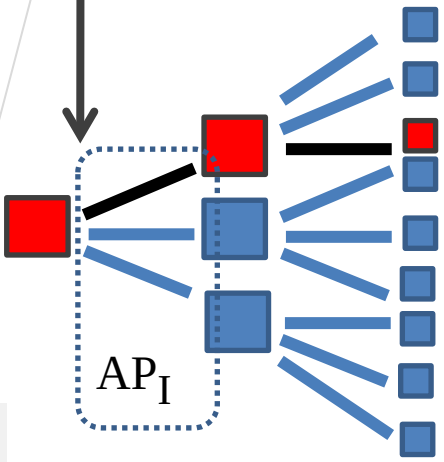

- 24/7 **coupon redemption**
- {BTC → Fees}, on-demand
- PTs ("Prospective Transactions"), always 0-conf, differ as $f(AP_i)$.

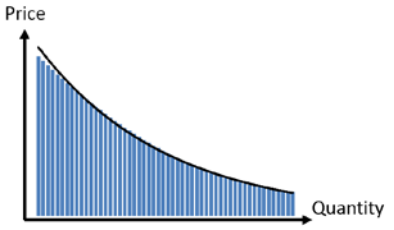

[iv] The Block Tree

- Does NOT include the txns themselves, only the Access Policies.
- Therefore, this tree is fully defined.
- As a concept, it isn't necessary to write out the whole tree (it would be huge).

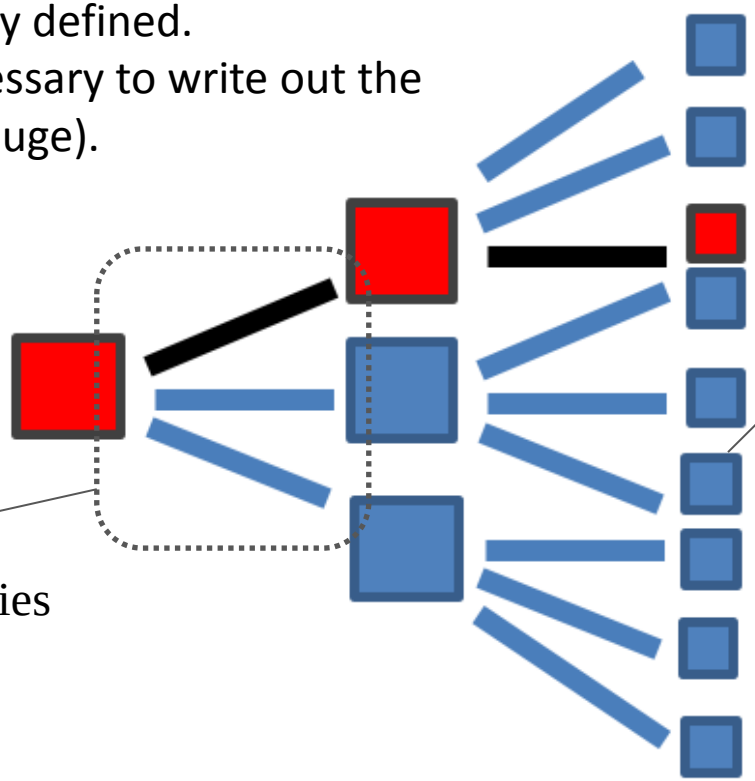

The Point

Not a “market” process.

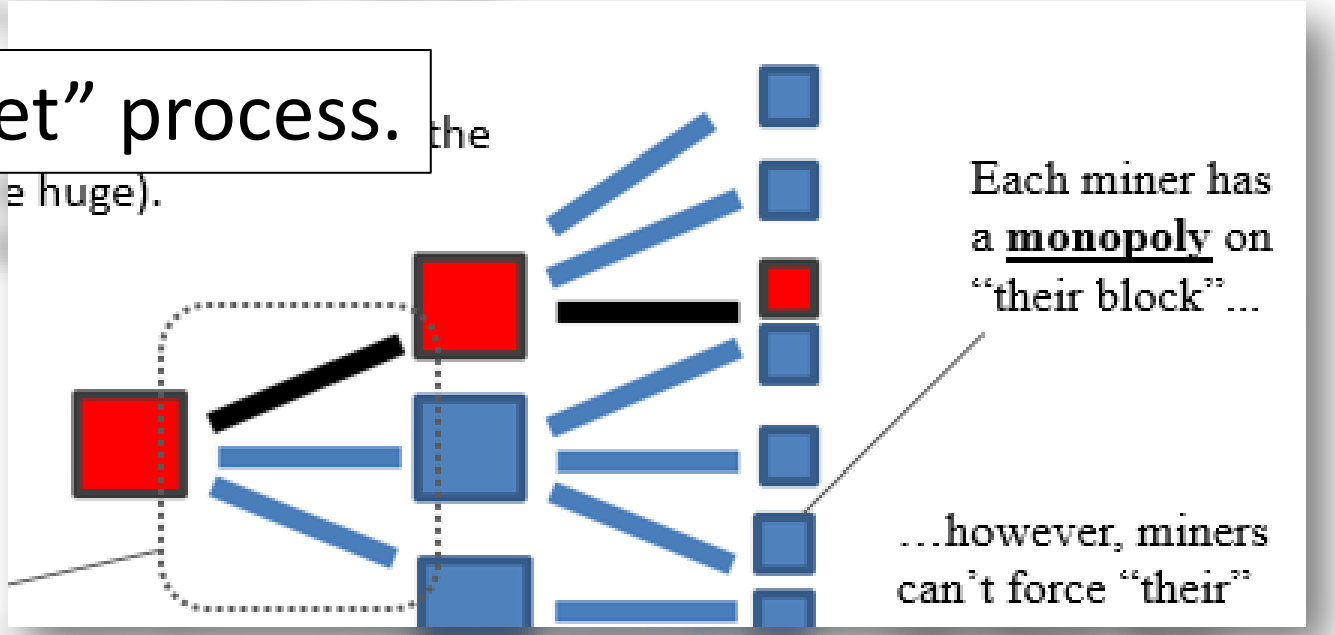

PC vs. M -- Important Dichotomy

Hashrate: Perfect competition.
Access-Policies: Serial monopoly.

?

[iii] The "Market" for Block Access

Miners

Users

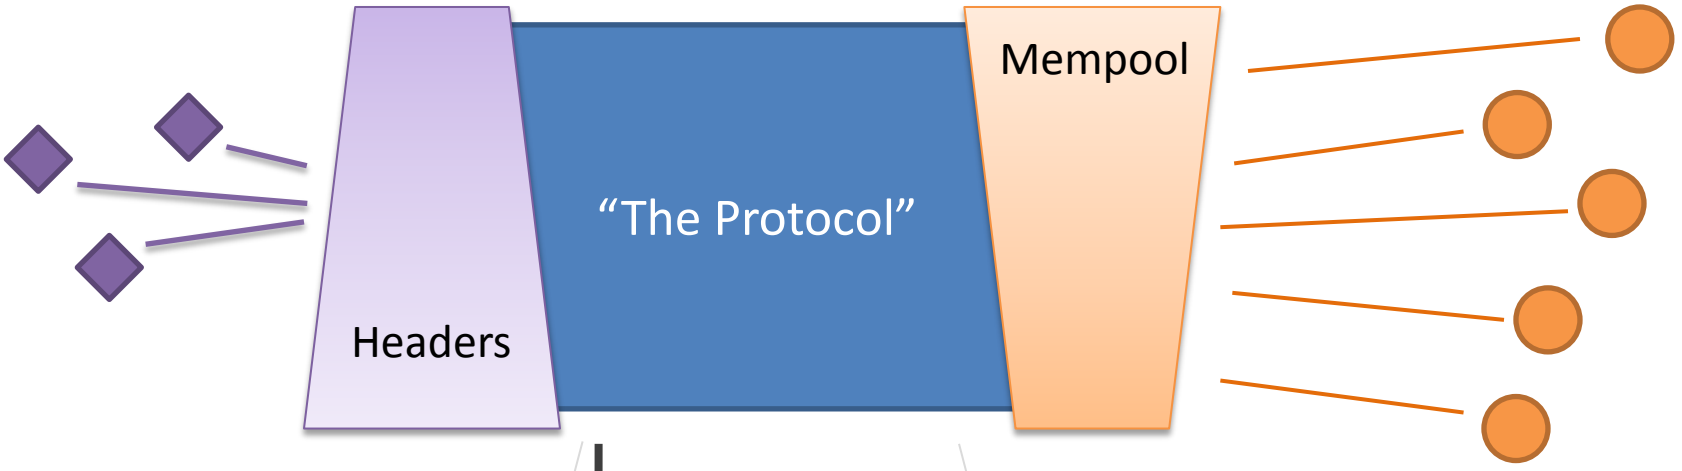

- **Contractors**, "work for BTC"
- {H + AP → BTC}
-
-

Protocol offers a fixed* contract, perfect competition.

Every ~10 mins…

Access Policy: Serial Monopoly

AP_I

- 24
- {B
- PT ions"),
- alv AP_i).

Perfectly reliable BTC conversion service.

BTC → PT Fees

(Perfect competition)

"Access policy" -- maps PTs to In/Out.

Reminder

How interdependent are miners?

2. Fees
   (#4 x #4)

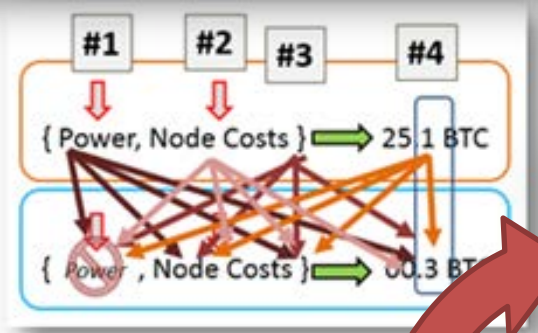

- Sub-Agenda
  - i. A "Normal" Market
  - ii. Abnormalities in Our Case
  - iii. Implications of these Abnor
  - iv. The Block Tree

1. Three Different Perspe ves
2. Bitcoin's Transition (to Equilib
3. The "Market" for Block-Access
4. Coase vs Folk -- Miner Coordination
5. Demand Curve Calculus

4. Coase vs Folk

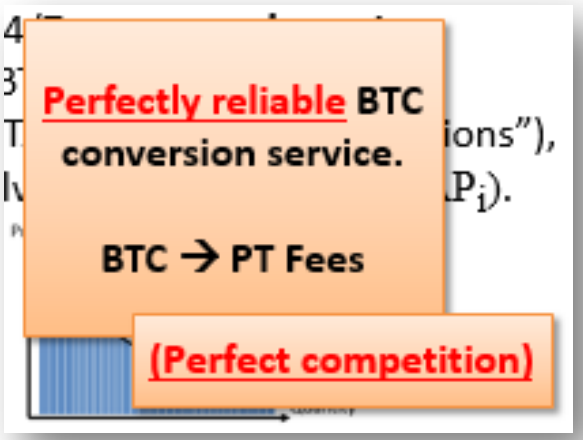

Coase Conjecture

Monopolies lose, if customers are patient.

Ronald Coase is a remarkable modern economist in the sense that he is independent thinking, rigorous, creative, with ideas that are applicable and explain the world around us –in other words, the real thing. His style is so rigorous that he is known for the Coase Theorem, an idea that he posited

11/29/15 N. N. Taleb.

Coase Conjecture [edit]

Another important contribution of Coase is the Coase Conjecture: an informal argument that durable-goods monopolists do not have market power because they are unable to commit to not lowering their prices in future periods.

because the monopoly is, in effect, in price competition with itself over several periods and the consumer with the highest valuation, if he is patient enough, can simply wait for the lowest price. Thus the monopolist will have to offer a competitive price in the first period which will be low.

Perfect Competition Wins

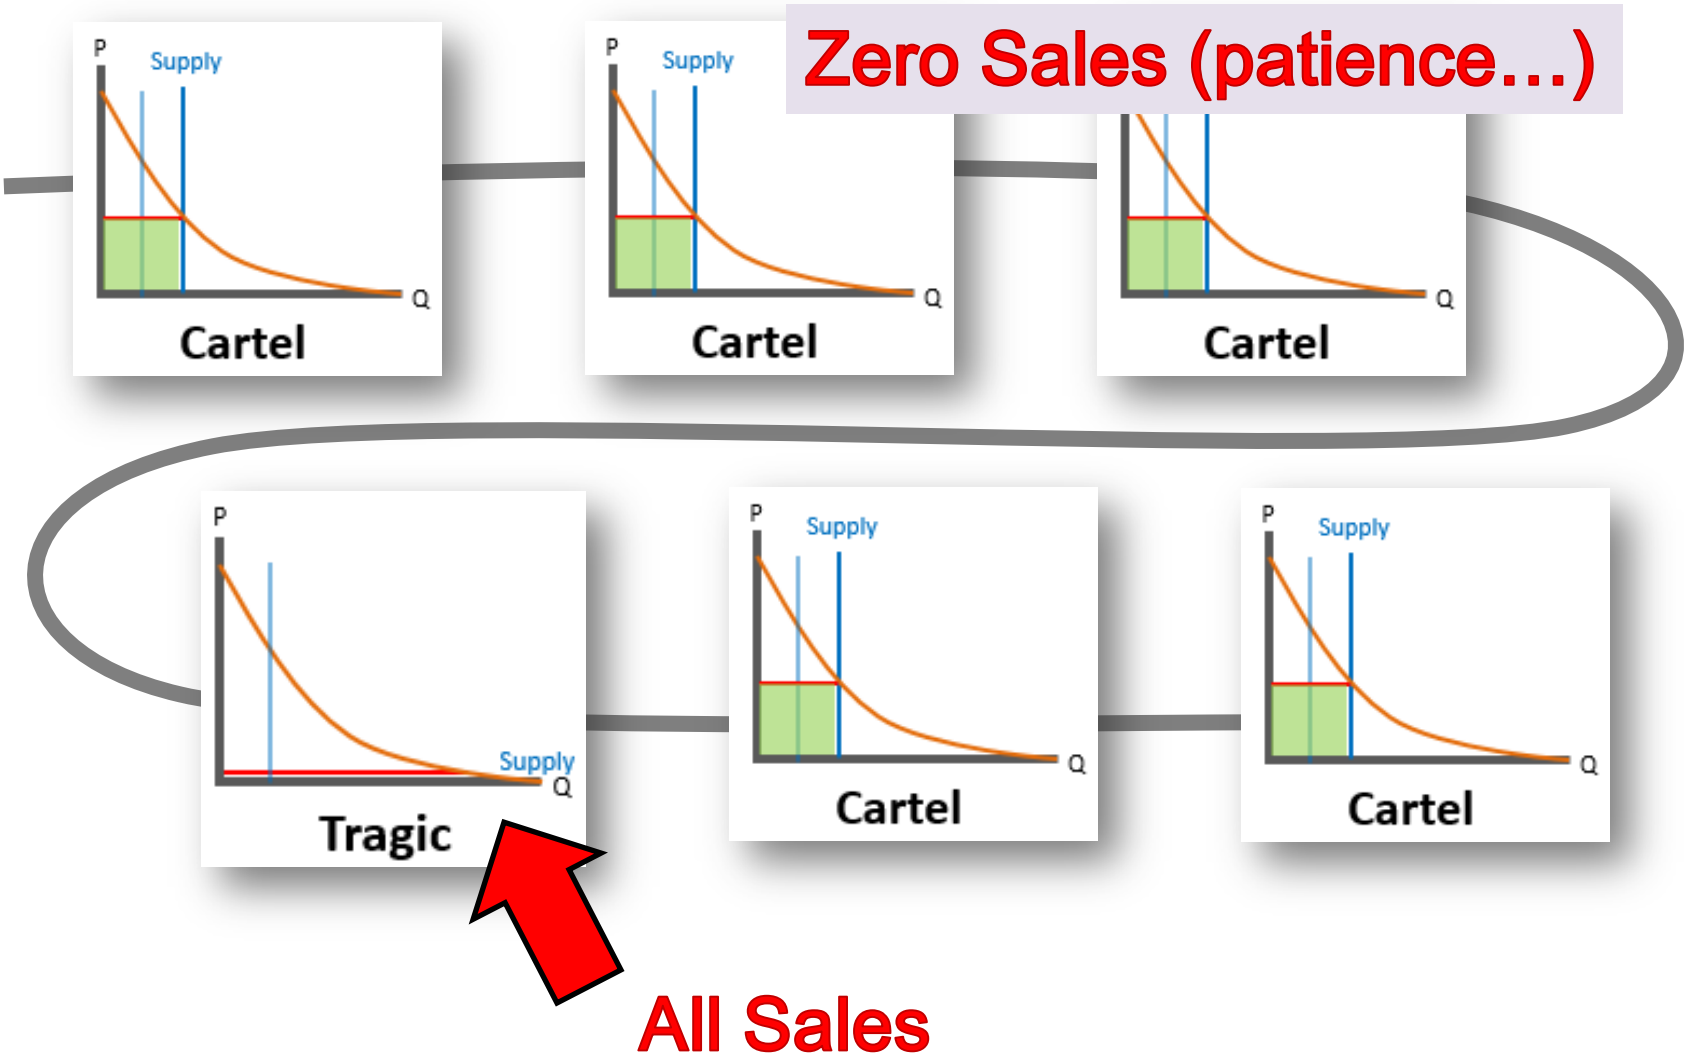

Folk

Coase Conjecture [edit]

Another important contribution of Coase is the Coase Conjecture: an informal argument that durable-goods monopolists do not have market power because they are unable to commit to not lowering their prices in future periods.

Are they really *unable* to do this?

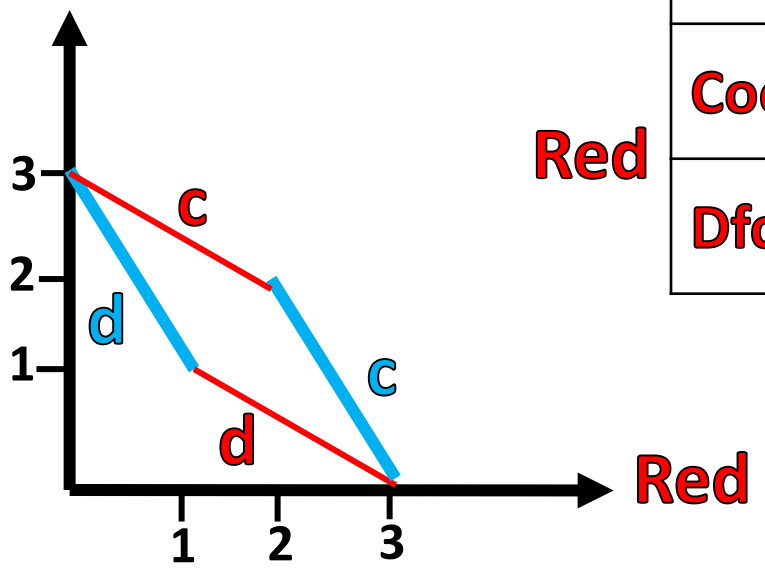

		Blue: Coop	Blue: Dfct
Red	Coop	2, 2	0, 3
Red	Dfct	3, 0	1, 1

-- Repeated Game.

Future Punishment → Today's Cooperation

Folk Basics

Future Punishment → Today's Cooperation

If…
[1] Game is long.
[2] Players are patient.

Then..
[3] Players can be punished (in the future).
Result…
[4] Mutually, all can hit any region better than "punish eq".

Simple rule:
[1] Choose Eq*
[2] Select Eq* until someone doesn't, Punish defectors.

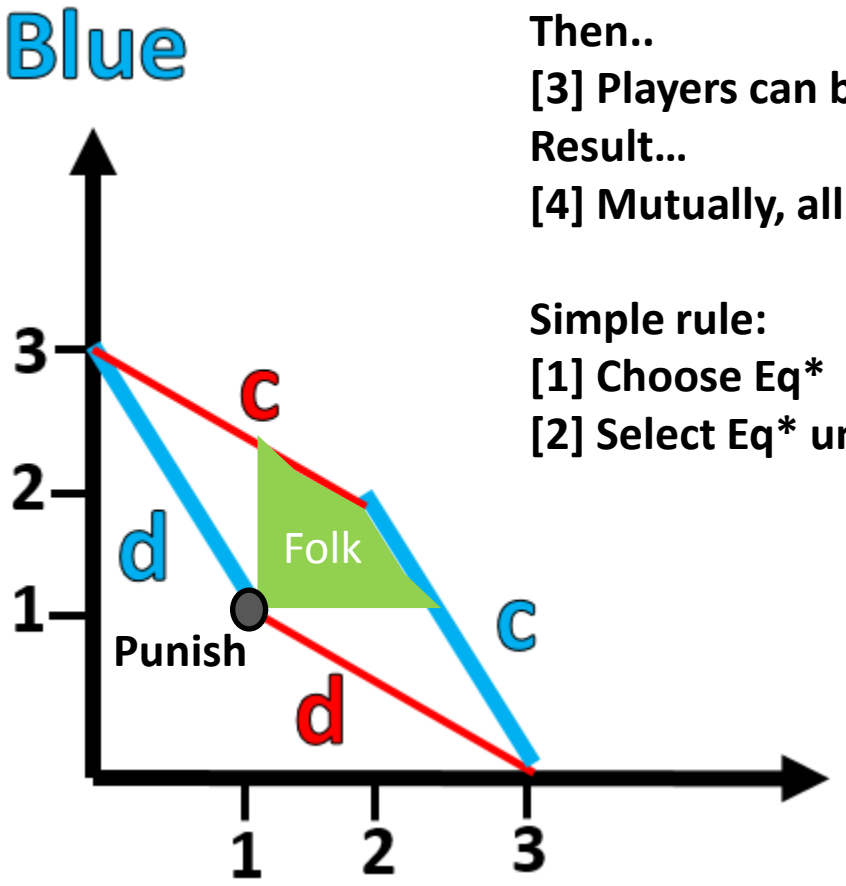

		Blue: Coop	Blue: Dfct
Red	Coop	2, 2	0, 3
Red	Dfct	3, 0	1, 1

Long Run Equilibrium

If…
[1] Game is long.
[2] Playe

Then..
[3] Playe
[4] Mutu **…ish eq".**

Simple:
[1] Choose Eq*
[2] Select Eq* until someone doesn't, Punish defectors.

If punishment is possible, cooperation is enforceable.

If value-transfers are possible, total co-op value is maximized.

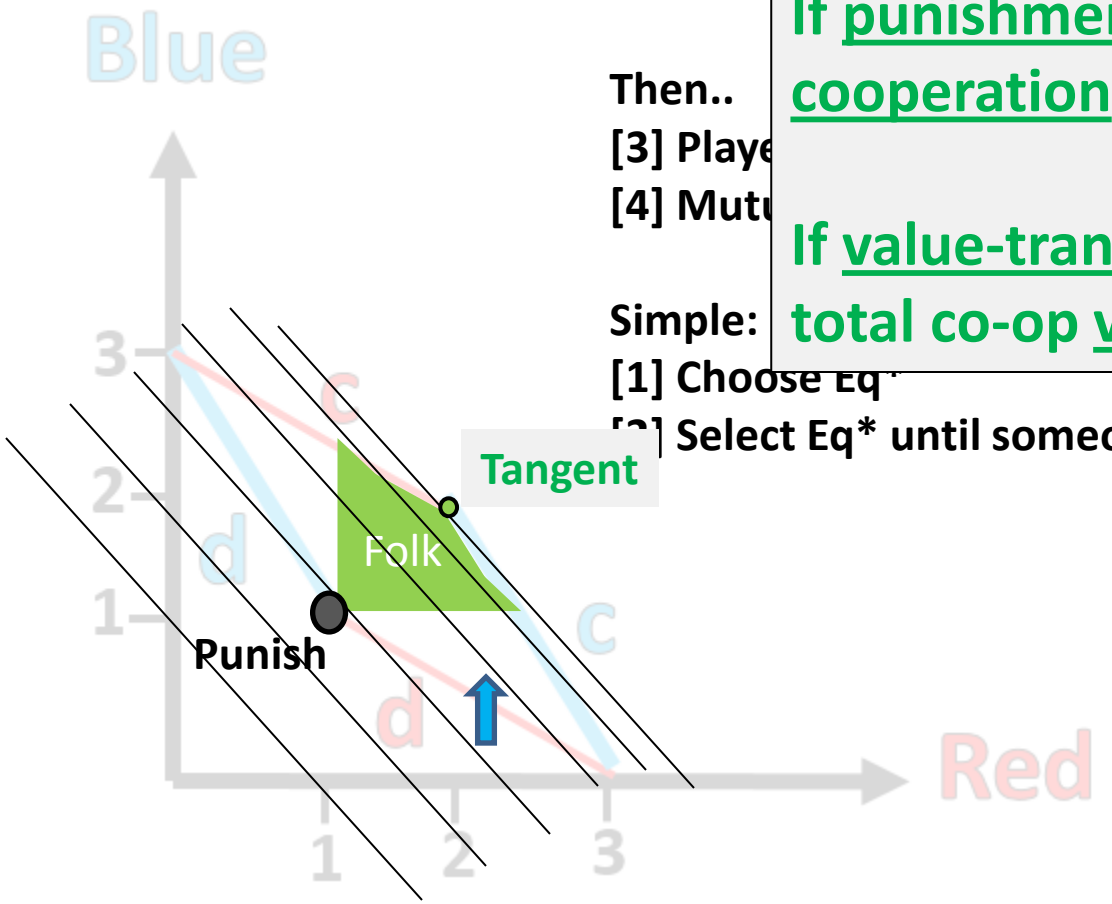

		Blue: Coop	Blue: Dfct
Red	Coop	2, 2	0, 3
Red	Dfct	3, 0	1, 1

Perfect Competition ~~Wins~~ Loses

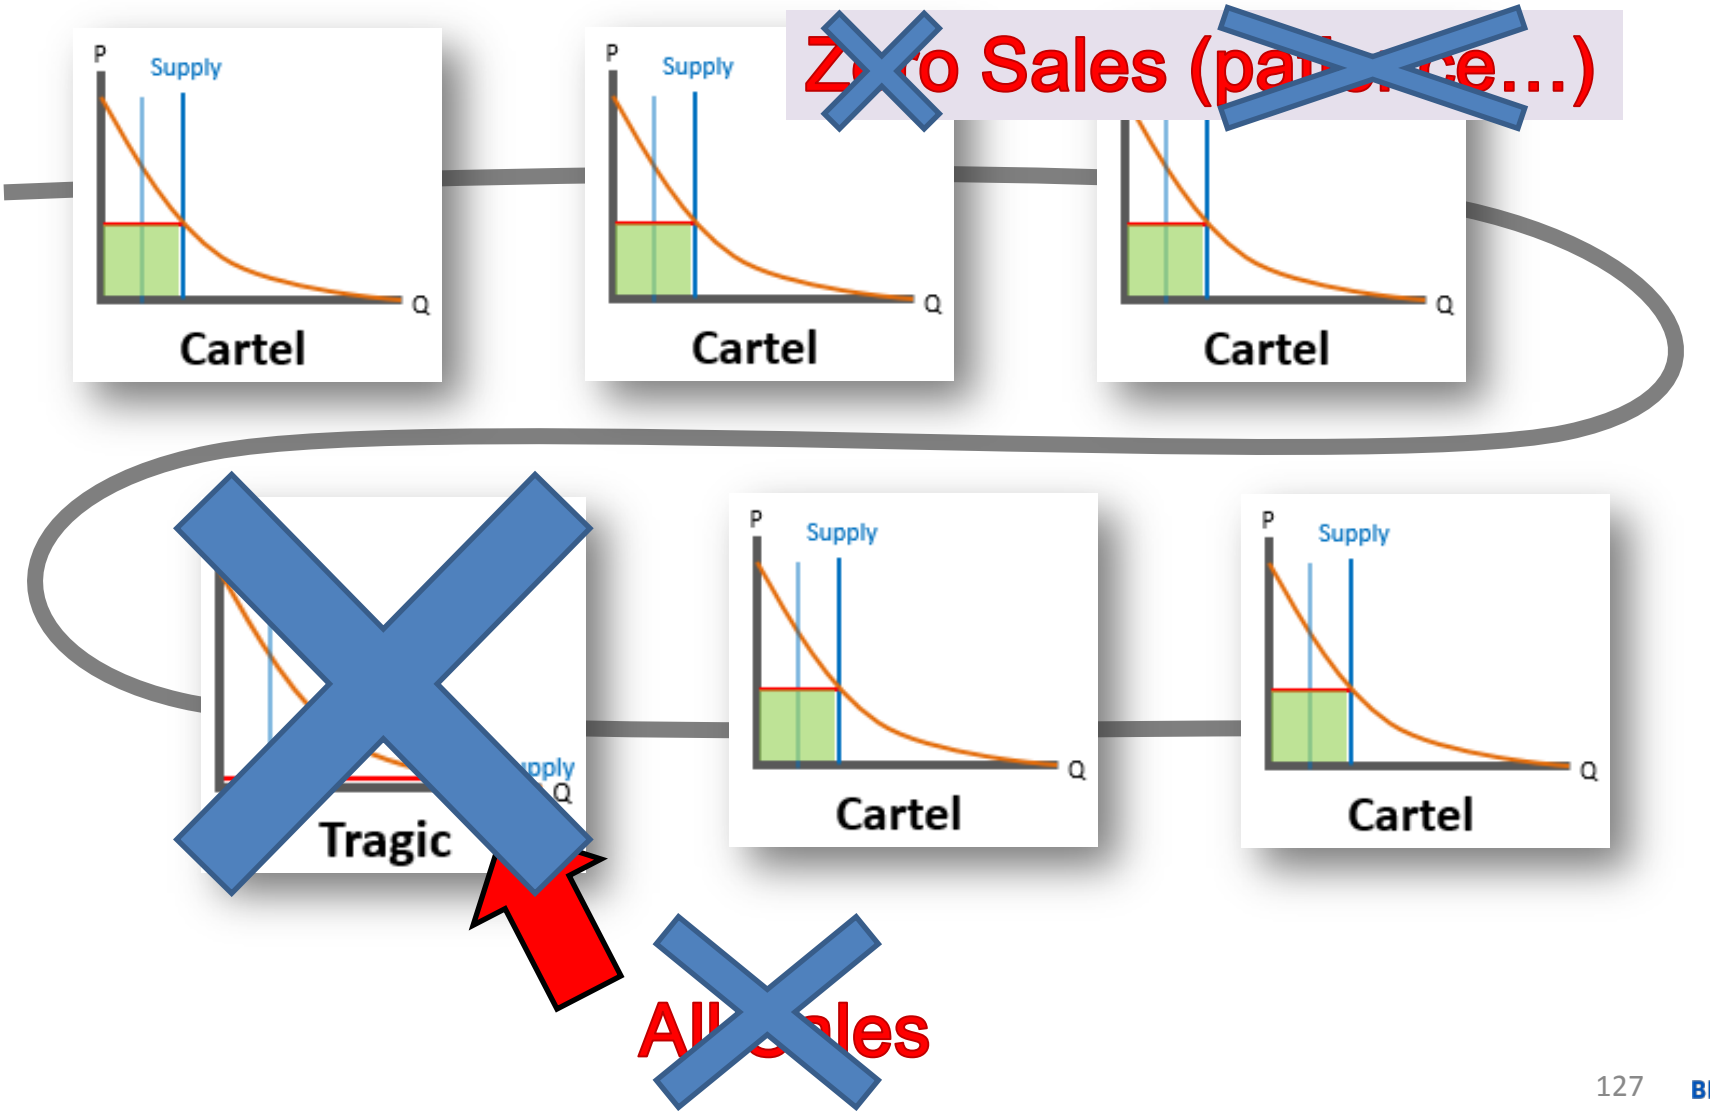

Perfect Competition ~~Wins~~ Loses

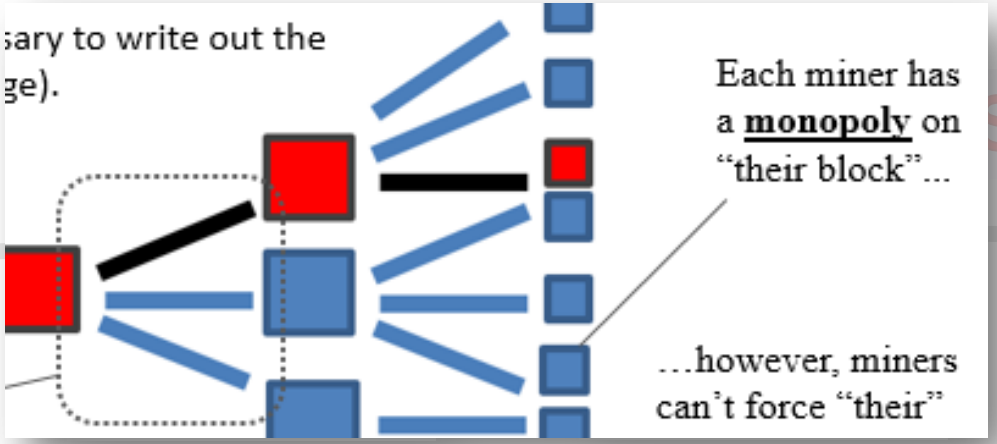

Access Policy: Serial Monopoly

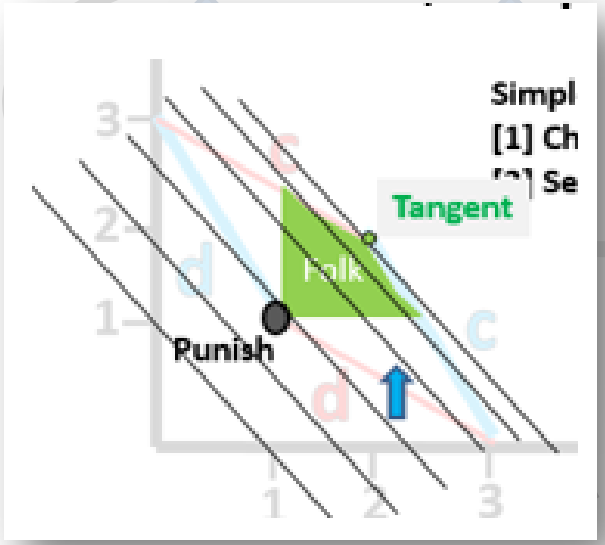

Coase Conjecture [edit]

Another important contribution of Coase is the Coase Conjecture: an informal argument that durable-goods monopolists do not have market power because they are unable to commit to not lowering their prices in future periods.

Need: predefined software rule.

“Revenue” Fork

- **Typically,** soft forks are used *to upgrade* Bitcoin’s software – features, bugfixes, resource improvements.
- **This** is a fork *to manage business policy* of Bitcoin, –specifically, to **optimize tx fees**.

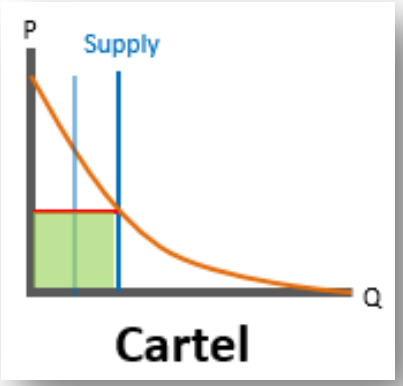

Simple:
[1] Choose Eq*
[2] Select Eα* until someone doesn’t
, Punish defectors.

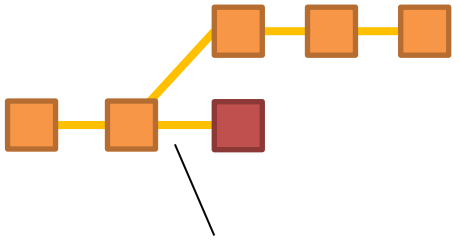

- “Should” happen, vs. “will” happen.

2. Fees

(#4 x #4)

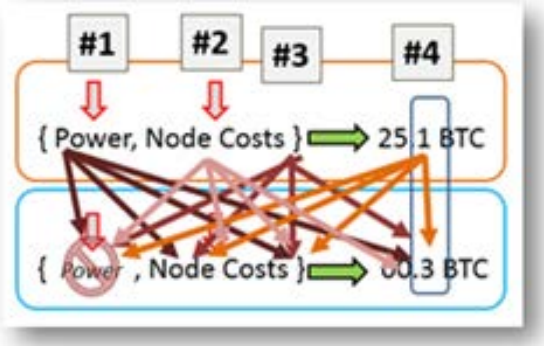

Section Agenda

- Review (Completed)

1. Three Different Perspectives
2. Bitcoin's Transition (to Equilibrium)
3. The "Market" for Block-Access
4. Coase vs Folk -- Miner Coordination
5. Demand Curve Calculus

[5] Demand Curve Calculus

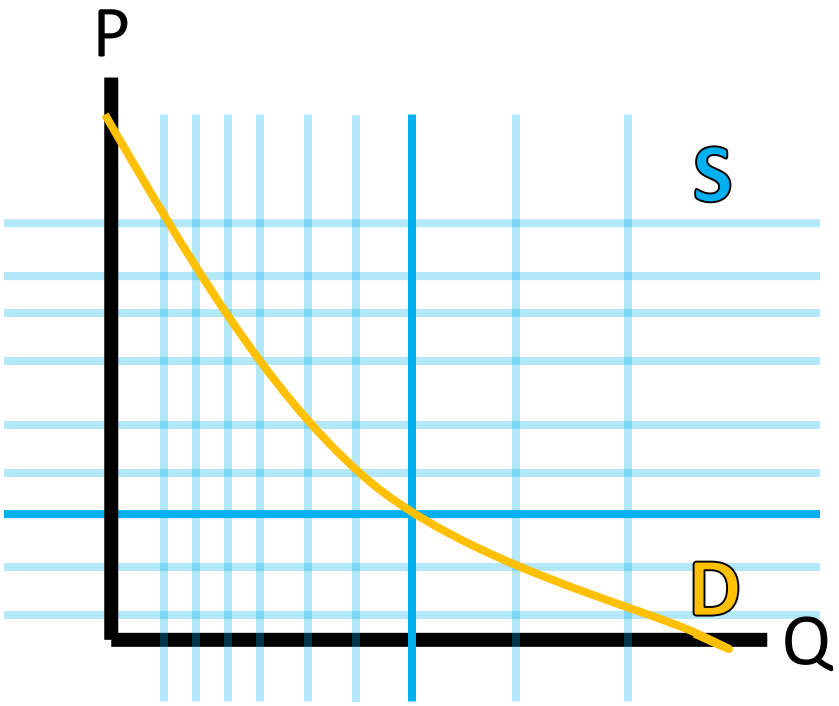

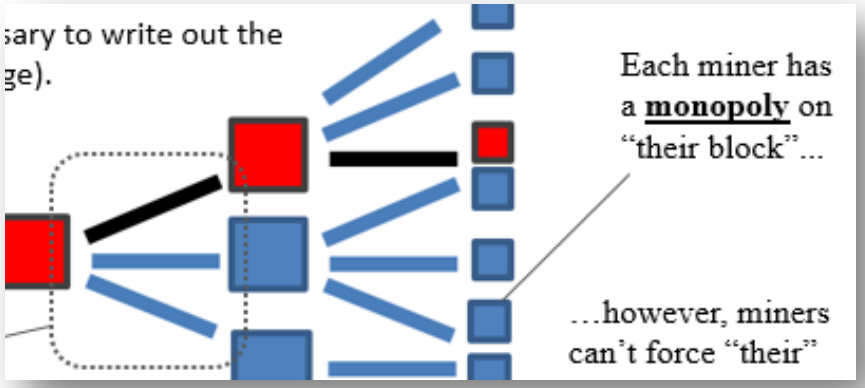

Need: predefined software rule.

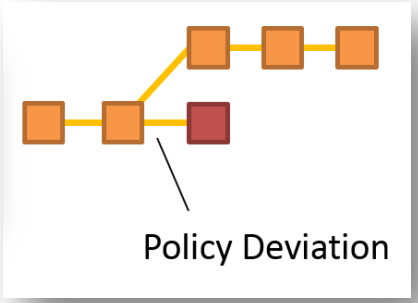

- Miners can choose (p*, q*).
- And they can choose as a group.

Caveat

- This model assumes, for simplicity, that *demand is relatively constant within a ~1 week period*. So, it doesn't apply well to contemporary Bitcoin, where txns enter in real-time.
- (Because, when the sun is over the Pacific, everyone's asleep, demand should be lower).

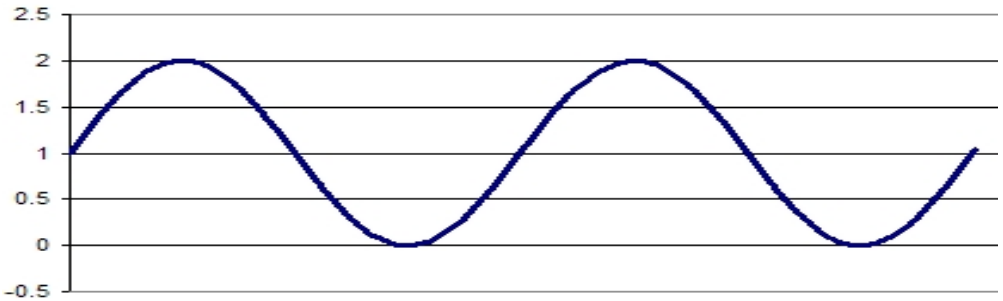

- Instead, this model approximates a future where there is a "constant backlog of txns". Such is the expected behavior under LN.

Where are revenues maximized?

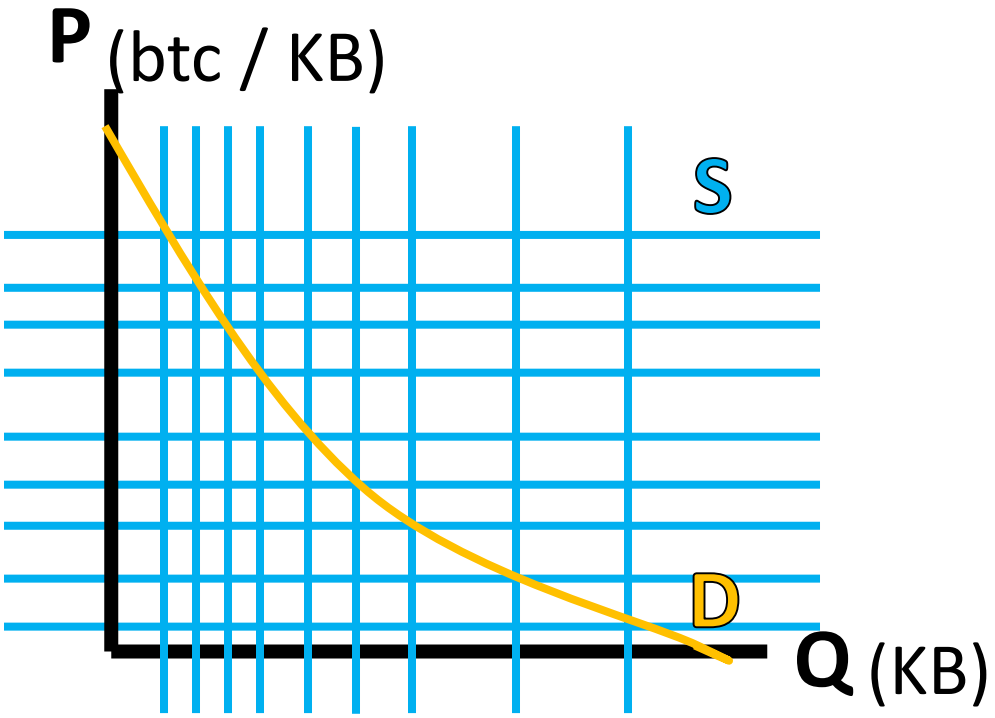

Total Revenues = [btc / KB] * [KB]

…let's [1] explore, then [2] select…

Where are revenues maximized?

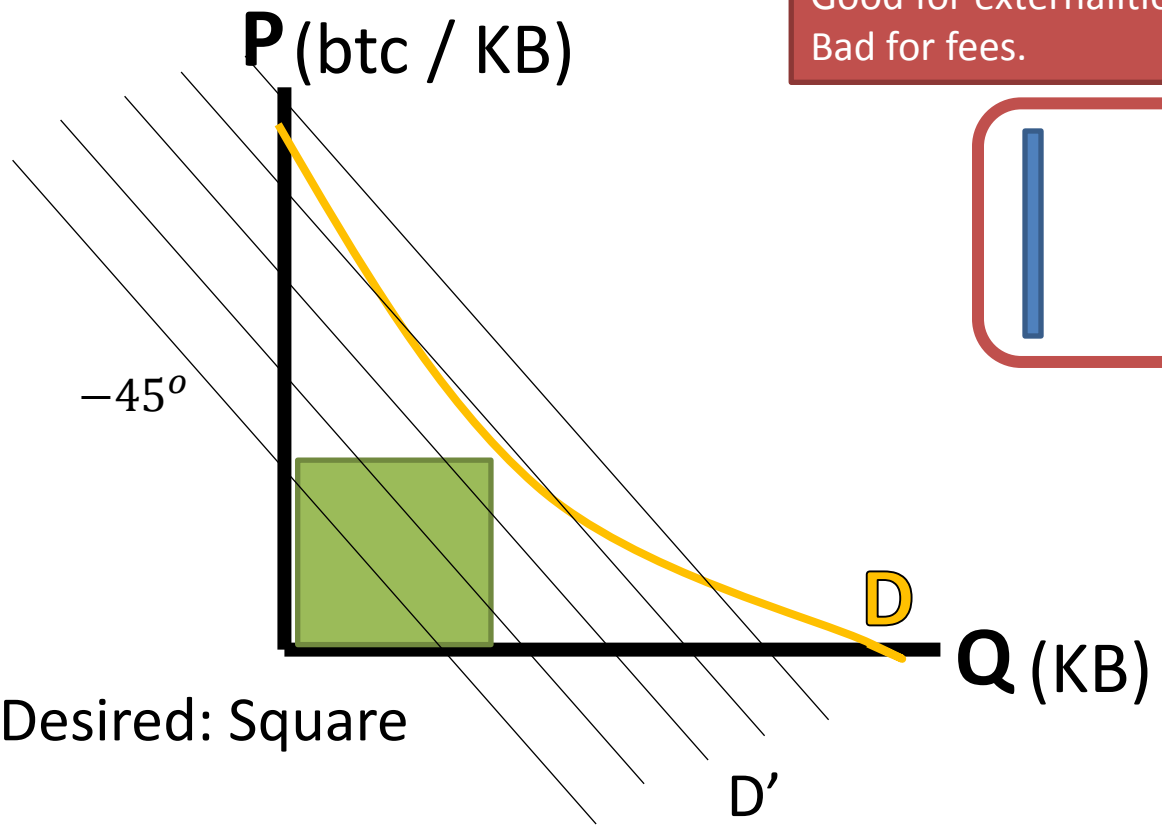

Desired: Square

Good for externalities,
Bad for fees.

Medium for externalities,
Good for fees.

Bad for externalities,
Bad for fees.

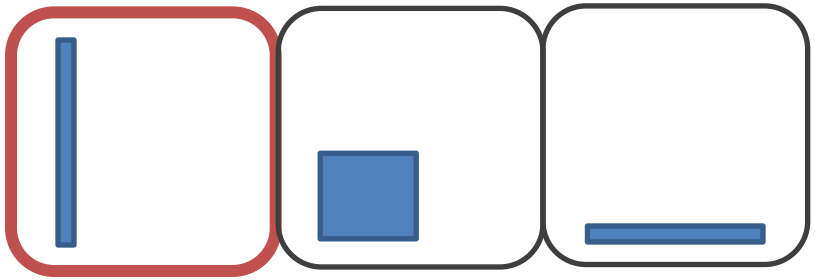

Miners *want* a "medium" limit.

***Demand* is the trade-off.**

Peter Todd Legendary

Funding network security in the future
April 14, 2013, 11:12:09 PM #1

of what happens with the blocksize it's important in the long term: without the block limit we can expect transaction fees to fall to the marginal costs of a transaction, which means the fees aren't paying for any security at all, on the other hand, with a small blocksize limit, as I've been arguing for, you still run the risk that off-chain transaction systems get 'too good' and so few transactions actually happen on-chain that security still isn't being paid for.

Calculus is clear for straight lines

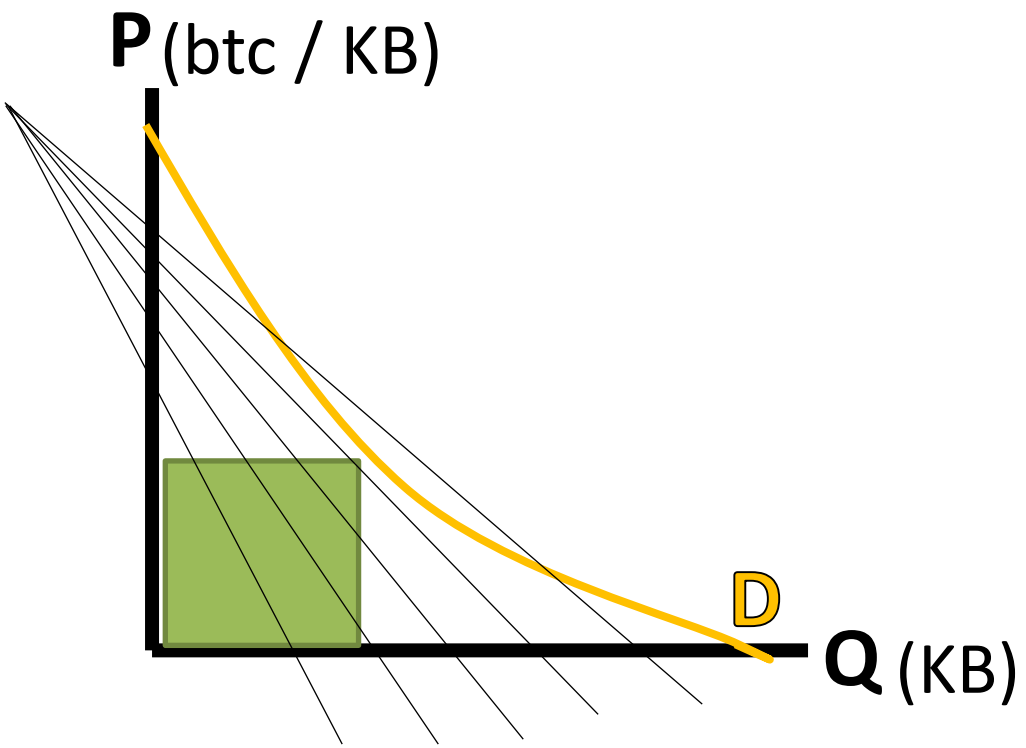

$$\boldsymbol{A = xy}$$

$$\frac{dA}{dx} = y\,; \frac{dA}{dy} = x$$

$$if\colon\ y(x) = b\ - 1x\,,$$

$$A_{ii} = x * (b - x)$$

$$\frac{dA_{ii}}{dx} = b - 2x; x^* = \frac{b}{2}$$

Therefore: $\boldsymbol{y^* = x^* = \frac{b}{2}}$

$$but, if\colon\ m \neq 1,$$

$$\frac{dA_{iii}}{dx} = b - 2mx$$

$$\boldsymbol{x^* = \frac{a}{2m}; y^* = \frac{a}{2}}$$

…a trivial problem.

Calculus is also clear, for parabolas

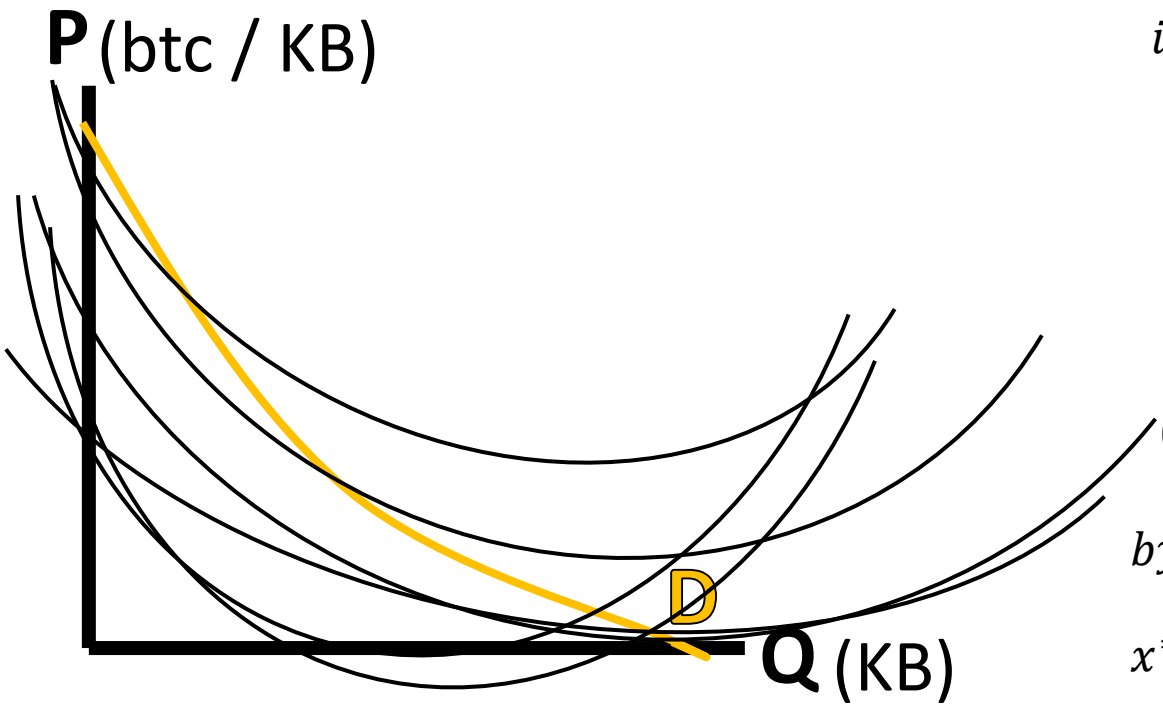

$$A = xy$$

$$if:\ y(x) = ax^2 + bx + c$$

$$A = x * (ax^2 + bx + c)$$

$$\frac{dA}{dx} = 3ax^2 + 2bx + c$$

$$0 = (3a)x^2 + (2b)x + c$$

$$by\ quadradic\ formula,$$

$$x^* = \frac{-2b - \sqrt{(4b^2 - 12ac)}}{6a}$$

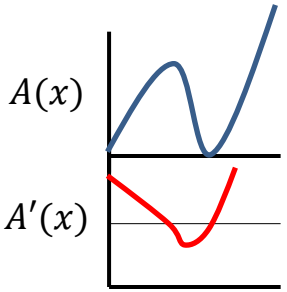

We know to select the minus, because we know x* will be to the left of the vertex (for +a, -b).

Tricky Shape(s)..are possible

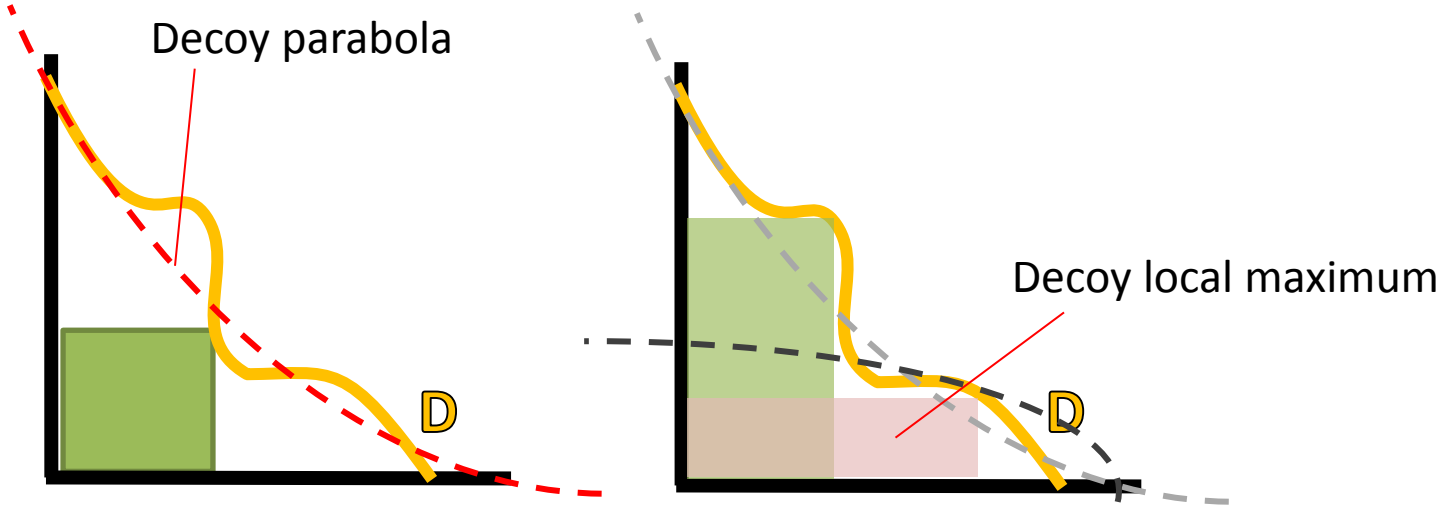

Higher-order (more curves),
more complex.

Possible “Greedy” Rule

1. Adjust Q by 1%.
2. Measure P’s response.
3. If +1% yields <-1%,

   Elif -1% yields <+1%,

   Reverse Course

   Else, Keep Going.
4. Enhancements:
   - Jump magnitude as f(dQ)
   - Averaging (s. difficulty adj)
   - Subjective Miner ‘Leap’
   - Tolerance

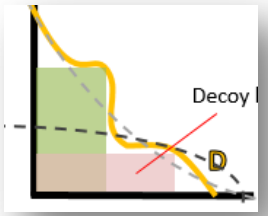

Note: though Q is fixed, tx-selection is not.

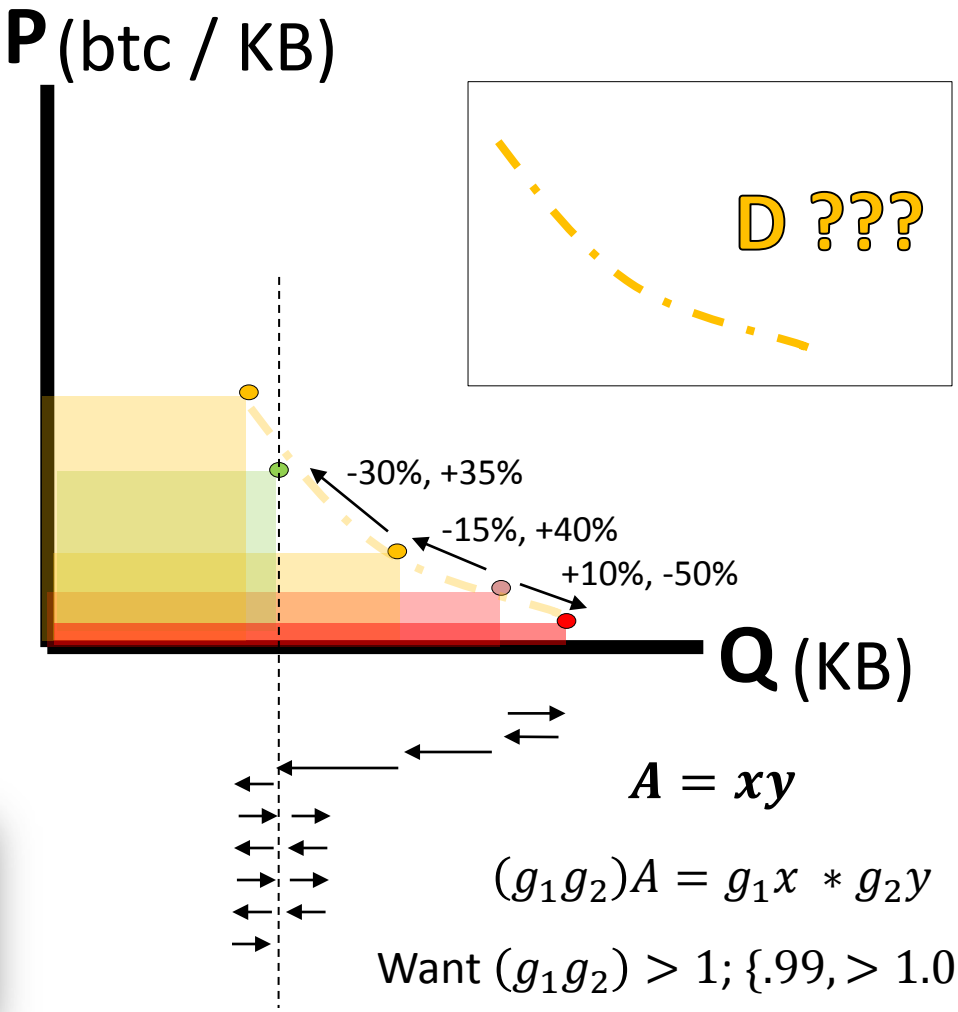

$$A = xy$$

$$(g_1 g_2)A = g_1 x * g_2 y$$

Want $(g_1 g_2) > 1; \{.99, > 1.01\}$

Conclusion: Can Revenues *Fall*, if Q is Controlled by Miners?

- It seems: no. Miners are *willing* and *able* to vary Q s.t. revenues are maximized.

2. Fees

(#4 x #4)

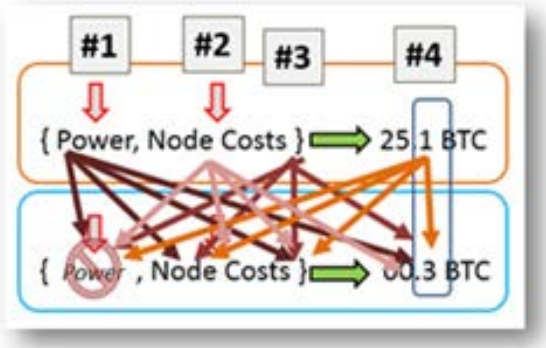

Characteristic	SC Demand Affects Mainchain...	Total Ec
Substitutes	...by removing need for a BTC tx.	?
Independent	...(not at all).	Good –
Compliments	...by inducing need for a BTC tx.	Great !

- My conclusion: *from a fee perspective*, **it is safe** to allow miners to use sidechains to increase Q.
- Even safer, if sidechains have distinct purpose.

Finally: Orphan Commentary (11)

- Concept of ‘orphaning’ is intimately related to our highlighted issues – **bandwidth** and **fee market**.

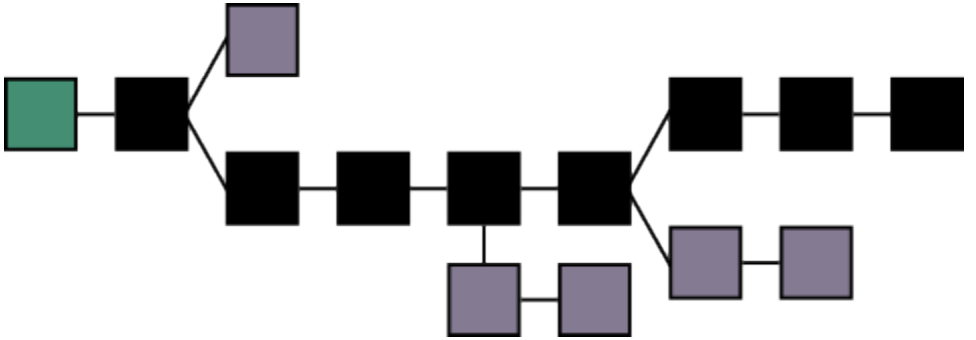

- In fact, it is intimately related to BFT!
- The **game-theoretic problems** which Bitcoin solves are:
  - When someone says “I didn’t get your transaction message” are they **lying**?
  - When someone says “This is *everything* the network decided, while you were gone.” are they **lying**?

One Argument

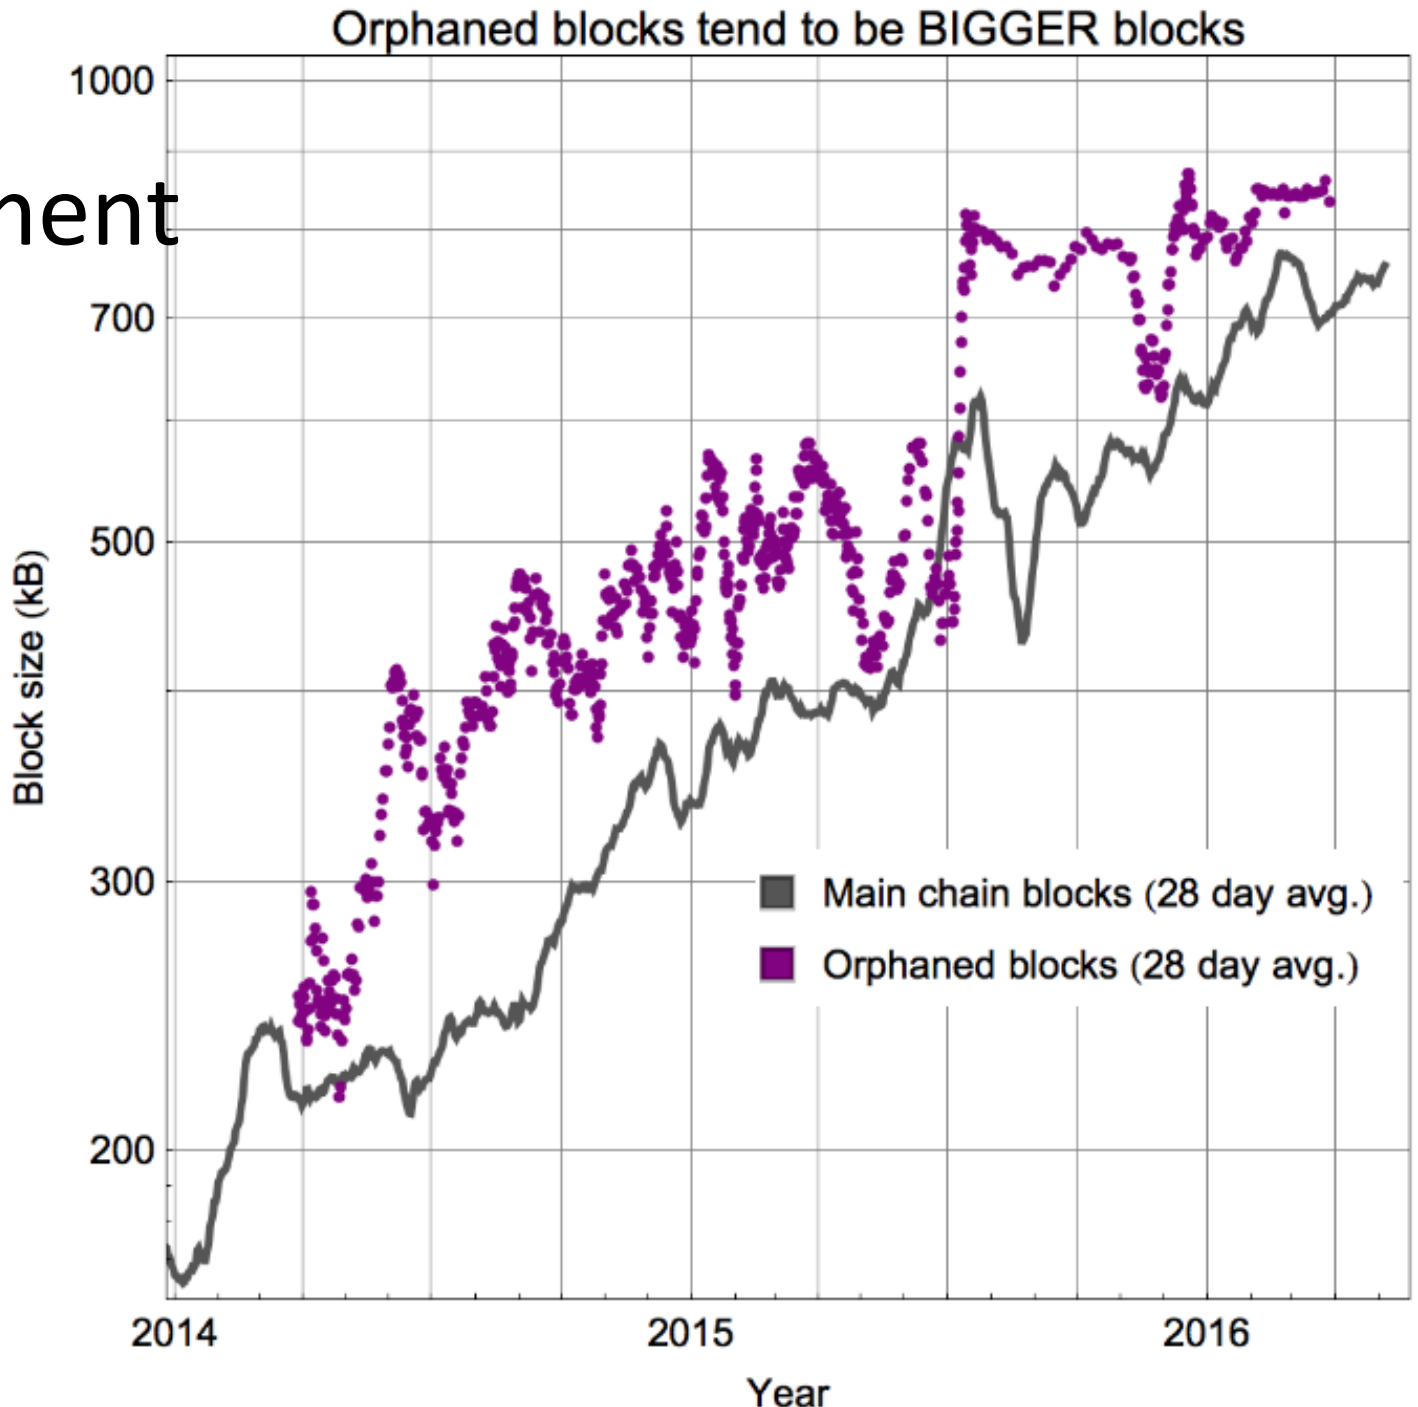

One Argument

Orphan Cost

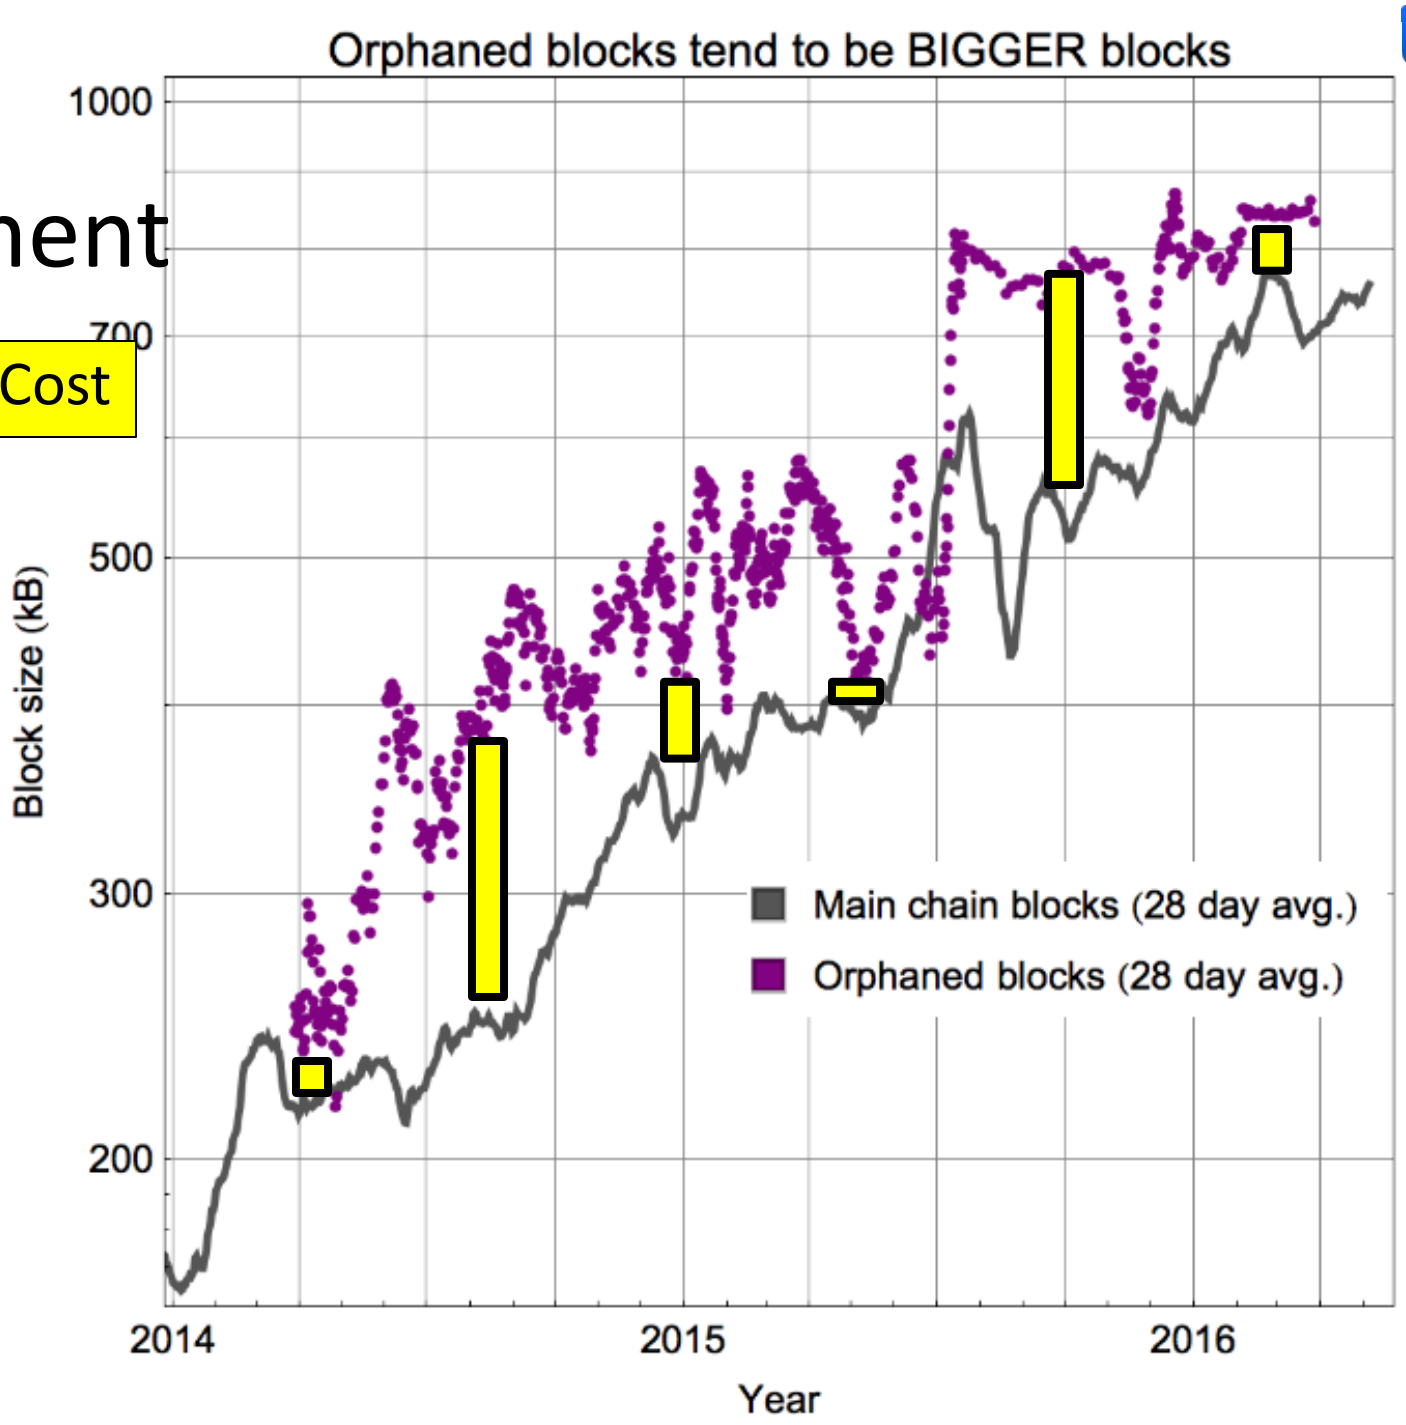

Recall

- My arguments were:
  1. SPV mining can eliminate orphaning.
  2. Orphan costs do not have a significant effect on ‘supply’. Under SPV / scheduled blocks, orphan costs would be zero.

Orphaned Blocks, n=2

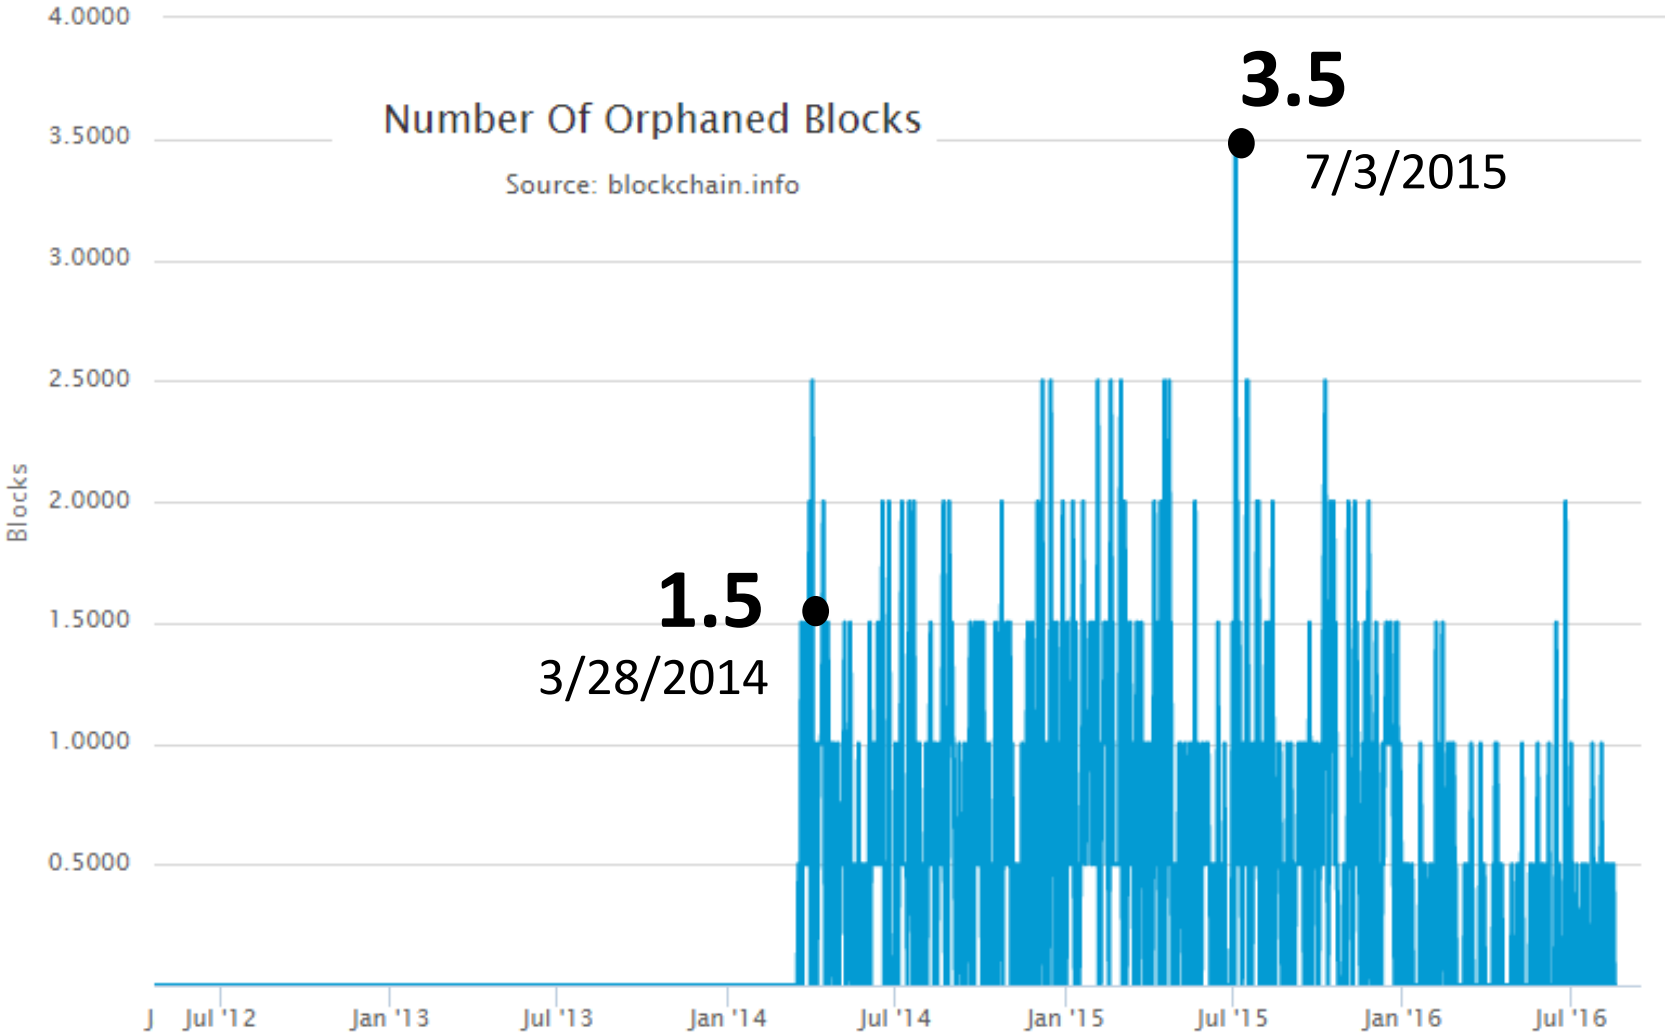

Orphaned Blocks, n=28

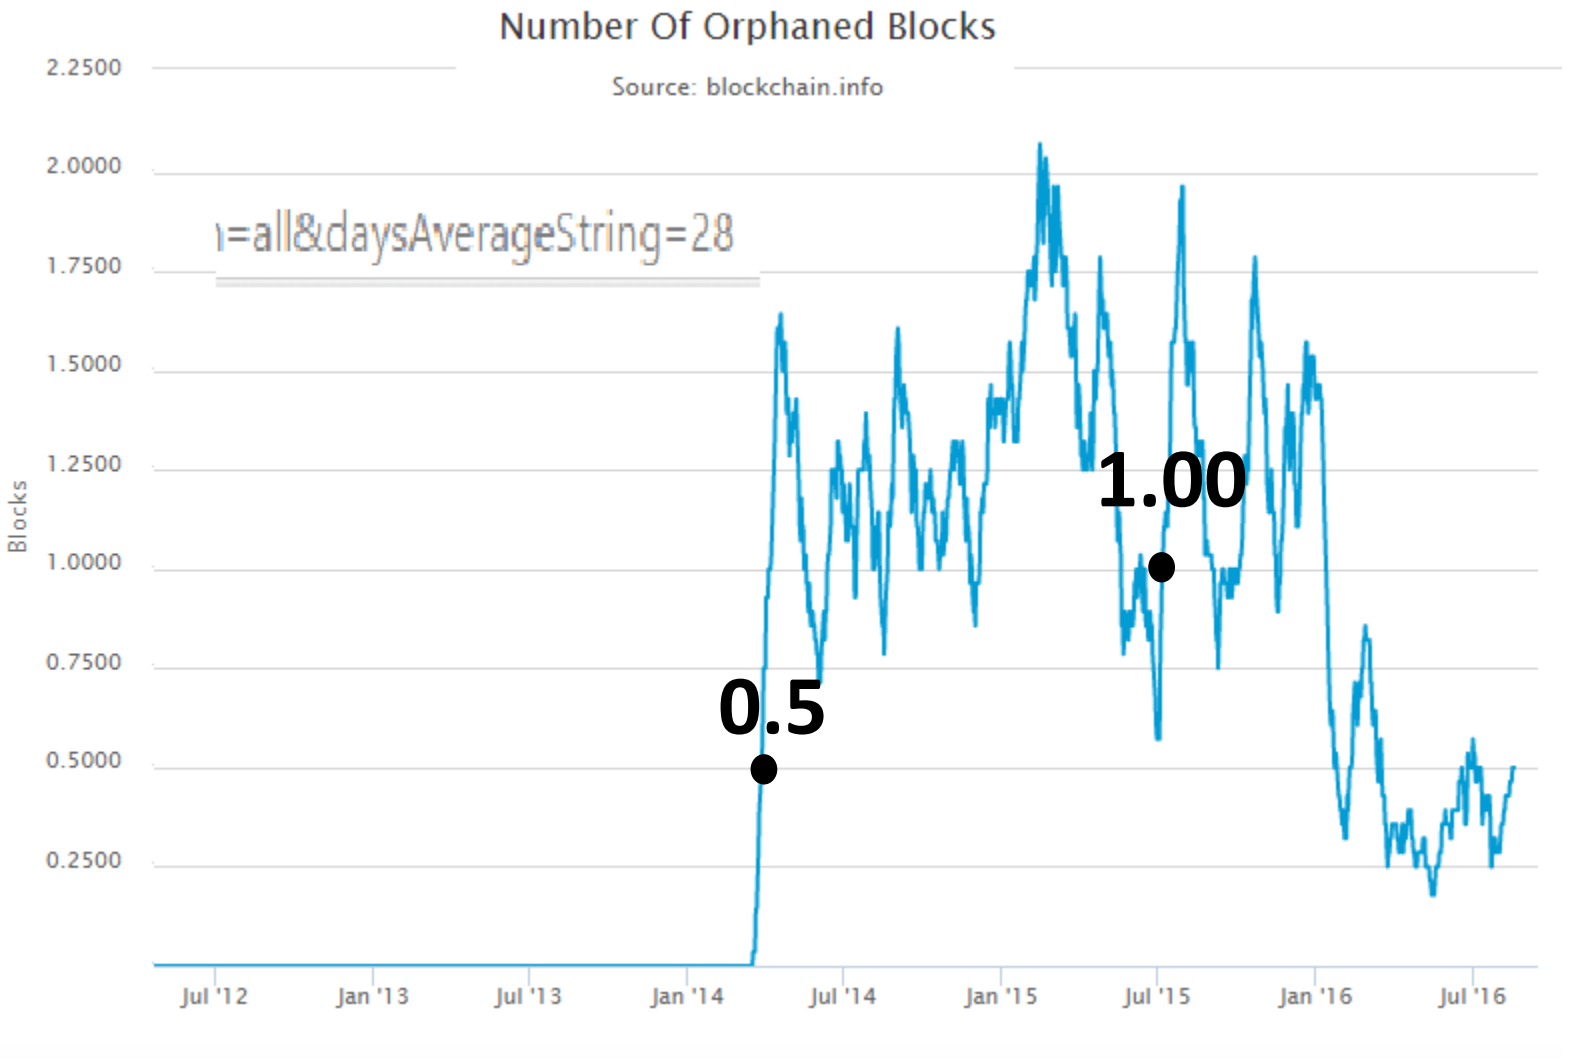

Orphaned Blocks, n=200

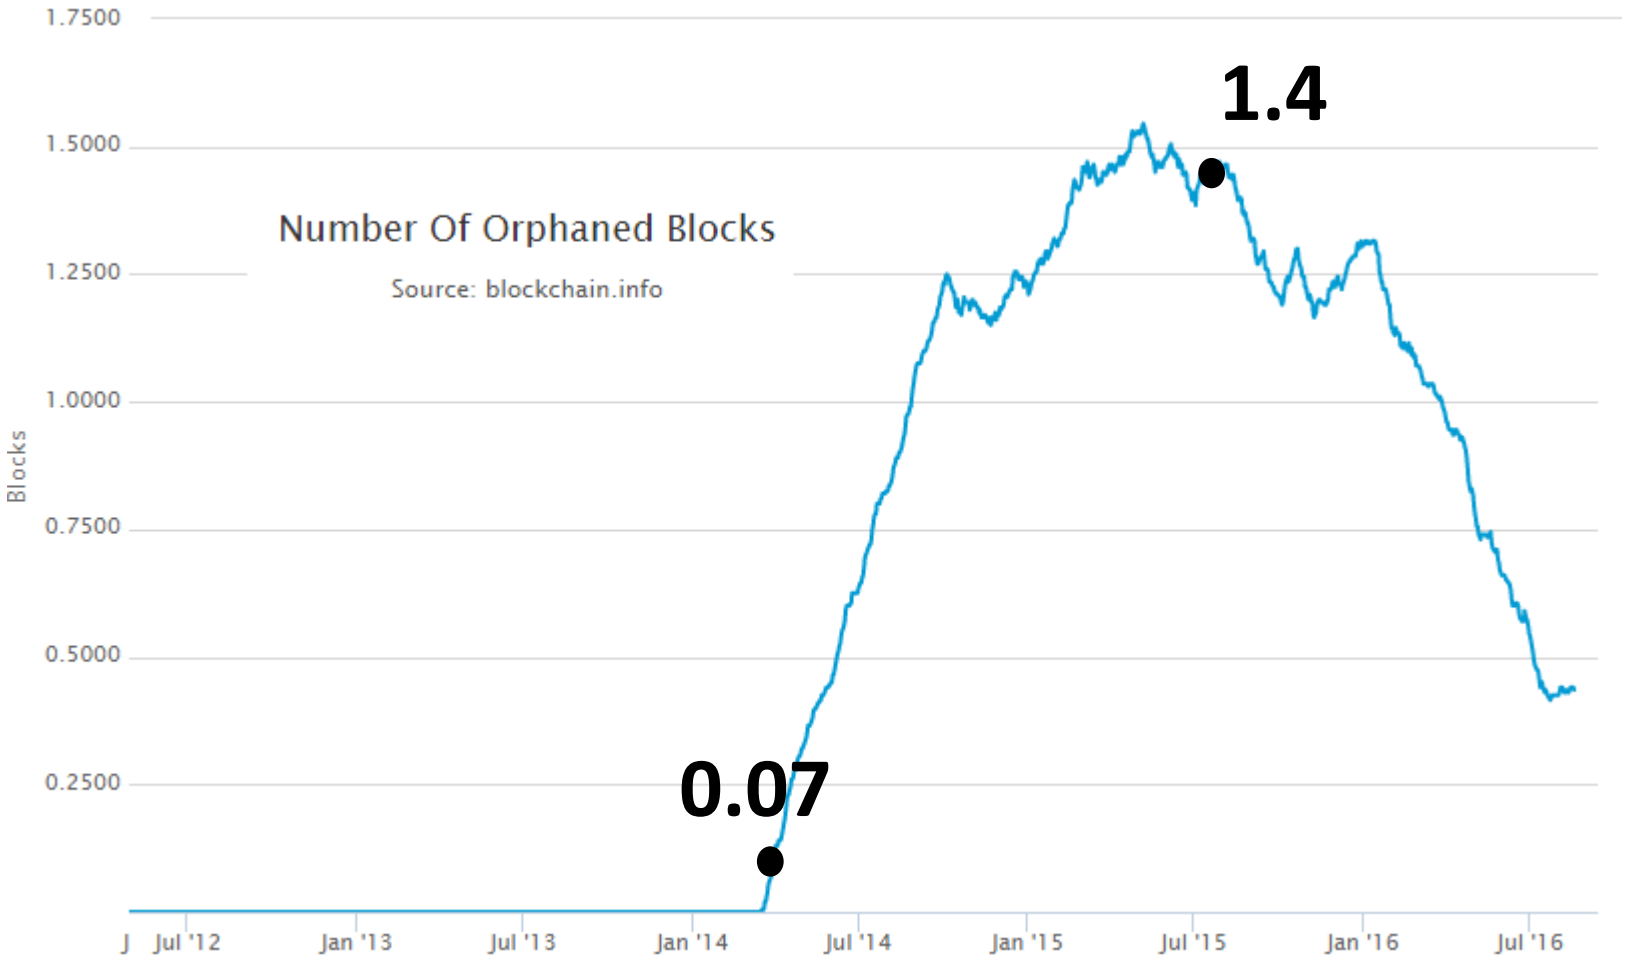

Block Size

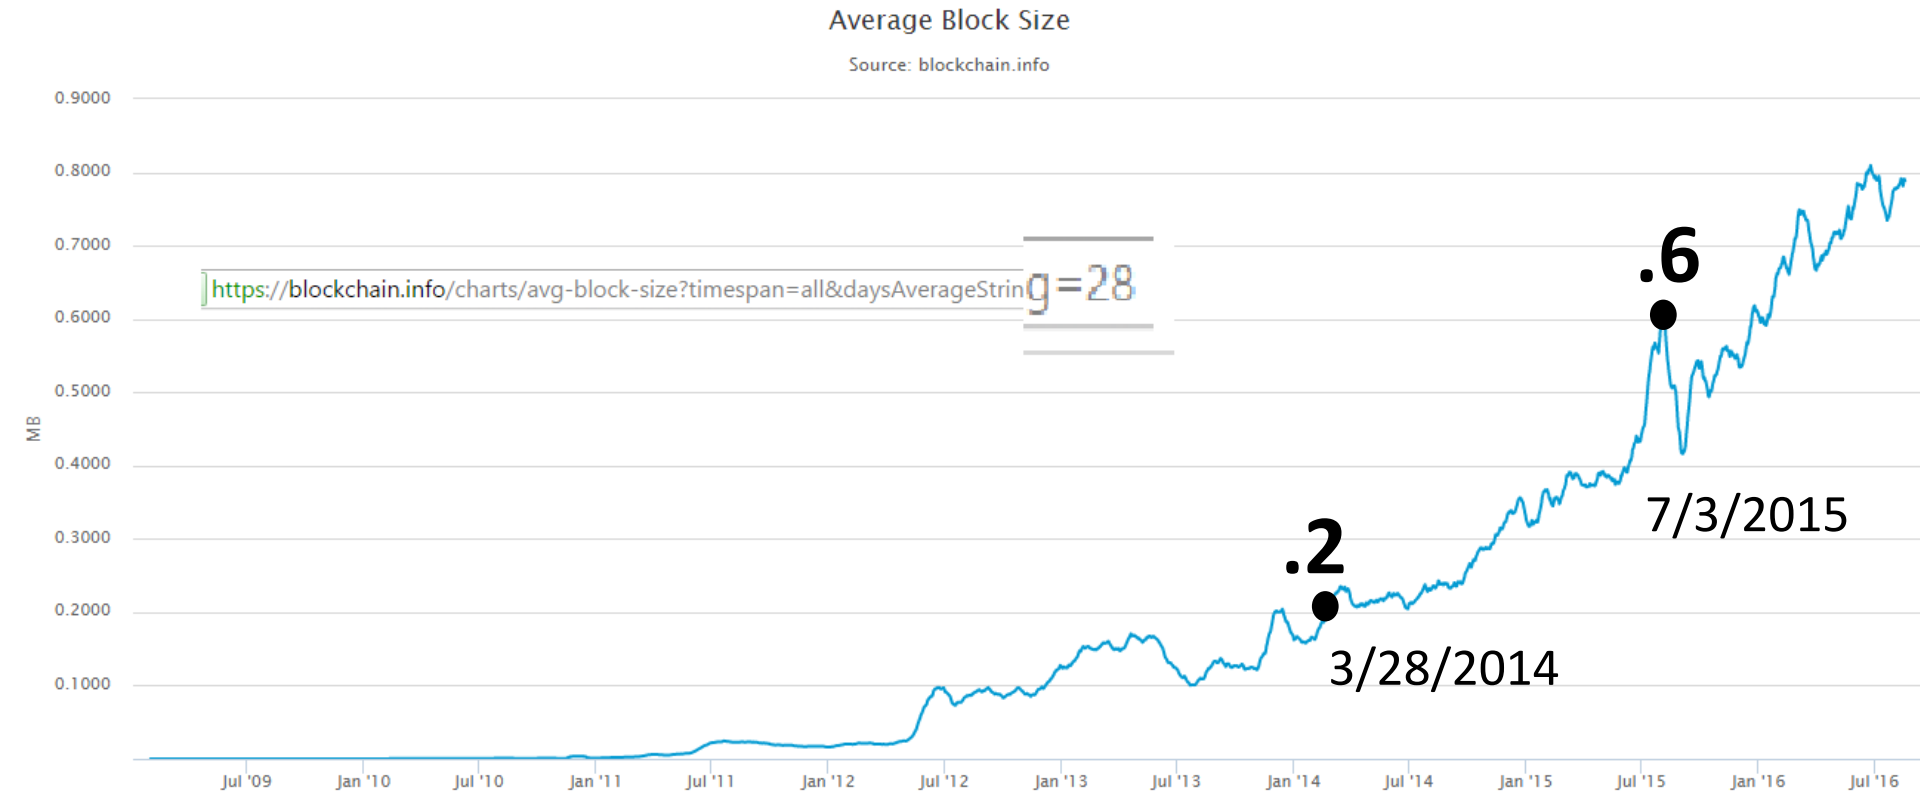
Average Block Size
Source: blockchain.info
https://blockchain.info/charts/avg-block-size?timespan=all&daysAverageStrin
g=28
MB
0.9000
0.8000
0.7000
0.6000
0.5000
0.4000
0.3000
0.2000
0.1000
Jul '09
Jan '10
Jul '10
Jan '11
Jul '11
Jan '12
Jul '12
Jan '13
Jul '13
Jan '14
Jul '14
Jan '15
Jul '15
Jan '16
Jul '16
.2
3/28/2014
.6
7/3/2015

Block Size, n=200

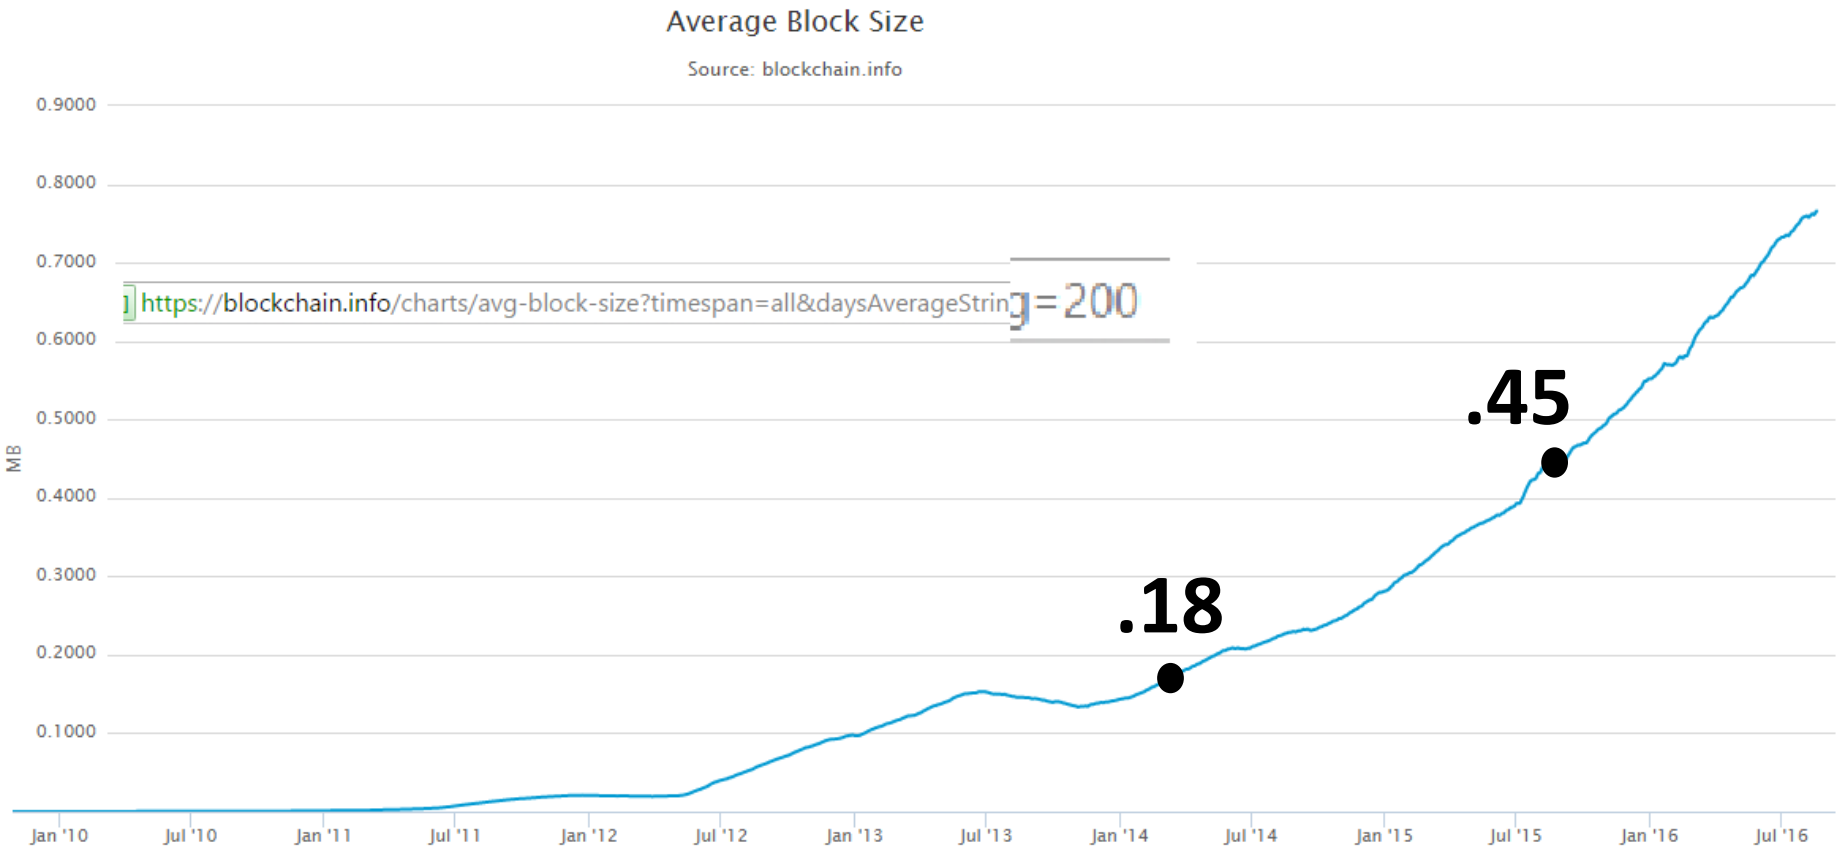

Conclusions

1. Around 3/2014, blocks surged past a 0.2 MB for the first time. At around the same time, orphaning increased significantly.
2. It is known that *some* miners were SPV-mining in July 2015 (re: unexpected PoW-chainsplit).
3. Furthermore, around that time, orphaning decreased sharply…*despite* blocks that were bigger than ever.
4. Likely culprit: SPV mining.

Note: During the worst 200 day period, the orphan-rate averaged **1.5 per day**. ~1% at worst.

Orphaning: Small, transient phenomenon?

2016

Co-founder, BitMain
~16% hashrate

Recall

- My arguments were:
  - SPV mining can eliminate orphaning.
  - Orphan costs do not have a significant effect on ‘supply’. Under SPV / scheduled blocks, orphan costs would be zero.
- What problems has SPV mining caused for Bitcoin?

Conclusion: Sidechain Safety

1. Node Costs
2. Relative to other SF/MM
3. Docile Miners
4. Tame vs. Aggressive
5. Fees and Bandwidth

Just don't run the software.

Node Costs
- Opt-In
- Internalized
- Anti-Fragile

Mining
- "There's only one SH
- How can different ch common network?

namecoin

No one cares!

My Mining Experience?

Interlude: The Docile Miner

I'm still worried

"Tame" Sidechains
1. To Trespass, Need:
1. 51% Miners (else
2. Interested, Fee-Payi
3. Code to have been
- res to Bitcoin.
- Cont s are firewalled – opt in, and don t affect each other.
- Contracts are managed, to maximi BTC value.

"Aggressive" Sidechains
affected. Bitcoin
other for the dur
3. Fundamental Q:
mining activities?
- All th
- Permis ss Innovation.
- No debate er precise "split" of validation resources (MB, SigOps).

Stop stealing my space!

Instant Block Propagation
High, Stable Demand

- **SPV Mining**
- Bloom Filters
- FIBRE
- UTXO Commitments
- **Block Scheduling***
- Clever Transaction Strategies

A block header contains these fields:

Field	Purpose
Version	Block version number
hashPrevBlock	256-bit hash of the previous block header
hashMerkleRoot	256-bit hash based on all of the transactions
Time	Current timestamp as seconds since 1970-01-0
Bits	Current target in compact format
Nonce	32-bit number (starts at 0)

Bandwidth

Fees

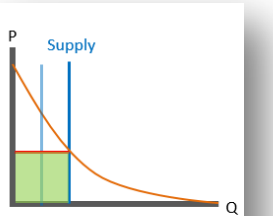

Cartel

Maximized

Thank You!

Paul Sztorc

Sept 2nd, 2016

Cartel isn't That Bad!

- Adaptive
  - When demand rises, it relaxes the constraint.
  - When demand falls, it tightens the constraint.
- Robust
  - Moving the constraint has a cost:
    - Moving Left: Non-included tx fees.
    - Moving Right: Tx fees paid to other miners.
- LN Synergy
  - May allows sale of a new "type" of tx demand.
  - Interacts with LN favorably.

Scheduled Blocks

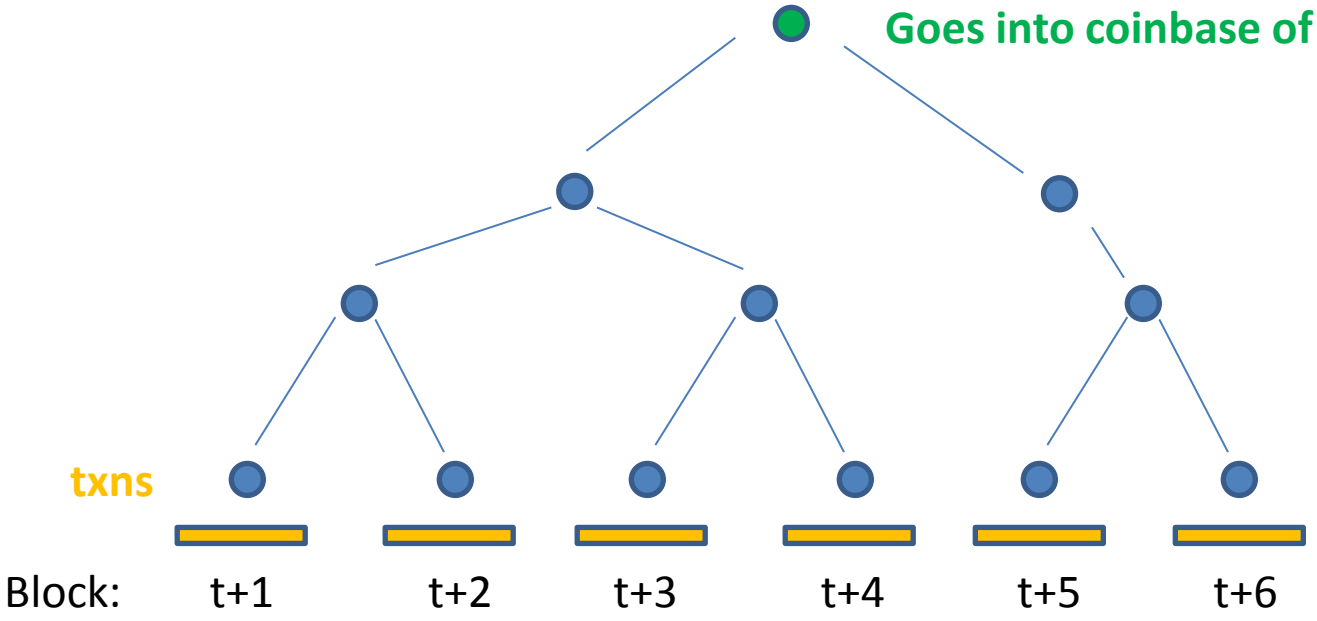

- Miner embeds commitment to a tree – this tree lists "suggestions" for future block contents.
- Suggestions are optional, but miners have motive and opportunity to take them.
- Canonical ordering of txns into block may help.
- For next "10" blks, suggestions are mandatory.

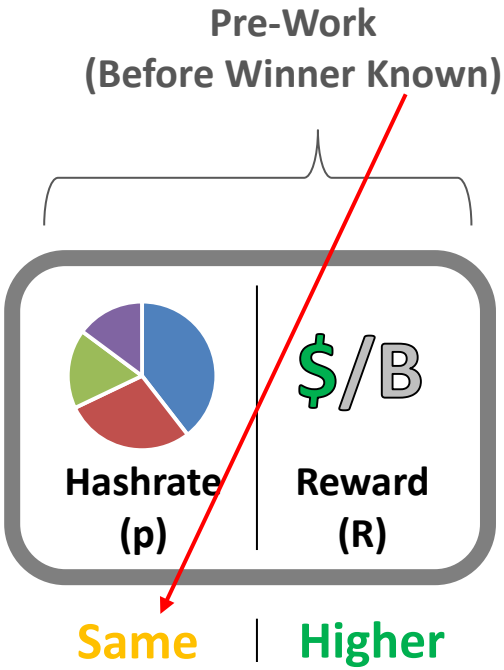

Scheduled Blocks

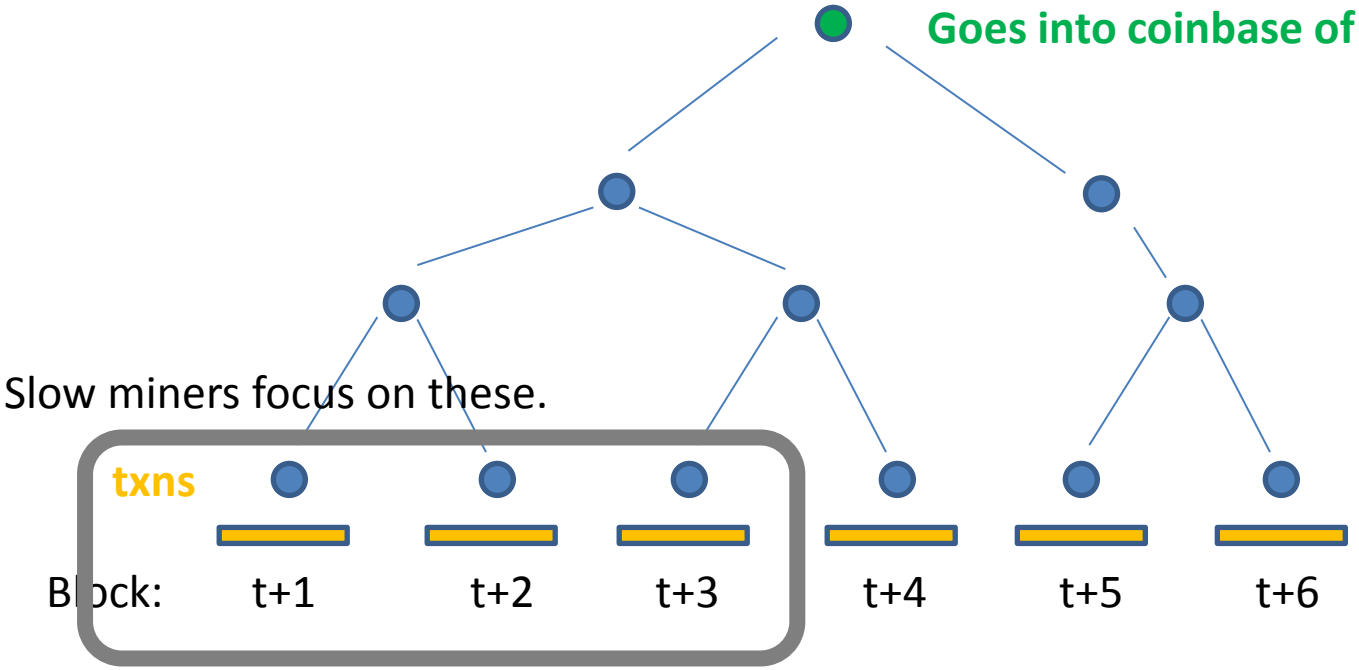

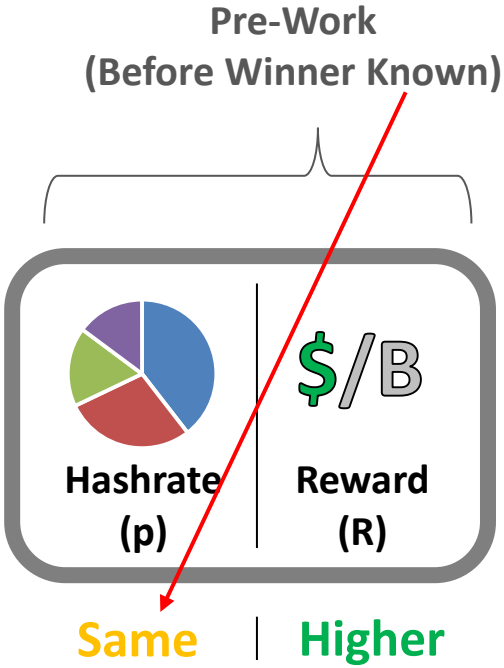

- Miner embeds commitment to a tree – this tree lists "suggestions" for future block contents.
- Suggestions are optional, but miners have motive and opportunity to take them.
- Canonical ordering of txns into block may help.
- For next "10" blks, suggestions are mandatory.

LN Synergy

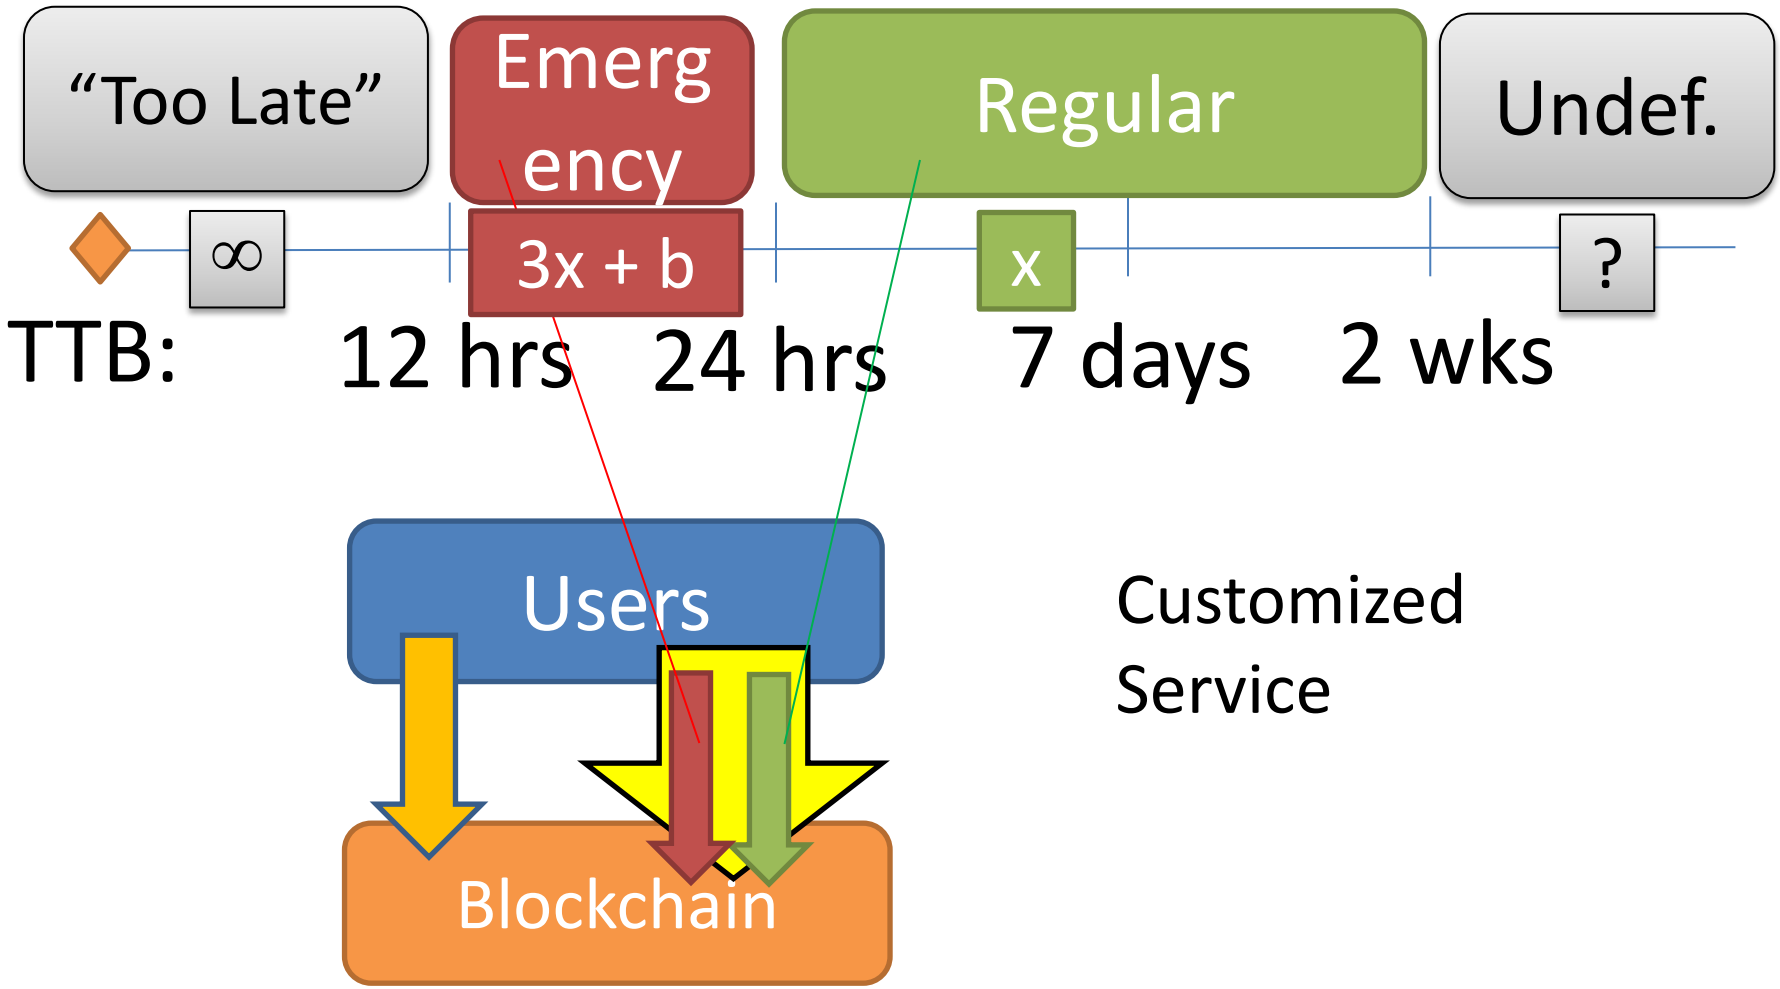
"Too Late"
Emerg ency
Regular
Undef.
∞
3x + b
x
?
TTB:
12 hrs
24 hrs
7 days
2 wks
Users
Blockchain
Customized Service

How?

- Commitment in coinbase.
- Merkle Tree, 2 ^ 11 = 2016 commitments
- So, miner of block #10,400 chooses the txns included in block #(10,400 + (12 $_{hr}$ * 6 $_{h/b}$)).

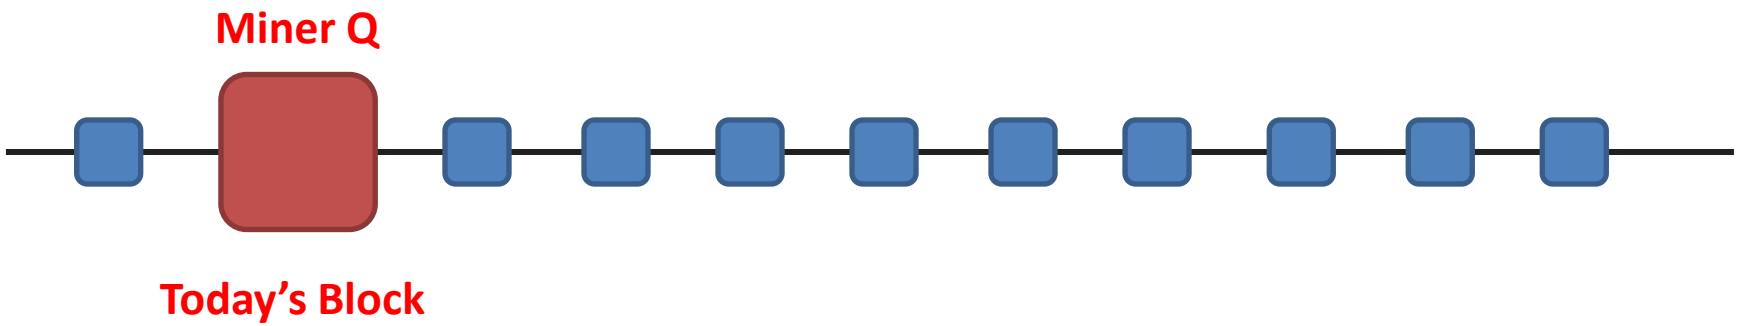

Bandwidth “Storage”

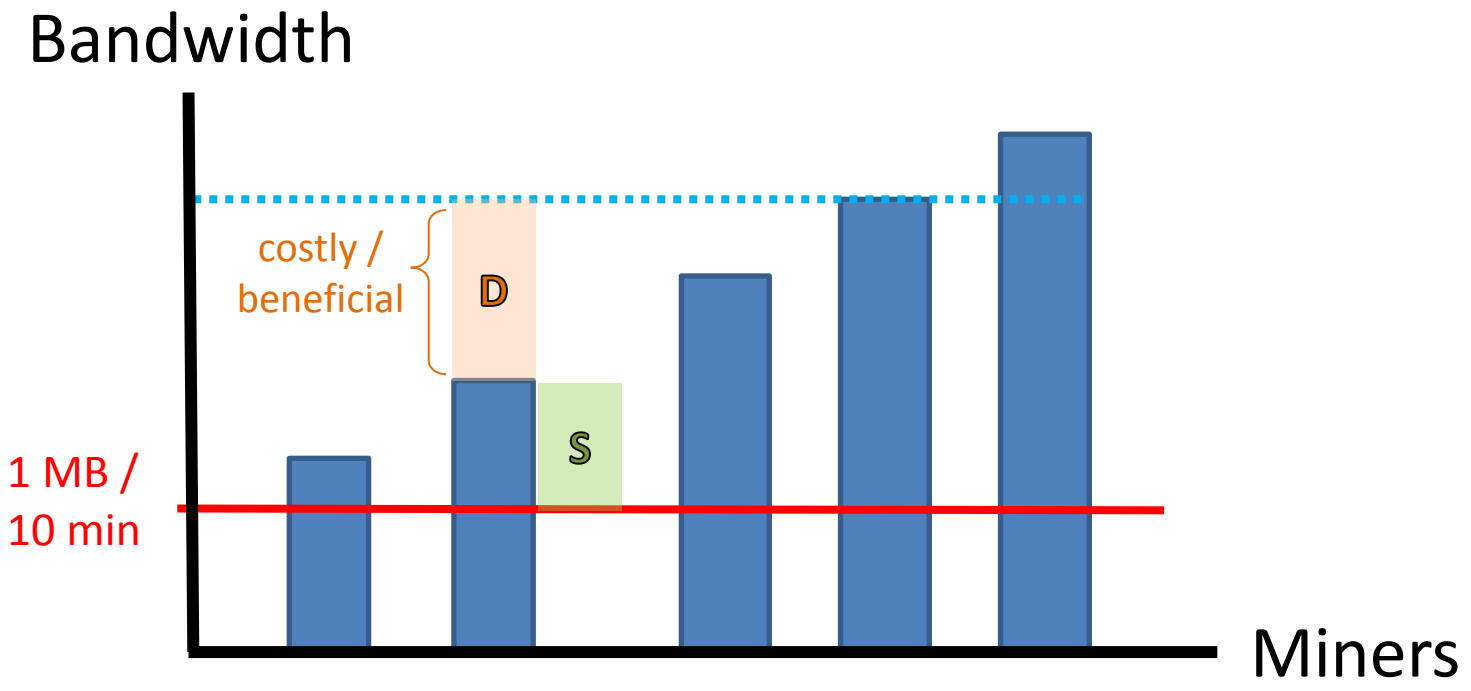

D =

1. “Someone” is unable to hash.
2. Bad for “someone”.
3. Resources/Strategy - not to be “someone”.

Bandwidth “Storage”

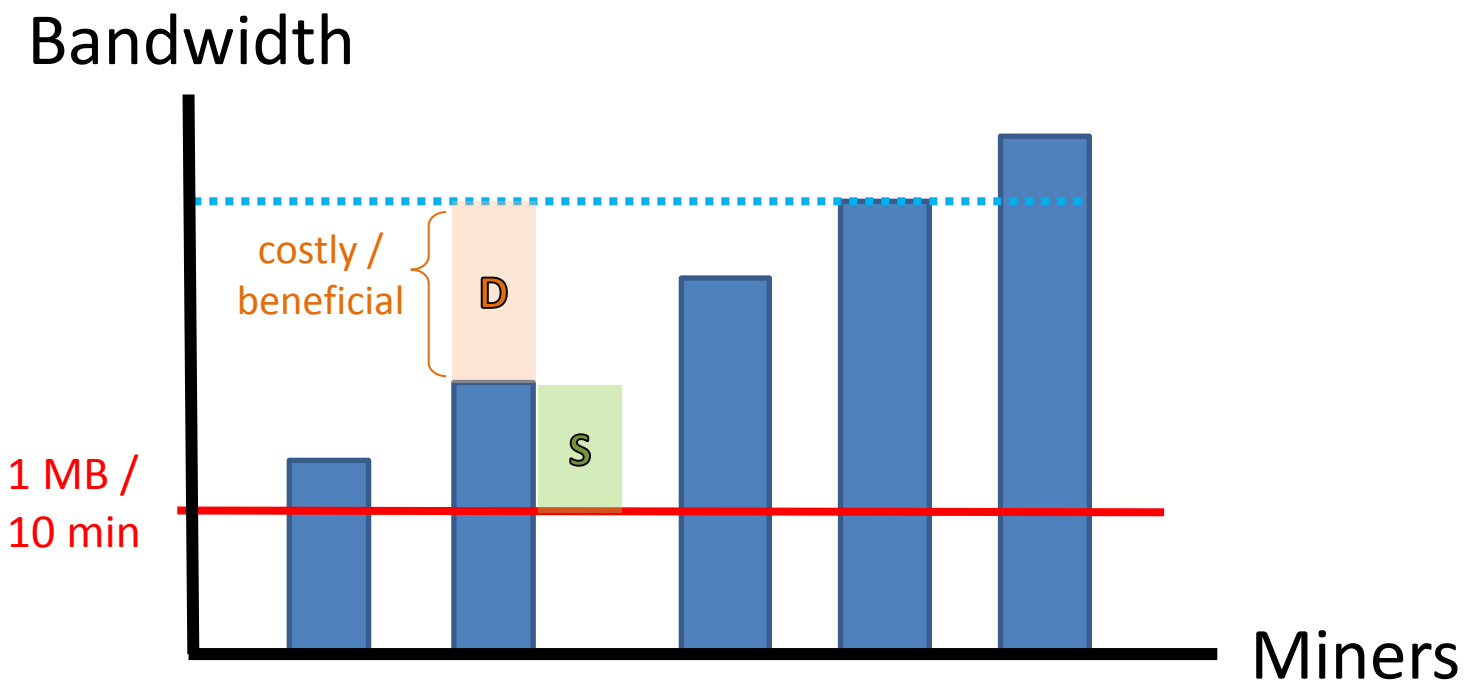

S =

1. Useful during periods of “fast blocks”.
2. “Recovery speed”.

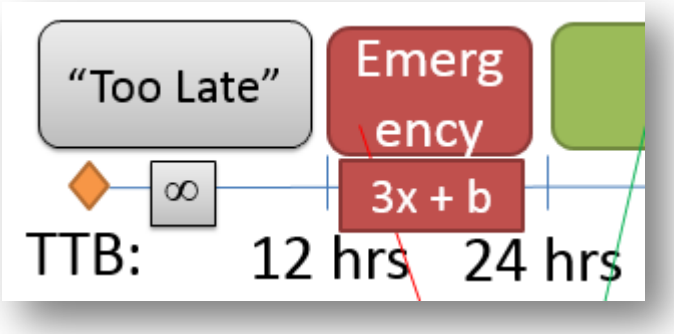

Is this idea any good?

- Might be worth investigating, because this wasn't possible before:
  - Miners would include all known txns
  - No desire to 'buy' a slot which won't be good
- On this, Miners have no reason to betray each other.
  - Preblock, they don't know who will get the tx fees for each block.
  - PreBlock, they lack 'prevblock hash' and can't.
  - So, their Pr(x) is fixed [hashrate], but X payoff isn't.
- Helpful
  - Blocks propagate near-instantly. (No PP).
  - Surplus bandwidth can be "stored up".

SC Context

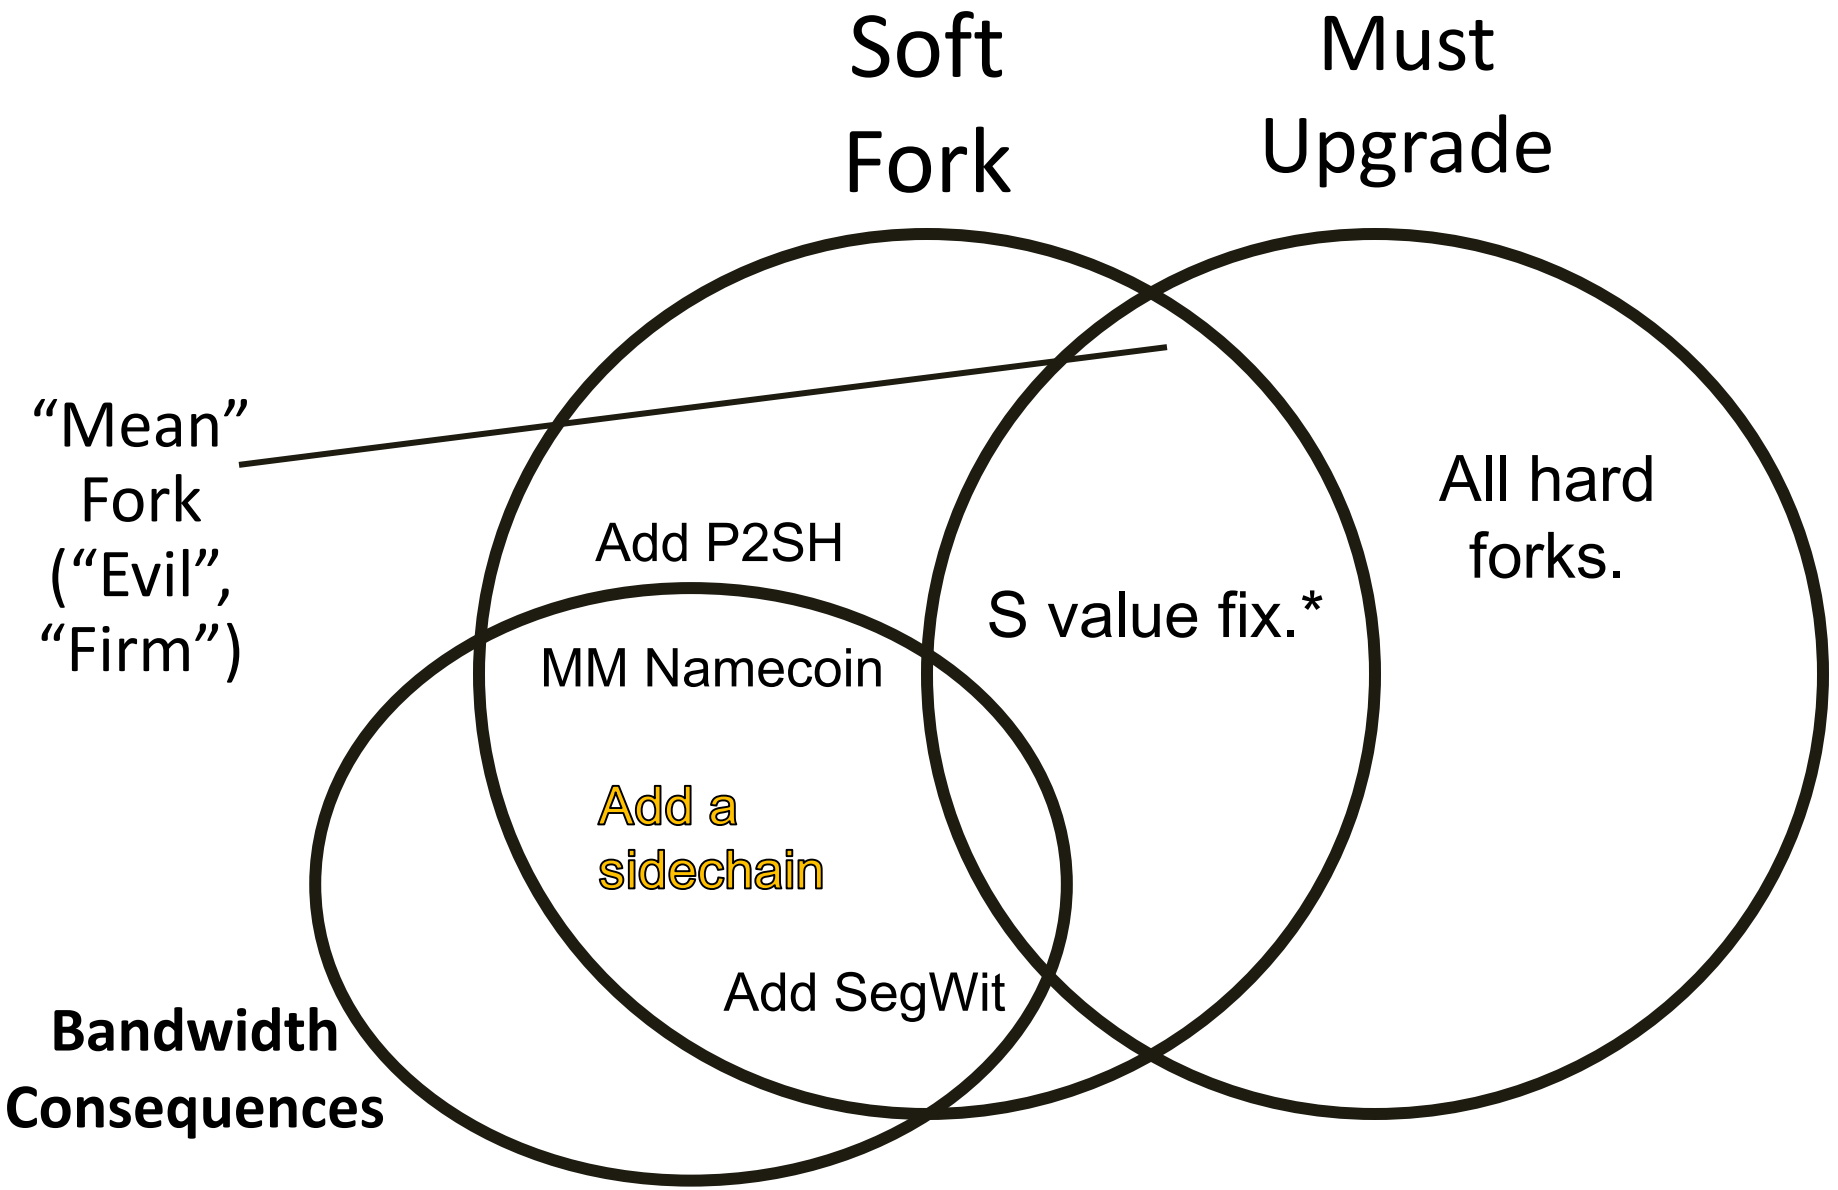

* Had to upgrade to send, but not to receive.