

Why Bitcoin will Fail to Scale?

Nikhil Malik, Manmohan Aseri, Param Vir Singh, Kannan Srinivasan
Tepper School of Business, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213
{nmalik1, maseri, psidhu, kannans}@andrew.cmu.edu

While Bitcoin has garnered enormous attention with its promises of inexpensive, fast and *trust-less* payments, it falls way short of the scale provided by banks. Bitcoin throughput is limited because its ledger accepts a block of fixed *maximum* capacity that can accommodate approximately 2200 transactions every 10 minutes. The Bitcoin community is actively considering technology upgrades to increase the block capacity. Intuitively, one would expect that increasing the block capacity would solve this scaling problem. However, in this paper, we show that increasing the block capacity would be futile for scaling Bitcoin.

We analyze the strategic interactions of miners, who are heterogeneous in their computational power, and users, who are heterogeneous in the value of their transactions, using a game-theoretic model. We show that a relaxation of throughput congestion can facilitate large miners to tacitly collude – artificially reducing the actualized throughput via the strategic partial filling of blocks to receive higher transaction fees. This strategic partial filling is sustained if the computing power of the smallest colluding miner is larger than a threshold. In addition, Bitcoin can only serve 50% of the demand if the colluding group power is above a threshold. We provide empirical evidence of such strategic partial filling of blocks by large miners of Bitcoin.

We show that a technological intervention, such as banning large miners, can eliminate collusion. However, this also makes the system less secure. A strategic adversary faces a trade-off between earning money as a miner and carrying out a double-spend attack to steal a large amount of money in one shot. Such an attack becomes preferable as revenues from mining decrease in the absence of collusion. Therefore, surprisingly, collusion indirectly protects the Bitcoin system by increasing miner revenue. Overall, we show that there is an economic limit to the scalability of Bitcoin. On the one hand, increasing Bitcoin’s capacity invites collusion; on the other hand, if the collusion is suppressed, the system becomes insecure, which drives away the demand regardless. Our analysis raises antitrust concerns with respect to Bitcoin and suggests at least some of the anticipated competitive gains from decentralization in Bitcoin may be difficult to realize.

Key words: Bitcoin, Blockchain, Miners, Collusion, Repeated Game, Fees, MicroPayments, Scalability, Security.

1. Introduction

Bitcoin has been touted as a revolutionary technology that would disrupt the traditional payments industry (Popper 2017, Tasca 2018). Bitcoin’s market capitalization stood at \$100 Billion in Oct 2018 (Zaitsev 2018), and it is being used for transactions worth \$1 Billion on a daily basis (BlockchainExplorer 2018). In June 2018, Bitcoin was used for a \$300 Million payment completed in 10 minutes at a fee of 4 cents (Blockchain.com 2018). In comparison, transactions through banks as intermediaries

are costly, and cross-border transactions take several days to complete (TransferWise 2018). For example, Chase Bank charges a transaction fee of \$40 on cross-border wire transfers from the US (Chase 2018). Merchants pay credit card networks (e.g., VISA and MasterCard) 0.5-2% for domestic retail sales (Dwyer 2018). In recent times, these slow and costly banking intermediaries have also faltered at providing security against data hacks (McMillan 2018, Glazer and Farrell 2018) and frauds (Glazer 2018). As a result, Bitcoin has garnered enormous attention with its promises of inexpensive, fast and *trust-less* payments (Jordan and Kerr 2018).

Traditional banks maintain a ledger of transactions. Users make payments by instructing banks to add desired transactions onto this ledger. Banks guarantee integrity (no one can spend more than their balance) and security (outsiders cannot steal balances). Bitcoin maintains a similar ledger, but it is maintained by consensus among a peer-to-peer network of Bitcoin users. Bitcoin's current technology is capable of adding only up to 3 transactions per second (Eyal et al. 2016) to this ledger, compared to VISA's 5000 transactions per second (Yli-Huumo et al. 2016). In December 2017, this limited throughput fell well short of demand; a vast majority of users were turned away, and a large number of users were unable to spend any of their Bitcoins. On Dec 21, 2017, under severe congestion, Bitcoin witnessed users offering up to \$54.90 to obtain preferential treatment on this limited throughput (Buntinx 2017).

Bitcoin's ledger consists of a chain of blocks. Blocks with a limited size of 1 MB (~ 2200 transactions) are added every 10 minutes. This low frequency of block addition is required to propagate a new block over a large network to keep all participants in sync. As a result, to increase the transaction throughput, Bitcoin developers have been aggressively debating technology upgrades for increasing the block capacity. Segwit2x was an attempted upgrade of Bitcoin to 2 MB blocks in Nov 2017 (Bitcoin Wiki 2017). Meanwhile, Bitcoin Cash was launched as a competing platform in Aug 2017 with an 8 MB block size (Wilmoth 2018). In this paper, we ask a fundamental question – Would increasing the block capacity allow Bitcoin to scale?

Counter-intuitively, we show that increasing the block capacity would be futile in terms of scaling Bitcoin. Bitcoin's ledger is maintained by thousands of participants, called "miners", who use special software to validate transactions. These miners in large numbers provide ledger integrity and security. Miners spend computational resources to validate transactions for which they collect a fixed block reward plus any transaction fees offered by the users. Although transaction fees are optional, miners can choose which transactions to process and prioritize those that pay high transaction fees. As a result, users tend to offer high transaction fees when the block capacity is well below the demand to obtain preferential treatment. However, when the demand competition decreases, users do not face

the same risk of being left out. As a result, they offer near-zero transaction fees. We show that under such a scenario, a large number of miners could enter a tacit collusion to artificially lower the effective block size to raise transaction fees. They would strategically partially fill the blocks by including only those transactions which offer high transaction fees. We identify conditions under which such tacit collusion is sustainable. This strategic partial filling is sustained if the computing power of the smallest colluding miner is larger than a threshold. A large miner invests more in computational power and, as a result, wins blocks more frequently. These large miners are most suited to sustain collusion via a long-term punishment threat. In contrast, small miners win so infrequently that they are not deterred by future punishment and thus deviate from collusion. We also show that Bitcoin can only serve 50% of the demand if the aggregate power of the colluding group is above a threshold.

There are two aspects of the Bitcoin's mining process that are necessary for such a collusive equilibrium. First, the public ledger allows anyone to identify which miner added a particular block and their choice of full or partial block fill. This allows other miners to verify if a colluding miner has deviated. Second, the miners are in the game for perpetuity. The group can punish the deviating miner by getting into a transaction fee war in perpetuity, making all miners worse off, which ensures that no colluding miner deviates.

A block size increase will be futile in the presence of such a collusive equilibrium. The rationale behind block size increase is to reduce transaction processing congestion. However, this would reduce transaction fees offered by the users. Faced with low transaction fees, miners could collude and artificially lower the effective block size. The Bitcoin community can respond to miner collusion via technology interventions. Banning large miners is one potential intervention to make collusion unsustainable. We show that while eliminating collusion (via banning miners) may increase throughput and lower transaction fees, this makes the system less secure from double-spend attacks. In a double-spend attack, a miner sacrifices block revenue for a small probabilistic shot at stealing a payment value. This sacrifice of block revenue is a small price to pay once the collusion is eliminated. Powerful miners would then prefer to launch double-spend attacks rather than earning these small revenues. Eliminating collusion removes artificial capacity constraints; however, the security threat drives away the demand for transactions. Thus, these collusion and security threats place economic bounds on Bitcoin's scale.

A closer examination of Bitcoin's ledger confirms the above intuition – miners indeed partially fill their blocks when demand competition is weak. Figure 1 shows block sizes pushing on the ~ 2200 transaction ceiling around Dec 2017. Since then, block sizes ($\sim 500 - 1500$ transactions) in 2018 have been well below their maximum limit. This seems to suggest insufficient interest to make payments

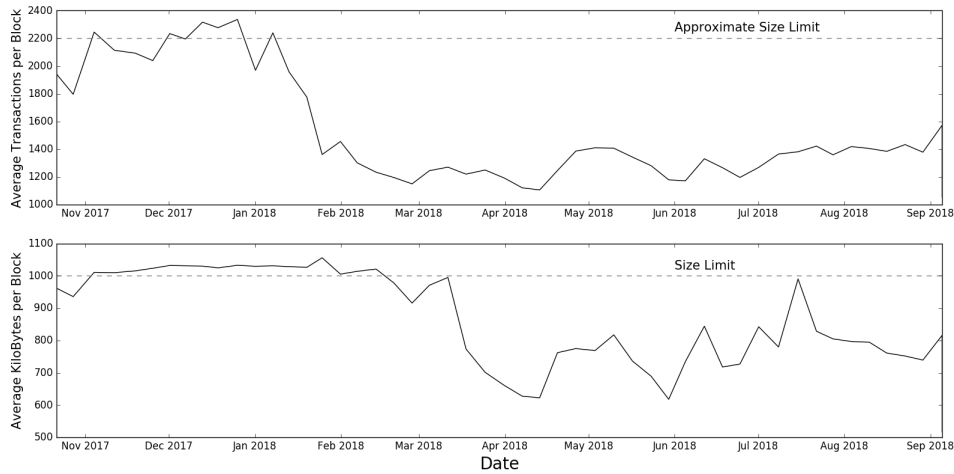


Figure 1 The average block size has been significantly below full capacity since Jan 2018, shown in terms of the number of transactions (top) as well as the actual size in KB (bottom).

via Bitcoin. In reality, user payments are still being turned away as empty blocks are added to the ledger. Numerous partially filled blocks have recently been observed regularly, even when tens of thousands of transactions were waiting to be processed. Figure 2 shows a very specific instance on the 31st of September 2018 when AntPool (which is a very large mining pool) added partially filled blocks. The first two blocks used up less than 10% of the block capacity, while more than 1000 transactions were in the queue. It is critical to note that very few transactions in waiting were offering a high fee (> 15 satoshi¹/Bytes) when these blocks were added. The third block utilized its full capacity when a significant number of high-fee-offering transactions were in the queue.

Figure 3 reports the criticism received by Antpool on Twitter for partially filling their blocks when thousands of transactions were waiting to be processed. Figure 4 provides the reply of an Antpool representative who admits to this practice and responds to the criticism by saying that their actions are within the rights provided by the consensus mechanism of Bitcoin.² We will show that the strategy of sacrificing present capacity for future earning surges is rational for large miners only. In

¹ 1 Bitcoin = 1×10^8 satoshi

² An alternate explanation for the partially filling of blocks could be related to the computation time advantage. A miner is required to prepare a new block with a set of awaiting transactions. This requires a small amount of computation time to validate all newly added transaction. Some miners prefer to forego the verification and therefore the transaction fees by preparing a small block to get started on the puzzle solving as soon as possible. The validation time and puzzle solving time trade-off should be equally beneficial for all miners. Partial block filling by a select few large miners contradicts this explanation.

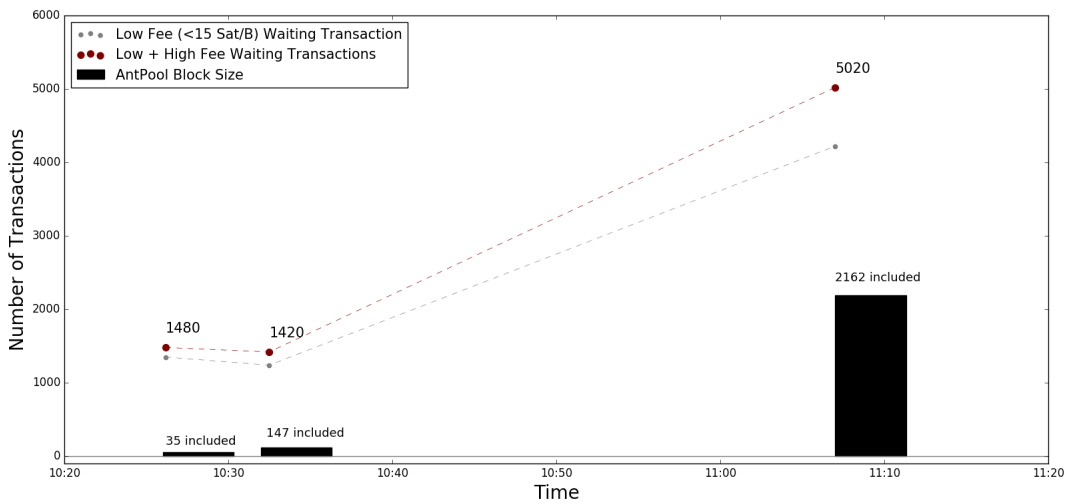


Figure 2 Three blocks mined by Antpool on the 31st of Sep 2018. The first two blocks used up less than 10% of the block capacity, while more than 1000 transactions were in the queue. The third block utilized its full capacity when a significant number of high fee (> 15 satoshi/Byte) paying transactions were waiting, as indicated by the dark gray region.

line with this, Figure 5 shows that AntPool (which currently holds 20% of the overall mining power) has been under-filling blocks compared to all other miners.



Figure 3 Two instances where a large bitcoin mining pool (AntPool) was criticized on Twitter for mining partial blocks (298 KB and 240 KB) while more than 60000 transactions were waiting in queue.

Our work has three major contributions. First, we provide economic bounds on the scalability of Bitcoin. This shows the vulnerability of decentralized platforms in terms of serving user demand efficiently without a leader at the helm. Limited research and industry debate are focused on technological challenges with respect to large block propagation over P2P networks. Second, we are one of the first to simultaneously model rational decision choices for users and miners. Existing research

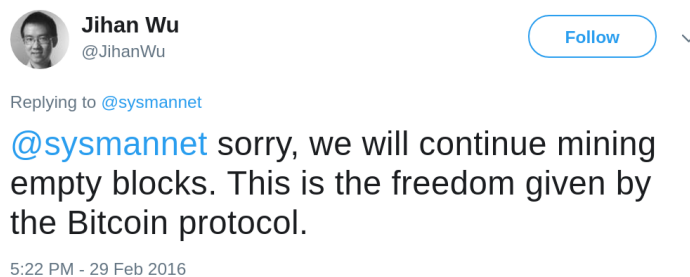


Figure 4 The figure shows a reply by an Antpool representative with regard to criticism concerning the partial blocks.




Figure 5 In periods of high demand (Pre-Jan 2018), AntPool fully fills their blocks like any other miner. In periods of low demand (Post-Jan 2018), AntPool has under-filled blocks compared to all other miners.

largely focuses on strategies for one of these agents (usually miners) in isolation. We show how user preferences (banks vs Bitcoin) impact miner earnings and how miner block filling impacts user fees and security. These externalities are highlighted because we model the two sets of agents in a single game. Finally, we provide rationale to recent surprising events on major public blockchains and hopefully provide clarity to future directions for these platforms. In addition to these major contributions, we add to the growing literature on information security economics (Gao et al. 2013, Cezar et al. 2013, Kannan and Telang 2005, Arora et al. 2007, August et al. 2014, Hsu et al. 2012, Dey et al. 2018). This literature focuses on the strategic interactions of firms and hackers. Our work is a unique setting whereby a firm is replaced by decentralized miners and users. Our work also intersects with literature on P2P platforms (Asvanund et al. 2004, Wei and Lin 2016, Li and Agarwal 2016). Johar et al. (2011) study participant equilibrium sharing responses to lower congestion in

P2P platforms. In Bitcoin, miners control the supply side; they instead prefer to increase congestion in this P2P setting.

In the next section, we describe the Bitcoin ledger and consensus mechanism in greater detail and highlight relevant prior literature on the economics of Bitcoin. Section 3 presents the model to illustrate Bitcoin's transaction fees as an outcome of user competition to earn space on the limited ledger throughput. We extend this base model to show the equilibrium collusion strategy sustained by a group of large miners. Finally, Section 4 discusses potential methods to break the collusion and the perilous security implications.

2. Background

Bitcoin attempts to offer ledger integrity and security without an intermediary. Its ledger of transactions is maintained by a group of *miners*. Anyone can join this miner group using a computer connected to the Internet. Every miner retains his own copy of the ledger and keeps it in sync with other copies. Miners validate and add a block of new transactions on average every 10 minutes. Miners compete to solve a computationally expensive cryptographic puzzle (Nakamoto 2008). The cryptographic puzzle used by Bitcoin requires the miners to find a SHA256 hash, which satisfies a difficulty condition, of their block of new transactions. Each block also stores the hash of the previous block, which ensures that blocks are added chronologically. The fastest solver gets to add the next block on top of the existing chain. This miner broadcasts his new block to the network. All other miners update their copies of the ledger after validating the puzzle solution as well as the block content. If a majority of miners accept the block, the state of the ledger is seen to be updated. These block addition steps are repeated infinitely to grow the ledger.

Unlike banks, Bitcoin does not set a fee. Users seeking transaction processing enter an auction to find a spot on the limited throughput (transactions per sec) of the Bitcoin ledger. Miner participation is incentivized by these transaction fees.

Figure 6 depicts a simplified example of Bitcoin's ledger. Each block records details of transactions, including the sending and receiving of account holder information, transaction fees offered by each sender for a transaction, and the cryptographic signature of the sender, which is needed to verify that the transaction is genuine. The block also records a transaction called the "generation transaction". The winning miner is awarded with a block reward (newly minted Bitcoins), which is recorded as a generation transaction. The balance in an account is determined by adding up all transactions that appear in the unique longest chain of blocks. Because Bitcoin is a peer-to-peer network and because block propagation takes time, two miners who are far from each other in the network can

often find a solution to the puzzle at approximately the same time. Both miners could propagate the solution to peers. This may lead to parallel chains. This situation is resolved endogenously. A chain with support from the miners with the largest combined computing power grows faster than others, eventually becoming the longest chain. This is because higher computing power leads to faster puzzle solving and therefore more block additions³. Once a longer chain of blocks emerges, all shorter chains are abandoned.

Users make payments by broadcasting desired transactions to all miners. Because of the fixed block size, not all transactions pending at a given time can make it into the next block. The consensus mechanism of Bitcoin provides a miner full autonomy over which pending transactions to include in a block. Users hope that a sufficiently high fee offer will incentivize any miner to include their transactions over other pending transactions. The winning miner collects these transaction fees as well as a fixed block addition reward. This dual block revenue and the exhausted computational resources incentivize puzzle-winning miners to scrutinize transactions carefully. If they add an invalid transaction (e.g., a user spending more than their balance), other miners will reject their block, wiping out the block revenue.

In addition to adding valid transactions with high transaction fees, a miner could increase his earnings by winning puzzles more frequently. As a miner invests more in computing power, they become increasingly faster at solving the mining puzzle. A select few miners have gained access to high-quality ASIC hardware specialized for Bitcoin mining. These large miners make up a large proportion of the total computing power. The remaining computing power comes from a large crowd of small miners utilizing inferior hardware (e.g., GPUs and CPUs). The mining network overall has become faster at solving puzzles due to increases in computational power devoted to mining Bitcoins. This computational power has increased due to the joining of new miners and existing miners upgrading their hardware. The puzzle difficulty is therefore adjusted endogenously to keep the average time to a solution at 10 minutes. A natural question arises – why not have 1 minute between blocks? The puzzle winning miner needs to transmit a new block to the entire mining network. Because of the peer-to-peer nature of the network, blocks propagate slowly, and a short duration gives too little time for all miners to sync with the new state of the ledger. A shorter time between blocks could lead to the creation of multiple parallel chains at any time. In addition, users

³ In an extreme setting, a miner with 51% of all computing power can unilaterally create the longest chain without agreement from other miners. Such a powerful miner can add any invalid transactions that they wish. However, outsiders will lose faith in Bitcoin in such a case, thus making any Bitcoins stolen by the miner worthless. Such large miners are typically economically invested in Bitcoin's success by buying expensive computing hardware. An outsider who wants to hack (steal user balances) the ledger requires a prohibitively expensive investment to gain control. Thus, a group of anonymous miners are able to offer the integrity and security of Bitcoin's ledger.

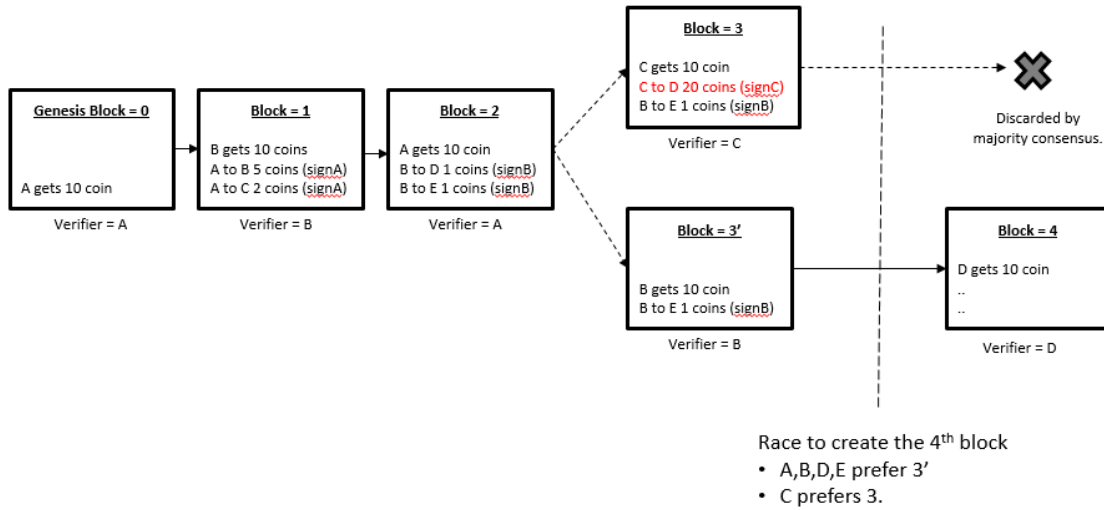


Figure 6 Left to Right: The 0th or Genesis block records a single fixed block reward transaction. The 0th, 1st and 2nd blocks are mined by miners A, B and A, respectively. These miners are quickest to solve the three puzzles, respectively. Each miner receives 10 coins as a fixed reward for block creation. They also collect fees offered by respective transactors. The 3rd block is proposed by C. Because of an invalid transaction whereby C attempts to spend 20 coins they do not have, this block is rejected by all other miners. An alternate 3rd block is then proposed by the 2nd fastest puzzle solver B. The chain is extended on top of this block by a majority of the mining community. The rewards earned by miner C on their proposed block do not form a part of the longest unique chain.

would need to wait for a longer chain to emerge to determine whether their transactions have made it to a block in the longest chain. Hence, reducing the time between blocks would not necessarily reduce the effective transaction validation time. As a result, increasing the block capacity is being considered by the Bitcoin community as one of the key ways Bitcoin could potentially scale.

Prior Literature: The economics of Bitcoin has gained attention relatively recently. Huberman et al. 2017 discuss a congestion pricing mechanism that drives Bitcoin’s transaction fees. Our work draws from some of their findings but focuses more on rational equilibrium choices for (user and miner) agents. Biais et al. 2018 model MPE strategies for miners to show equilibria with and without forks. We study simultaneous miner actions in the context of block filling instead of fork choices. Cong et al. 2018 study equilibrium mining pool sizes. Malinova and Park 2017 highlight the unique public address identity on Bitcoin, which allows users to consistently exhibit their actions while maintaining anonymity. This is a crucial feature that allows small miners to sustain a grim trigger collusion in our model.

The academic literature has discussed a few adversarial attack strategies on the mining network. A double-spend attack (Sompolinsky and Zohar 2015) attempts to re-spend a single coin by temporarily

expanding a large amount of computing power to create forks in the Bitcoin chain. A selfish mining attack (Nayak et al. 2016) attempts to repeatedly discard blocks created by honest miners to gain an undue mining reward. An eclipse attack (Natoli and Gramoli 2017) spreads mis-information on the peer-to-peer network by isolating certain target honest miners. All these attacks would require an extremely powerful miner and would likely result in a long-term loss of faith in Bitcoin for the users. We show a stable collusion strategy performed by small miners that simply extracts economic rents rather than interferes with the ledger’s integrity. We contrast gains from collusion versus attacks for an individual with high computing power. The literature quantifying security threats and hacks on blockchains is limited. Gervais et al. 2016 compare gains from double-spend attacks under different blockchain upgrades via simulations. Rosenfeld 2014 touch upon the economics of a double-spend attack.

In regard to firm collusion in repeated interactions, the management literature has a lot to offer. Miners in our context offer the exact same product, i.e., space on blocks. Rothschild 1993 show the ease of collusion when products are substitutes, while Chang et al. 1991 and Ross 1992 show the opposite. Häckner 1994 shows easier collusion with vertically differentiated products but the opposite with horizontal differentiation. Thomadsen 2007 study collusion as a function of firm homogeneity. Miner homogeneity impedes collusion in our model. However, this is more an outcome of very small individual miners when the network is forcibly homogenized. A different strand of marketing literature (Athey and Ellison 2008, Zhu and Wilbur 2011) studies advertisers who participate in ad space auctions and a handful of publishers (e.g., Google) that offer this space. There are similarities with our user-side auction for payment completion. However, unlike publishers, who have numerous pricing tools to differentiate their ad space, miners have very few levers to maximize their revenues.

3. Model

We model Bitcoin’s blockchain as an infinitely repeated block creation game between miners and users. In Section 3, we first obtain the equilibrium transaction fee offered by users and the endogenous entry of miners to participate in the mining network. Next (in Section 3.1), we examine if an individual large miner can be better off by performing partial block filling. In the absence of this large miner, we inspect if a group of small miners can achieve the same outcome. In Section 3.2, we obtain the conditions for a stable tacit collusion.

Miners are modeled as heterogeneous in quality of computing hardware. Only a few miners have access to high-quality ASIC hardware, while most miners can only access lower quality GPUs or CPUs. We assume that when arranged in decreasing order of hardware quality, a miner with rank m

has a mining power $h(m)$, measured in number of cryptographic *hashes* performed per second. This hardware distribution $h(m)$ over rank is a monotonically decreasing convex function⁴ with domain $m \in [0, \infty)$ and $h(\infty) = 0$. The convexity represents a long tail, i.e., a large number of potential miners own a small CPU or mobile computation device⁵. A large number of hashes per second means a greater likelihood of finding the Bitcoin mining puzzle solution faster than other miners. A miner expects to win the mining puzzle with a probability equal to their share of the total hash power in the network. As more miners enter the network, each miner's puzzle winning probability decreases, as they all share the same pie. We model miner entry as a zero-barrier event, i.e., miners enter until the marginal miner makes zero profits. Miners with better hardware (high h) will naturally crowd out worse hardware (low h) because of their competitive advantage.

Bitcoin users are modeled as heterogeneous in value of transaction $v \sim U[0, V_{max}]$ ⁶. Assume that N users want to complete their payments every period. Each user can either complete his transaction using Bitcoin or an off-chain fiat option. The off-chain option (e.g., VISA and MasterCard) allows near instantaneous completion of the payment. However, Bitcoin transactions need to wait for block addition on the chain, which takes a fixed time. This decreases the user's utility by a factor of δ (< 1). Users also face an uncertainty in successful inclusion of their transaction on the blockchain. Their transaction may not be picked up on the next block. In this case, they revert back to the off-chain option. The user's utility from the available choices is given by equation (1). The off-chain option is modeled as a proportional (ρv) fee structure. This off chain fee structure is driven by costs (e.g. liquidity guarantee, fx margins, money laundering checks) that depend on payment value⁷. Further Shy and Wang 2011 show why traditional intermediaries achieve greater profits via proportional fees. The on-chain fee f is endogenously determined by the demand and capacity of the blockchain.

$$U(v, f) = \begin{cases} v(1 - \rho), & \text{off chain,} \\ (v - f) \times \delta, & \text{on chain included,} \\ v(1 - \rho) \times \delta, & \text{on chain excluded.} \end{cases} \quad (1)$$

Let \mathcal{F} represent the fees corresponding to the transactions included by a miner on a block. The winning miner earns the sum of fees ($\sum_{f \in \mathcal{F}} f$) offered by transactions that she includes in her block. A miner also earns a constant block reward (B). The platform has a hard upper limit on the number

⁴ $h(m) = \lambda e^{-\lambda m}$ is one potential distribution.

⁵ The convexity can also be interpreted as heterogeneous access to electricity, i.e., ability to run more of the same hardware for a given cost.

⁶ V_{max} is the largest payment a user wants to make, say \$10 Mn. In practice, very high value payments may be rare. A different distribution (e.g. $v \sim exp(\lambda)$) does not change our findings.

⁷ More recently, fintech start-ups have started to offer fixed fees options by avoiding these variable costs. We do not model these smaller players.

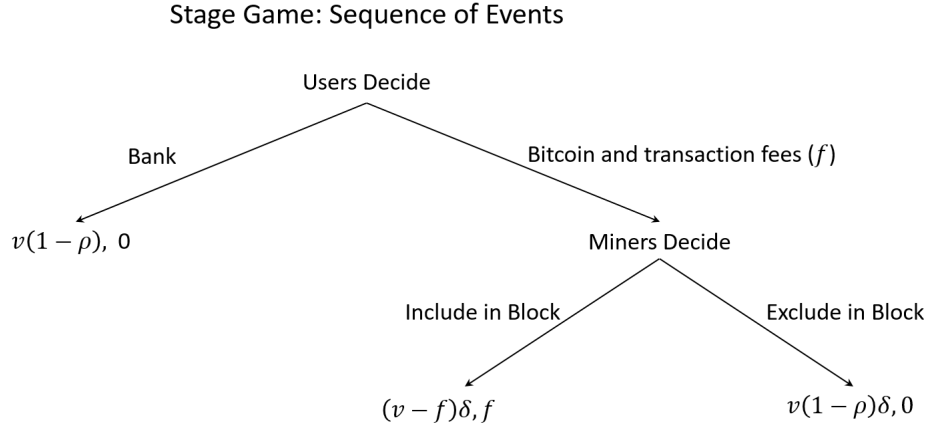


Figure 7 The figure shows the single period extensive form game between the users and miners that is infinitely repeated at every block creation period.

of transactions that can be included; we let n_F represent this capacity. This platform capacity n_F is considered as an exogenous and static choice by the blockchain designer. A passive miner is defined as one who uses up the full block capacity n_F to cater to user demand. Let $\gamma = \frac{n_F}{N}$ represent the ratio of capacity to demand. Let \mathcal{M} represent the set of miners that decide to enter and mine on the Bitcoin network. Then, the expected profit of a miner with computing power h_j and cost c_j can be written as

$$E[\pi(\mathcal{F}, h)] = \begin{cases} (B + \sum_{f \in \mathcal{F}} f) \times (h / \sum_{j \in \mathcal{M}} h_j) - c_j, & \text{Enter,} \\ 0, & \text{Do Not Enter,} \end{cases} \quad (2)$$

where “Enter” (resp., “Do Not Enter”) represents a miner’s decision to enter (resp., not enter) the mining network. Figure 7 represents the sequence of events in a block creation game. Miners and users take sequential steps to create a single block on the chain. Users decide between off-chain and on-chain options. Users who take the on-chain option broadcast their transaction to all the miners with a fee offer. All users make this choice simultaneously with rational expectations of other user and miner actions. Next, miners prepare a block independently by selecting transactions from the waiting queue up to the blockchain’s capacity. The block prepared by the puzzle-winning miner is added to the blockchain.

We define the following constants, which will be used throughout the paper:

$$\gamma = \frac{n_F}{N}, \quad \alpha_h = \frac{\delta + \rho - 1}{\delta \rho}, \quad \alpha_l = \alpha_h \left(\frac{1}{\gamma} - 1 \right), \quad \text{and} \quad \beta = \frac{(1 - \gamma)(\alpha_h - \alpha)}{\alpha}. \quad (3)$$

We first investigate subgame perfect equilibrium (SPE) strategies for a single stage using backward induction. Miners will pick the top n_F transactions in the order of fee offer if the demand N is higher than the capacity n_F . In an auction-like manner, users determine their fee offer based on rational




Figure 8 Users with transaction value higher than v_0^* transact on the Bitcoin network.

beliefs of user heterogeneity. We focus on a pure strategy equilibrium whereby a user's fee offer strategy is an increasing function $f(v_1) \geq f(v_2)$ if $v_1 \geq v_2$. Users with large payment values v pay a higher fee ρv for the off-chain option. These users stand to gain more by avoiding this proportional bank fee. Consequently, the winning miner includes the top n_F transactions in the order of their values. Let v_0 be the transaction value such that there are exactly n_F users that transact a greater value. Since N transactions are uniformly distributed between 0 and V_{max} , we have

$$\begin{aligned} v_0^* &= V_{max} \left(1 - \frac{n_F}{N}\right), \\ &= V_{max}(1 - \gamma), \end{aligned} \quad (4)$$

where γ is defined in (3).

All users with the top n_F values ($v \geq v_0$) offer a fee such that the remaining users ($v < v_0$) are driven out. The marginal user with $v = v_0$ is indifferent between the on-chain and off-chain option. Thus, we have

$$(v_0^* - f_0^*)\delta = v_0^*(1 - \rho),$$

where f_0^* is the transaction fee offered by the user with transaction value v_0^* . Using the above equation and expression of v_0^* from (4), we obtain

$$f_0^* = V_{max}\rho\alpha_h(1 - \gamma), \quad (5)$$

where α_h is defined in (3)

Let R_0 represent the revenue earned by miners. Then, we have

$$R_0 = f_0^*n_F = V_{max}(1 - \gamma)\rho\alpha_h\gamma N. \quad (6)$$

Clearly, R_0 is a concave quadratic function of γ , with a maximum value at $\gamma = \frac{1}{2}$. Since we know that $\gamma = \frac{n_F}{N}$, the miner revenue maximizing capacity of Bitcoin is $n_F = \frac{N}{2}$.

Proposition 1 *Bitcoin favors the entry of high-value users. A large capacity enables more users to transact on the Bitcoin network at a lower fixed fee.*

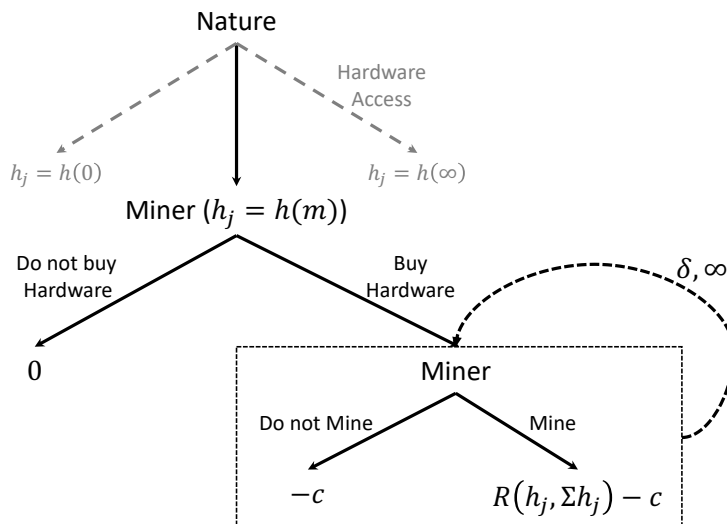


Figure 9 Miner j first realizes access to hardware computing power h_j such that the miner of rank m has access to the computing power of $h(m)$. Next, the miner decides whether to buy the hardware for mining. If they buy, it is only rational to mine every block period thereafter. The period revenue is based on one's own power relative to everyone else's power. The same payoff is repeated infinitely with a discount factor δ .

Table 1 Power and costs for popular computing hardware. ASICs are more expensive but have exponentially higher power.

	$h(m)$	$c(m)$	
	Hash Power (GH/sec)	Power Consumption (Watt)	Hardware Cost (USD)
AntMiner S9 (ASIC)	14000	1375	2400
Avalon Batch 1 (ASIC)	66.3	620	1300
NVIDIA GTX 460 (GPU)	0.127	340	200
Intel Corei-5 (CPU)	0.014	95	82

- Bitcoin fee is near zero when the capacity is greater than the demand ($n_F > N$).
- Bitcoin fee revenues are maximized at the capacity $n_F = \frac{N}{2}$.
- Bitcoin fee increases with the proportional rate (ρ) levied on the off-chain option.

Miners join the mining network to earn these transaction fee revenues. Miner m has access to hardware with hash power $h(m)$. Faster hardware is also more costly to purchase and run. Hardware $h(m)$ costs $c(m)$ (USD/period). This periodic cost $c(m)$ is a combination of running costs (e.g., electricity) and the amortized cost of the hardware purchase. Table 1 provides a small comparison of computing hardware.

The best hardware (ASICs) costs up to 10 times more than a CPU but can deliver thousands of times more hash power. The hash power delivered is disproportionately larger than the higher cost.

If this were not true, simply purchasing multiple CPUs would be preferable to buying a single ASIC.

$$m_1 < m_2 \implies h(m_1) > h(m_2); \quad c(m_1) > c(m_2); \quad \frac{h(m_1)}{c(m_1)} > \frac{h(m_2)}{c(m_2)} \quad (7)$$

The running cost component does not change any of our analysis. Without loss of generality, we set this cost to zero going forward. We focus on the upfront cost of hardware purchases, which is amortized periodically at c (USD/period). This is equivalent to an upfront fixed cost of $\frac{c}{1-\delta}$. This upfront fixed cost plays a significant role in miners' decisions to participate in Bitcoin mining.

As more miners join the network, the revenue pie for each individual miner shrinks. The marginal miner is the one who earns just enough fee revenue to equal his/her cost. Let m_0^* represent the rank of marginal miners and H_0^* represent the total mining power in the network. The revenue earned by the marginal miner is given by the total block fee $f_0^* n_F$ multiplied by the miner's probability of creating any given block. This probability is a ratio of their own computing power h_j relative to the networks overall power. Networks overall power is the sum of power for all miners between 0 and m_0^* i.e., $\int_0^{m_0^*} h(m) dm = H(m_0^*)$. All miners that possess more computing power than the marginal miner stand to earn greater revenues. These miners will all decide to buy the hardware and mine every period. All miners that possess inferior computing power would not be able to earn sufficient revenues to make up for the hardware costs. These miners prefer to refrain from mining. Any miner that decides to buy their hardware at the start prefers to continue mining every period. We express *Do not mine* as a possible choice, which is always worse off than mining once the block creation subgame is reached. This implicitly places an exit cost on miners. In the current simple setting, no miners want to exit. We highlight the importance of this exit cost in later sections.

We have a unique solution to identify the marginal miner m_0^* as long as $f_0^* n_F > c(0)$. The uniqueness is guaranteed since $\Lambda(m_0) = \frac{h(m)}{H(m)c(m)}$ is monotonically decreasing from 1 to 0 on its domain $m \in [0, \infty)$ ⁸.

$$(B + f_0^* n_F) \times \frac{h(m_0^*)}{H(m_0^*)} = c(m_0^*); \quad \text{where} \quad H(m_0^*) = \int_0^{m_0^*} h(m) dm \quad (8)$$

This simplifies to,

$$m_0^* = \Lambda^{-1} \left(\frac{1}{B + f_0^* n_F} \right) \quad (9)$$

The equilibrium number of miners that enter m_0^* is an increasing function of the block reward B and fee revenue $f_0^* n_F$. This block reward is decreased by half every four years and is expected to

⁸ $\frac{h(m)}{c(m)}$ is monotonically decreasing, as discussed above. Additionally, $H(m)$ is a cumulative sum thus increasing in m .

Table 2 The main notation for our analysis

Notation	Description
v	Value of a transaction.
V_{max}	Maximum value of a transaction.
f	Transaction fee offered by a user.
ρ	Proportional transaction fee charged by a bank.
δ	Discount factor delay in verification of a transaction.
c	Per period amortized cost of mining.
n_F	Fixed capacity (block size) of Bitcoin, decided by its designer.
n_P	Size of a partially filled block, decided by a colluding miner.
N	Total number of transactions needing verification (per unit block time).
γ	Ratio of Bitcoin capacity to its demand.
α	Proportion of computing power relative to the total mining network.
α_l	Minimum power for a single miner to profit from partial blocks.
$\hat{\alpha}$	Minimum power for an individual miner in a colluding group.
α_m	Power of smallest colluding miner.
α_s	Power of smallest free riding miner.

vanish eventually. Thus, going forward, we set this block reward to zero. In current state a non zero block reward does not change our findings but simply reduces a miners sensitivity to the fee revenue component. This block fee revenue is zero at $n_F = 0$. Additionally, from earlier, the fees f_0^* are zero at $n_F = N$. At these extremes with block reward at zero ($B = 0$), $m_0^* = \Lambda^{-1}(\infty)$ collapses to zero. No miners are willing to join if there are zero revenues. For a given non-zero block revenue ($f_0^* n_F$), the marginal miner block winning probability is $c(m)/f_0^* n_F$.

Proposition 2 *The number of miners and corresponding network computing power increase with increasing mining revenues. The mining power collapses at zero capacity ($n_F = 0$) and full capacity ($n_F = N$).*

- *Number of miners and network computing power decrease with high hardware running cost c (USD/block period).*

3.1. Strategic Mining

We have shown that miner revenue is maximized when $n_F = \frac{N}{2}$. However, this implies that half the users seeking transactions are turned away. This is why the Bitcoin community is considering increasing the block capacity. In this section, we investigate whether increasing the block capacity over $\frac{N}{2}$ will help scale Bitcoin. We consider the case whereby the Bitcoin designer has set Bitcoin's

capacity above the fee-revenue-maximizing value, i.e., $n_F \geq \frac{N}{2}$. We investigate if an individual strategic miner, with α proportion of the total network hash power, sees a profitable deviation by starting to create partially filled blocks. These partial blocks include only the top $n_P (< n_F)$ transactions. Such a strategy could potentially create an artificial capacity constraint. A smaller capacity means that individual users need to beat greater competition to have their transaction included. It is initially unclear if the increased fee competition among users will make up for fewer transactions being included on the block. We will show that a strategic miner with at least α_l (defined in (3)) proportion of the total computing power is able to extract excess revenues via partial block filling. A miner smaller than α_l creates blocks too infrequently and, therefore, is not able to create sufficient congestion to make up for their sacrifice of block capacity.




Figure 10 The users with the highest values ($v \geq v_h$), unwilling to risk being left out of the immediate block, will pay higher fees f_h . A middle tier of users ($v \in [v_l, v_h]$) will prefer to take the risk rather than pay a higher fee. The lowest tier ($v \in [0, v_l]$) will continue to be left out of the chain as before.

All users who attempt to transact on chain were successful without exclusion in our base model, where the partial filling of blocks was not allowed. Users had rational expectations of deterministic miner actions. They had to beat exactly $N - n_F$ other users to be included. However, users now face uncertainty in whether the next block will be mined by a strategic or passive miner. They have three choices: (1) offer zero fee and take the off-chain option, (2) offer a low fee that beats only $N - n_F$ other users hoping for a full block (n_F), or (3) offer a high fee that beats $N - n_P$ other users to be included irrespective of a full (n_F) or partial (n_P) block. We focus on a pure strategy equilibrium whereby the user’s choice of fee ($\{0, f_l, f_h\}$) tiers is an increasing function of value. The users with the highest values stand to lose the most $v(1 - \rho)(1 - \delta)$ if excluded. High-value users ($v \geq v_h$) unwilling to take this risk pay higher fees f_h . They are included in the immediate block, deriving a payoff $\delta(v - f_h)$. A middle tier of users ($v \in [v_l, v_h]$) will prefer to take the risk rather than pay a higher fee. A lower fee offer of f_l derives a payoff of $\delta(v - f_l)$ with probability $(1 - \alpha)$ or a fall back off-chain payoff of $\delta v(1 - \rho)$ with probability α . The lowest tier ($v \in [0, v_l]$) will continue to be left out of the chain as before. The three segments of users are illustrated in Figure 10.

The transaction tiers depend on block sizes (n_F or n_P) not on α . But indirectly, α changes how aggressively are users willing to compete to be part of the top tier. Being in the middle tier, exposes a user to a risk. If α is high at 0.5, then every 2nd block is partially filled. A middle tier user has a 50% chance that their transaction will be ignored. If alpha was lower at 0.05, there is only a 5% chance that their transaction is ignored on the next block. Thus a greater alpha leads to higher risk that the user will be ignored and would eventually need to pay proportional fee on the outside option. A higher risk in turn creates more aggressive competition to move from the middle tier to top tier.

We now proceed to obtain the equilibrium expressions of v_l , v_h , f_l , f_h , and n_P . We know that the value of v_h (resp., v_l) should be such that there are exactly n_P (resp., n_F) transactions that have transaction values above v_h (resp., v_l). Thus, we have

$$v_h = V_{max} \left(1 - \frac{n_P}{N}\right), \quad (10)$$

$$v_l = V_{max} \left(1 - \frac{n_F}{N}\right). \quad (11)$$

The user at $v = v_l$ is indifferent between using Bitcoin and banks. Thus, we have

$$(1 - \alpha)\delta(v_l - f_l) + \alpha\delta v_l(1 - \rho) = v_l(1 - \rho).$$

From the above equation, we obtain

$$f_l = \frac{v_l(\delta + \rho - 1 - \alpha\delta\rho)}{(1 - \alpha)\delta}. \quad (12)$$

Similarly, the user at $v = v_h$ is indifferent between taking a risk and offering a fee of only f_l or not taking the risk and offering a fee of f_h . Thus, we have

$$\delta(v_h - f_h) = \delta(1 - \alpha)(v_h - f_l) + \delta\alpha v_h(1 - \rho).$$

From the above equation, we have

$$f_h = \alpha\rho v_h + (1 - \alpha)f_l. \quad (13)$$

Substituting the value of v_l in (12), we obtain

$$f_l = V_{max}\rho\beta\alpha(1 - \alpha). \quad (14)$$

Let $\pi(n_P)$ represent the profit of the strategic miner. Then, we have

$$\pi(n_P) = f_h \times n_P = \left[\alpha\rho V_{max} \left(1 - \frac{n_P}{N}\right) + (1 - \alpha)f_l \right] n_P.$$

The strategic miner chooses the optimal value of n_P by maximizing $\pi(n_P)$. Let n_P^* represent this optimal value. Clearly, $\pi(n_P)$ is a concave quadratic function of n_P . Thus, using the first-order condition, we obtain

$$n_P^* = \frac{N}{2}(1 + \beta), \quad (15)$$

where β is defined in (3). Substituting the value of n_P^* in (10), we obtain

$$v_h = \frac{V_{max}(1 - \beta)}{2}. \quad (16)$$

Using (16), (14), and (13), we have

$$f_h = \frac{V_{max}\alpha\rho(1 + \beta)}{2}. \quad (17)$$

Finally, we now have the equilibrium expressions of v_l , v_h , f_l , f_h and n_P given by equations (11), (16), (14), (17) and (15), respectively.

We know that the strategic miner cannot choose a value of n_P higher than n_F , i.e., $n_P^* \leq n_F$. Using the expression of n_P^* from (15), this condition simplifies to

$$\alpha \geq \alpha_l. \quad (18)$$

This condition implies that for a miner with low hash power ($\alpha < \alpha_l$), the partial block filling strategy is always dominated by a full block filling strategy. This miner adds blocks too infrequently to seriously threaten users to increase their transaction fee offers and benefit from them.

Proposition 3 *A miner with at least α_l proportion of the total mining network's computing power is able to conduct a partial block filling strategy and extract undue revenues from Bitcoin transaction fees.*

$$\alpha_l = \alpha_h \frac{1 - \gamma}{\gamma} \quad (19)$$

- This strategic miner adds n_P^* transactions offering f_h^* fees each. All other blocks add n_F^* transactions offering f_l^* fees each.
- The revenue-maximizing choice of partial filling size, i.e., n_P^* , decreases with miner power α .

$$n_P^* = N(1 + \beta)/2; \quad f_h^* = V_{max}\rho\alpha(1 + \beta)/2; \quad f_l^* = V_{max}\rho\beta\alpha(1 - \alpha) \quad (20)$$

Remark: From the definition of β in (3), we know that $\beta = 0$ at $\alpha = \alpha_h$. Thus, in (15), at $\alpha = \alpha_h$, the value of β is 0, and therefore, the value of $n_P^* = \frac{N}{2}$. However, from Proposition 1, we also know

that the Bitcoin system generates the maximum revenue when the capacity is $\frac{N}{2}$. Therefore, it is not in the interest of a strategic miner to reduce the capacity of the Bitcoin network below $\frac{N}{2}$. Thus, $n_P^*(\alpha) = \frac{N}{2}$, $\forall \alpha \geq \alpha_h$. An extremely powerful strategic miner ($\alpha \geq \alpha_h$) can force users to either offer a high fee or not take a risk with the on-chain option at all. A miner who wants to add a full block n_F does not find any awaiting transactions beyond n_P paying a lower fee f_l . In other words, this miner can take complete control of Bitcoin and act like a designer while keeping the effective capacity at its optimal value, i.e., $n_P = \frac{N}{2}$. This partial filling strategy is irrelevant if the capacity on the chain is already constrained (say, $n_F \leq \frac{N}{2}$).

Consider a hypothetical situation whereby demand is 25% over capacity ($n_F = \frac{4N}{5}; \gamma = \frac{1}{5}$). A miner with power greater than 90.9% ($\alpha_h = \frac{\delta + \rho - 1}{\delta \rho} = 0.909$ for $\delta = 0.99; \rho = 10\%$) shrinks the artificial capacity to the revenue-maximizing level $n_P = \frac{N}{2}$ by filling 62.5% (n_P/n_F) of their block. A miner with as little as 22.7% of the total computing power ($\alpha_l = \alpha_h \frac{1-\gamma}{\gamma}$) of the network fills 78% ($n_P/n_F = \frac{1+\beta}{2\gamma}$) of their block and extracts undue revenues. A miner with less than 22.7% of the total computing power is unable to unilaterally engage in a partial block filling strategy.

Assume that the blockchain designer attempts to increase the block capacity to serve greater demand, say, $n_F = 0.9N$. The same miner ($\alpha = 0.227$) performs partial filling. These partial blocks are only 59% full in the second setting compared to 78% in the first setting. Furthermore, at $n_F = 0.9N$, a miner with as little as 9.09% of the total network power strategically performs under filling. They fill 60% ($n_P/n_F = \frac{1+\beta}{2\gamma}$) of their block. Therefore, an increase in block capacity allows an increasingly smaller miner to perform strategic under filling.

11 shows three scenarios for a large miner with power $\alpha = \alpha_h, \alpha_h/2, \alpha_h/4$. The largest miner is able to keep the effective block capacity at $N/2$. The effective capacity is given by $n_P * \alpha + n_F * (1 - \alpha)$ when the individual miners have less than α_h power. Smaller miners are not as effective but nevertheless stop the blockchain from serving the entire demand. While smaller miners have less impact on block capacity, such miners (say, 5% of the total computing power) are very likely to be present on the Bitcoin network. In practice, the accumulation of a high amount of computing power ($\alpha > 0.5$) may not be realistic for a single miner. This can be accomplished by a mining pool (e.g., Antpool). Although a mining pool allows for some coordination of mining resources, they do not have complete control. Individual participants can hop on and hop off at any time. In the next section, we discuss the possibility for a group of small miners to collude together credibly to achieve such computation power as a group. These groups need not be mining pools.

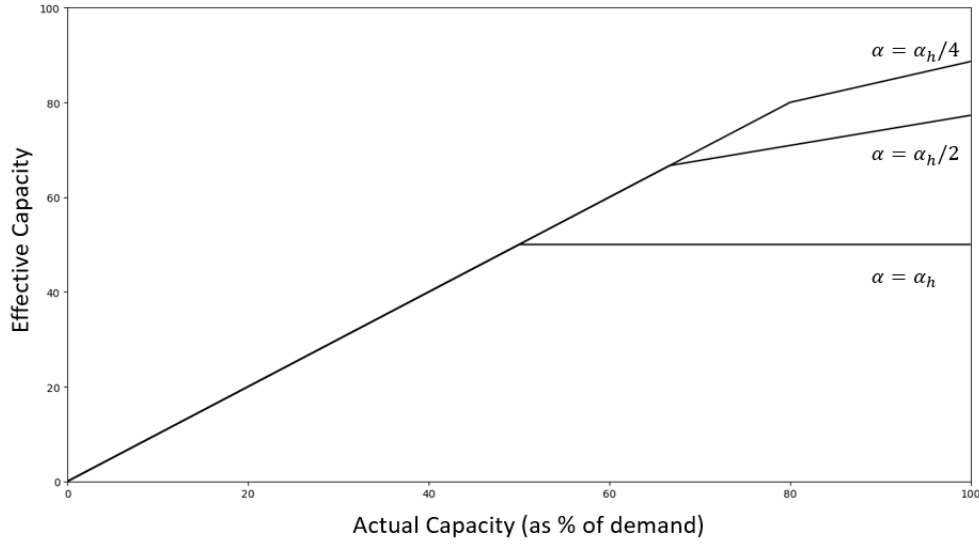


Figure 11 A very larger miner ($\alpha = \alpha_h = 0.909$) does not let the effective capacity increase beyond 50% of demand. A small miner ($\alpha = \alpha_h/4 = 0.227$) performs partial filling when the actual capacity is raised beyond 80%. They do not let the effective capacity increase beyond 86% of demand. ($\delta = 0.99$, $\rho = 10\%$)

3.2. Collusion

Partial block filling by a large strategic miner (with at least α_l fraction of the total computing power) achieves extra mining revenues. A group of small miners comprising the same computation power ($\sum \alpha_j = \alpha_l$) should similarly be able to achieve the same extra revenues. This requires miners to agree on the mutually beneficial partial block filling strategy. Such coordination is complicated by the low barriers to entry of new miners. In this section, we identify the conditions under which a tacit collusion is stable.

There is one key difference between the incentives of a single large miner and a group. The large strategic miner sacrifices non-zero fees. This creates artificial congestion and increased fee offers. The subsequent benefit of higher fee offers via this congestion is shared by the large miner as well as all other miners. The remaining miners ($1 - \alpha_l$) thus free ride on congestion created by the large miner. Full blocks added by free riders gather larger fee revenues (R_F) than partial blocks (R_P) added by the large miner (equation 21). A group of small miners is not incentivized to make the same sacrifice. Every individual miner prefers that others perform the partial block filling while they free ride on the resulting congestion. The mutually beneficial Nash equilibrium is therefore unstable in a single block creation game, i.e., all miners want a free ride.

$$R_P = f_h n_P^*, \quad (21)$$

$$R_F = f_h n_P^* + f_l (n_F - n_P^*). \quad (22)$$

Bitcoin's block creation is repeated infinitely among the same set of miners. The infinite repetitions allow miners to sustain a tacit collusion. They can mutually commit to partial block filling under the threat of future punishment. A single deviation by an individual miner would result in no future collusion.⁹ In a single-stage game, an individual miner has a profitable deviation. However, the deviation is no longer profitable when accounting for all future payoffs. We therefore focus on Subgame Perfect Equilibrium (SPE) miner strategies in an infinitely repeated block creation game.

The verifiability of a miner's action is a key requirement for tacit collusion. Colluding miners should be able to verify each other's partial block filling actions. New blocks on the Bitcoin network – full or partial – are public. Colluding miners can perform their actions using a unique public address to add their blocks. An alternate method is for colluding miners to pool their computational power. For example, Antpool is a group of miners who have decided to pool their resources together for mining. Each miner participating in a pool contributes a verifiable amount of computational power and submits blocks under the pool ID. The pool then distributes the rewards from winning blocks based on the work performed. Regardless of whether the miners gather together in a pool or perform their actions using a unique public address, they can still deviate. However, such a deviation is easily detectable. If the colluding miners use the same public address, transactions included in any block that they win will reveal if they are deviating. However, it is possible for the miner to use a different public address when cheating. In a pool, this would not be possible because the pool can verify the computational power contributed by the miner to the pool and also the transactions included by the miner in a block. Outside of a pool, a drop in the win rate of a large miner (identified by the public address) may indicate that he is cheating using another address to submit fully filled blocks. Further, a colluding miner does not need to verify every other miner's actions. Overall, if α_l fraction of the total computational power is assumed to be involved in collusion, a participating miner simply needs to check that on average α_l fraction of blocks are partially filled to ensure that no one is deviating.

Every miner prepares a block of transactions and then starts to look for the solution to the Bitcoin mining puzzle. Miners need to decide their filling action – full or partial – before starting to find the mining puzzle solution. A colluding miner considers a trade-off (equation 23) between (a) immediate profits from a partially filled block, i.e., $\pi_P := R_P - c$, followed by expected lifetime profits from collusion or (b) deviation to completely fill their block, i.e., $\pi_F := R_F - c$, followed by a lifetime of no collusion profits $\pi_0 := R_0 - c$. Recall that δ represents the time discounting factor for Bitcoin users,

⁹ Instead of the grim strategy, we could also study a punishment strategy whereby the miners terminate the collusion strategy for a fixed period (instead of perpetuity) that is sufficient to wipe out the profits of the deviating miner.

which can potentially be different from that for miners. Let δ_m represent the time discount factor for miners. A colluding miner of size α_j will not deviate from collusion if

$$\pi_P + \alpha_j \pi_P * \frac{\delta_m}{1 - \delta_m} \geq \pi_F + \alpha_j \pi_0 * \frac{\delta_m}{1 - \delta_m}, \tag{23}$$

which simplifies to

$$\alpha_j \geq \hat{\alpha}, \quad \text{where} \quad \hat{\alpha} = \frac{1 - \delta_m}{\delta_m} \times \frac{R_F - R_P}{R_P - R_0}. \tag{24}$$

This creates a different tradeoff for heterogeneous miners. A small miner ($\alpha_j < \hat{\alpha}$) mines blocks infrequently, say, once a month or year. They are less threatened by punishment far off in the future. This can also be seen as a small miner’s desire to make the most out of winning a mining puzzle once in a long while. A large miner sticks to the collusion strategy, as they expect frequent or near-term fee revenues.




Figure 12 The x-axis represents miners ranked by access to hardware. The y-axis represents the computing power of the respective miner. An example of a monotonically decreasing concave function $h(m) = \lambda e^{-\lambda m}$ with a long tail. Three categories of miners, from left to right, - (1) Colluding by partial block filling, (2) Free Riding by full block filling, and (3) Inactive miners. m^* and s^* represent the boundary between these groups in a collusion equilibrium.

The constraint above ($\alpha_j \geq \hat{\alpha}$) assures a colluding miner’s commitment to the collusion. Extra mining revenues from collusion invite new miner entry. Small miners – who would otherwise find it too costly to mine – join the network to free ride. This endogenous miner entry makes collusion non-trivial. Figure 12 illustrates our model of miner heterogeneity with respect to access to computing hardware. A large miner has access to hardware capable of faster hash calculations at the same cost. Given the favorable trade-off for the large miners, we focus on collusion among a group ($\sum \alpha_j = \alpha_l$) of the largest miners. We must ensure that no miner (colluding, free riding, or outside) deviates from their action in equilibrium.

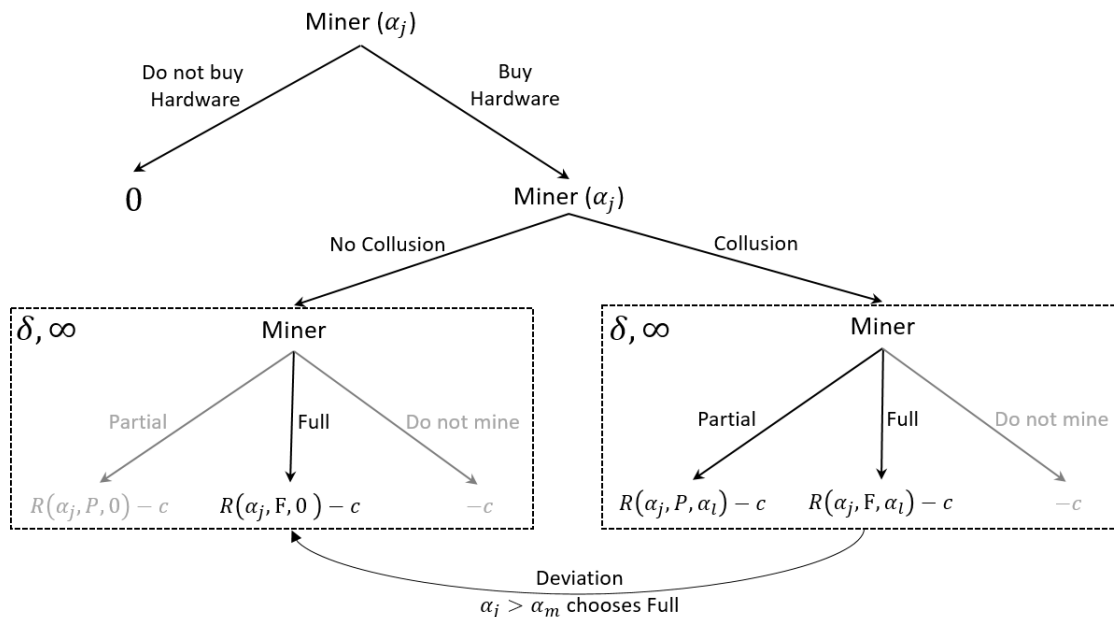


Figure 13 Miner j with access to computing power (α_j) decides whether to buy hardware. Next, the miner decides whether to follow the collusion strategy. The colluding group of miners chooses Partial fill. The free-riding group chooses Full. Deviation by a colluding group member leads to a no-collusion setting whereby everyone always fully fills their block. The block creation sub-games on equilibrium (right) and off equilibrium (left) are repeated infinitely with a discount factor δ_m .

Figure 13 depicts the sequence of choices for a miner. First, they decide whether to purchase the hardware. Second, they decide whether to follow the collusion strategy. The second choice is repeated over infinite block creation periods. These choices are made by all miners simultaneously. We want to identify conditions under which the top m^* miners collude and the next $(s^* - m^*)$ miners participate as free riders. In a subgame equilibrium, these participating miners should be willing to purchase the hardware and stick to the collusion strategy. All miners beyond s^* should be better off not purchasing the hardware. An individual (focal) miner only observes the actions of other miners once they participate in block creation. However, the individual miners have rational expectations with respect to equilibrium strategies followed by all other miners.

The focal miner first decides whether to buy the hardware. Under collusion equilibrium, a rational miner expects the top s^* miners to buy hardware. Next, they decide whether to follow the equilibrium collusion strategy or the no-collusion strategy. On the collusion path, a rational miner expects the top m^* miners to add partial blocks. On the no-collusion (off-equilibrium) path, a rational miner expects all active miners to add full blocks. On both the collusion equilibrium and the no-collusion off-equilibrium path, the focal miner exercises a choice of partial filling, full filling or no mining. The block creation sub-game on-equilibrium (right) and off-equilibrium (left) paths

Table 3 Single-period expected payoffs corresponding to three actions (Partial, Full, and No Mining) off and on collusion equilibrium paths.

	On Equilibrium (α_l)	Off Equilibrium ($\alpha_l = 0$)
Partial $R(\alpha_j, P, *)$	$f_h n_P$	$f_0 n_P$
Full $R(\alpha_j, F, *)$	$f_h n_P + f_l (n_F - n_P)$	$f_0 n_F$
No Mining $R(\alpha_j, 0, *)$	0	0

are repeated infinitely with a discount factor δ_m .

Single-block-fee payoffs are denoted by $R(*, *, *)$ with three arguments. The first argument represents the mining power of the focal miner. The second argument represents the focal miner block fill action, i.e., partial (P) or full (F). The third argument represents the colluding group power, i.e., collusion (α_l) or no collusion (0). Table 3 provides the single-period expected payoffs corresponding to all actions. Single-period payoffs are strictly better under the full block filling action for all miners. If the focal miner has power $\alpha_j \geq \alpha_m$ but decides to add a full block under collusion, they expect to be punished. All miners would move to the no-collusion sub-game if a single miner deviates from the collusion.

Table 4 lists all constraints that ensure that no miner has a profitable deviation in an SPE. For a large focal miner ($\alpha_j > \alpha_m$), constraint 1a ensures that they prefer to add partial blocks rather than a full block in the repeating block creation sub-game. This is fulfilled if the marginal miner satisfies $\alpha_m \geq \hat{\alpha}$. Constraint 1b represents their preference to buy the hardware at the start of the game. This is satisfied for the marginal miner making zero profits.

$$\alpha_m = \frac{c_m(1-\delta)}{R_P} \geq \hat{\alpha}; \quad \text{where} \quad R_P = f_h n_P \quad (25)$$

A miner with ($\alpha_m \geq \alpha_j \geq \alpha_s$) proportion of the total power free rides. The lower limit α_s denotes the smallest miner that joins the mining network. Constraint 2a represents the preference to free ride over joining the colluding group. Joining the colluding group would increase the power of the colluding group to $\alpha_l + \alpha_j$ and therefore the partial block revenues. If the marginal miner ($\alpha_j = \alpha_m$) is large enough, they may increase the partial block revenues $R_P(\alpha_l + \alpha_j)$ to be higher than the full block revenue $R_F(\alpha_l)$. In this section, we are interested in settings whereby even the largest miner is too small to perform partial block filling without a threat of punishment¹⁰. Large miners unilaterally perform partial block filling as shown in Section 3.1.

¹⁰ An equilibrium at α_l is only valid when individual miners are relatively small: $\alpha_m \leq R_P^{-1}(R_F(\alpha_l) - \alpha_l)$

Table 4 List of all constraints that ensure that no miner has a profitable deviation in an SPE. Three pairs of constraints (1a,1b), (2a,2b) and (3a,3b) correspond to three types of miners - (1) colluding miners, (2) free-riding miners and (3) inactive miners, respectively.

Focal Miner	Constraint
	1a $R(\alpha_j, P, \alpha_l) + \alpha_j \frac{\delta_m}{1-\delta_m} R(\alpha_j, P, \alpha_l) \geq R(\alpha_j, F, \alpha_l) + \alpha_j \frac{\delta_m}{1-\delta_m} R(\alpha_j, F, 0)$
$\alpha_j > \alpha_m$	1b $\frac{1}{1-\delta_m} R(\alpha_j, P, \alpha_l) - c_j \geq 0$
	2a $R(\alpha_j, F, \alpha_l) + \alpha_j \frac{\delta_m}{1-\delta_m} R(\alpha_j, F, \alpha_l) \geq R(\alpha_j, P, \alpha_l + \alpha_j) + \alpha_j \frac{\delta_m}{1-\delta_m} R(\alpha_j, P, \alpha_l + \alpha_j)$
$\alpha_m \geq \alpha_j \geq \alpha_s$	2b $\frac{1}{1-\delta_m} R(\alpha_j, F, \alpha_l) - c_j \geq 0$
	3a $\alpha_j \frac{1}{1-\delta_m} R(\alpha_j, F, 0) - c_j \leq 0$
$\alpha_j \leq \alpha_s$	3b $\frac{1}{1-\delta_m} R(\alpha_j, P, \alpha_l + \alpha_j) - c_j \leq 0$

Constraint 2b represents free-riding miners' preference to buy the hardware at the start of the game. This is satisfied for the smallest miner making positive profits.

$$\alpha_s \geq \frac{c_s(1-\delta)}{R_F}; \quad \text{where} \quad R_F = f_h n_P + f_l (n_F - n_P) \quad (26)$$

For a focal miner who stays out ($\alpha_j \leq \alpha_s$), constraint 3a represents their preference to not join as a free rider. Joining in as a free rider reduces the power of the colluding group below α_l and makes collusion unprofitable. Since the smallest miner makes zero profits when free riding, they are guaranteed to make negative profits when collusion breaks. Finally, constraint 3b represents their preference to join the colluding group. Similar to the free rider, this increases the power of the colluding group to $\alpha_l + \alpha_j$ and therefore the partial block revenues. These partial block revenues $R(\alpha_l + \alpha_j, P, \alpha_j)$ must be smaller than zero. This is automatically satisfied since $\alpha_s < \alpha_m$.

We now proceed to obtain the equilibrium expressions for the smallest colluding miner, denoted by m^* , and the smallest free-riding miner, denoted by s^* . For the smallest colluding miner, we need constraint 1b to become an equality. Specifically, we need

$$\frac{1}{1-\delta} R(\alpha_m, P, \alpha_l) - c_m = 0.$$

We know that $R(\alpha_m, P, \alpha_l) = \alpha_m \times R_P$. Thus, from the above equation, we have

$$\alpha_m \times R_P = c_m(1-\delta)$$

or equivalently

$$\alpha_m = \frac{c_m(1-\delta)}{R_P}. \quad (27)$$

From the definition of α_l , we know that

$$\alpha_l = \frac{H(m^*)}{H(s^*)}.$$

Thus, we have

$$H(s^*) = \frac{H(m^*)}{\alpha_l}. \quad (28)$$

From the definition of α_j , we also know that

$$\alpha_m = \frac{h(m^*)}{H(s^*)}.$$

Substituting the expression of $H(s^*)$ from (28), we obtain

$$\alpha_m = \frac{h(m^*)\alpha_l}{H(m^*)} = \Lambda(m^*)c(m^*)\alpha_l,$$

where $\Lambda(m) \equiv \frac{h(m)}{H(m)c(m)}$. Using (27), we have

$$\Lambda(m^*)c(m^*)\alpha_l = \frac{c(m^*)(1-\delta)}{R_P},$$

or

$$m^* = \Lambda^{-1} \left(\frac{1-\delta}{\alpha_l R_P} \right). \quad (29)$$

The monotonically decreasing function $\Lambda \equiv \frac{h(m)}{H(m)c(m)}$ ensures unique solutions. The marginal colluding miner earns zero profit and must be large enough such that future punishment is a credible threat.

$$\alpha_{m^*} \geq \hat{\alpha}; \quad \frac{c(m^*)(1-\delta_m)}{R_P} \geq \frac{1-\delta_m}{\delta_m} \times \frac{R_F - R_P}{R_P - R_0} \quad (30)$$

Free-riding miners present in equilibrium enter until colluding miners have exactly α_l proportion of the total computing power. Using (28), we have

$$s^* = H^{-1} \left(\frac{H(m^*)}{\alpha_l} \right), \quad (31)$$

where m^* is given in (29). In addition, the smallest free-riding miner must be large enough to make positive profits.

$$\alpha_{s^*} = \frac{h(s^*)}{H(s^*)} \geq \frac{c(s^*)(1-\delta_m)}{R_F} \quad (32)$$




Figure 14 A very large group ($\alpha = \alpha_h$) does not let the effective capacity increase beyond 50% of demand. They do not need to satisfy any $\hat{\alpha}$ constraint. A small group ($\alpha = \alpha_h/4$) does not let the effective capacity increase beyond 86% of demand. They only perform partial filling when the actual capacity is raised beyond 90%. The group is slightly less effective than a single miner because they also need to satisfy the $\hat{\alpha}$ constraint. ($\delta = 0.99$, $\rho = 10\%$)

Proposition 4 A group of miners with α_l proportion of the total mining power can sustain collusion if the mining power of the smallest colluding miner, i.e., α_m , is greater than $\hat{\alpha}$.

$$\frac{c(m^*)(1 - \delta_m)}{R_P} \geq \frac{1 - \delta_m}{\delta_m} \times \frac{R_F - R_P}{R_P - R_0} \quad (33)$$

- Under this collusion, all miners with greater than α_m proportion of the total mining power collude.
- The total number of miners s^* exceeds the total number of miners who joined under the no-collusion strategy m_0^* .¹¹

The subgame perfect equilibrium above focuses on a colluding group with exactly α_l power. All miners adding partially filled blocks ($\sum \alpha_j = 1$) is yet another SPE. In such a case, users either stay off the chain or offer a high fee (f_h); no one offers a low fee (f_l). As a result, a deviation in filling the block with low-fee-paying transactions is not an option. We have not observed such full collusion on the Bitcoin network. We do not provide a specific justification for one equilibrium over other equilibria; however, we focus on the α_l collusion as a more interesting and practical setting.

¹¹ Since $R_P > R_0$ and Λ^{-1} is monotonically decreasing, $s^* > m^* > m_0^*$.

In addition to these extreme cases, collusive equilibria with $\alpha_l \leq \Sigma\alpha_j \leq \alpha_h$ may also be possible. If a single miner with power $\alpha_{j'}$ deviates from such collusion, the remaining group is left with power $\Sigma\alpha_j - \alpha_{j'}$. Punishing the deviating miners requires this group to not engage in collusion permanently. This is not necessarily a rational strategy for the remaining $\Sigma\alpha_j - \alpha_{j'}$ group at this stage. They are better off colluding on partial block filling if $\Sigma\alpha_j - \alpha_{j'} > \alpha_l$. This rational strategy to collude with a smaller group does not constitute a threat to the deviating miner. The deviating miner is thus better off by continuing to free ride. In this paper, we do not validate or reject the existence of miner strategies that sustain such equilibria.

4. Intervention and Security

The collusion equilibria discussed in the previous section results in a smaller effective block capacity, greater fees and larger revenues for the mining network. The Bitcoin community wants to eliminate this collusion to serve the entire user demand at low fees. A simple intervention could be to enforce block filling. Such an intervention is easily defeated, as the colluding miners could simply fill their remaining blocks with artificial transactions. They could add zero-fee-paying transactions sending coins back and forth between their own accounts. Such artificial transactions would achieve the same objective as partial block filling by keeping out actual user demand. A more realistic solution could be to intervene in the tacit miner collusion itself. The Bitcoin community could attempt to block large miners to interrupt such collusion.

The existence of the hardware of the large miners ($> \hat{\alpha}$) helps sustain the collusion. Bitcoin's mining puzzle requires miners to perform SHA256 hash calculations. These calculations are computationally intensive, i.e., requiring high calculation speeds but minimal memory. Large miners deploy ASIC machines that are orders of magnitude faster than GPUs and CPUs at SHA256 hash calculations. A few blockchains – alternatives to Bitcoin – implement ASIC-resistant mining puzzles. ASIC-resistant mining puzzles are memory intensive, i.e., requiring substantial storage memory. ASIC machines deployed by large miners are efficient at computations but inefficient at memory-intensive operations. As ASICs become un-competitive, miners with GPUs and CPUs are able to enter the mining network. Since GPUs and CPUs are more widely accessible than ASICs, this change has the potential of making miners more equitable. In the context of our hardware distribution model, this effectively means that mining is now performed by large numbers of miners all in the tail of the distribution.

The hardware distribution functions $h(m), c(m)$ and $\Lambda(m)$ are now updated to $h'(m), c'(m)$ and $\Lambda'(m)$, respectively. A colluding group could still form, thereby adding up to the same (α_l) proportion




Figure 15 The x-axis represents miners ranked by access to hardware. The ASICs on the left are blocked by an intervention. A colluding group with α_l power would need to include a larger number of small miners. The new marginal miner (m^*) may not be willing to sacrifice block capacity under tacit collusion.

of the total computing power. The proportion of the total computing power for the marginal miner remains given by (30). The marginal miner must still remain greater than $\hat{\alpha}$. This is satisfied if

$$c'(m^*) \geq \frac{R_P}{\delta_m} \times \frac{R_F - R_P}{R_P - R_0} \quad (34)$$

The updated cost function $c'(m)$, similar to $c(m)$, is a decreasing function of m .

$$c'(m) = \begin{cases} 0 & \text{if } m \text{ in banned region} \\ c(m) & \text{otherwise} \end{cases}$$

Prior to the ban, the marginal miner's cost was large enough to sustain collusion. As high-quality mining hardware is banned, the new marginal miner moves toward the right (see Figure 15). The cost of an infinitely small hardware (say, hand-held mobile devices) is near zero $c'(\infty) = 0$. There exists a unique level of hardware ban beyond which the inequality above is shattered. In summary, banning large miners invites smaller hardware that can be operated at lower cost. A lower cost attracts a more equitable community of miners – a larger number of miners each expending a smaller computing cost. The new marginal miner may be too small ($> \hat{\alpha}$). The marginal miner may not be willing to sacrifice block capacity under tacit collusion. At first glance, this appears to be an effective method to eliminate collusion and recover the full block capacity. We will show that intervention (ASIC resistance or otherwise) to eliminate collusion threatens the security of user payments by inviting hacks.

4.1. Double-Spend Attack

A double-spend attack poses a threat to the validity of a transaction. Note that the Bitcoin ledger is a unique longest chain of blocks. As discussed earlier, when multiple chains emerge, the chain that attracts the majority of the mining power emerges as the longest and thus the accepted chain. Once

the longest chain emerges, all other parallel chains are abandoned. The transactions that appear in these shorter parallel chains but not in the longest chain are deemed invalid. An adversary with computing power ($\alpha_j = \theta$) can create a parallel chain. In addition, if the adversary has majority power ($\theta > 0.5$), his chain will always be the longest chain. Such an adversary essentially controls the blockchain. They could prevent new transactions from gaining confirmation, add invalid transactions, or even reverse transactions that were completed to double-spend their coins. This is well known as a 51% attack. If any miner were to achieve 51% of the aggregate computing power of the Bitcoin network, Bitcoin would lose all trust and value. As a result, it is not in the interest of any miner to accrue that much power.

A smaller adversary ($\theta < 0.5$) cannot arbitrarily add invalid transactions onto their block. Their block would be rejected in favor of a valid block. However, they can spend their account balance twice by creating parallel chains (Figure 16). The attack involves an adversary and a merchant. The adversary orders some service/product from a merchant and pays the merchant in Bitcoin. The adversary wants to receive the service/product but does not want to pay. The merchant would not deliver the service/product until the adversary made the payment. In a double-spend attack, the adversary makes a payment to the merchant, and the merchant delivers the product/service; then, the adversary has his transaction invalidated and receives his Bitcoins back. The attack would be successful if the adversary has the transaction recorded in one of the blocks, and the merchant subsequently delivers the product/service. The adversary then somehow convinces honest miners to choose a different chain a little later. The original payment to the merchant ends up on the chain that will be abandoned in future. The merchant has delivered a valuable service/product for a payment that never really occurred.




Figure 16 Adversary (A) sends a payment v to the Merchant (B). The merchant delivers a valuable service once this transaction is added onto a Bitcoin block. The adversary then manages to resend the same value v to another account (A'). The second transaction is recorded on the eventual longest chain.

We describe a simple double-spend attack strategy employed by an adversary. The attack is launched when the adversary has just mined a new block. The adversary decides to keep this block

private. Honest miners are attempting to add a new block in parallel, which will eventually create a parallel chain. We denote the adversary chain by a and the honest chain by h . State 0 represents the launch of the attack when the height of a is 1 and that of h is 0. The adversary's goal is to mine a private chain a with more blocks than the honest chain h . If the adversary releases their private chain, all miners will accept this adversary chain as the longest unique chain. The shorter honest chain would be abandoned. This is a risky strategy, with power $\theta < 0.5$ the adversary is likely to be left behind. If the adversary becomes left behind, they will lose the block fees earned on their abandoned blocks. If the small probability that their chain wins is realized, they want to make the most out of it by duping a merchant.

The adversary adds a payment of value v to one of their own accounts. They do not broadcast this spend publicly; instead, they add this to their private block. The adversary also broadcasts a payment of value v . This payment is meant to purchase valuable goods or services from a merchant. As $h = 0$, the adversary's payment has not been picked up on the publicly visible chain yet. The adversary must keep mining privately since the merchant does not deliver the valuable service before they see the payment added to the longest publicly visible chain.

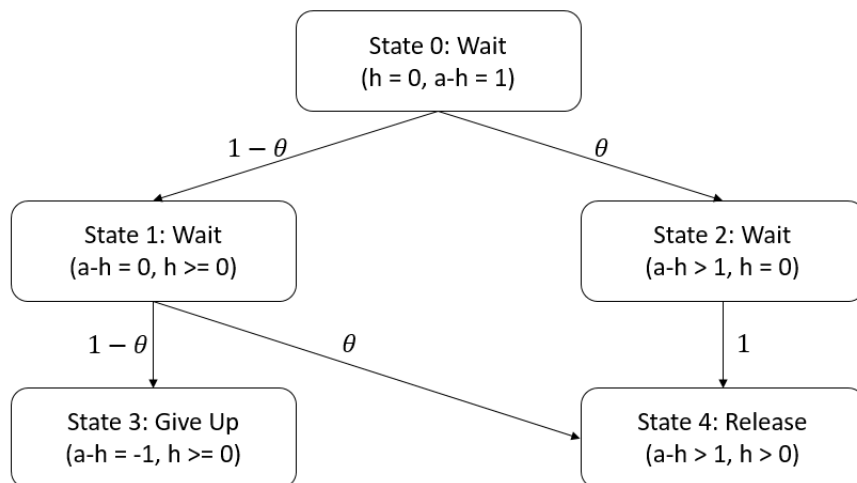


Figure 17 Adversary double-spend Markov decision process (MDP). Each node represents a state and the corresponding action for the adversary. State 0 is the initial attack launch state, and States 3 and 4 are terminal states. The edges represent the state transitions and the corresponding probabilities.

The next block is added to the honest chain with probability $1 - \theta$ landing to State 1. In this state, the honest and adversary heights are equal. If the adversary adds the next block (probability θ to State 4), they succeed. If the honest miners add the next block (probability $1 - \theta$ to State 3),

the adversary gives up, and the attack fails.¹² If the next block from the initial State 0 is added by the adversary, we land in State 2. The two heights are updated to $h = 0$ and $a = 2$. The attack is not complete since the broadcasted transaction has not been added to the honest chain yet. However, the adversary is certain to succeed in his attack at this point. He will release the longer chain once a single block is added to the honest chain. All miners will accept this adversary chain as the longest unique chain. The transactions on the adversary's blocks remain *valid*. When the ledger entries on the unique chain of blocks is added up, no balance seems to be double spent once the honest chain is ignored.

The adversary fails in this attack with a probability of $(1 - \theta)^2$. This corresponds to the probability of two consecutive blocks being added to the honest chain. As an example, an adversary with a 0.05 proportion of the total computing power has a success rate of $[1 - (1 - \theta)^2] = 9.75\%$. If the adversary succeeds, they obtain the double-spend value (v) in addition to the usual block fee. The adversary fails if honest miners add two blocks before the adversary is able to find a single block on top of their private chain. In this scenario, the adversary's private chain consists of a single block. This private block records a list of transactions that offer fees to the block creator (adversary). Unfortunately, this private chain is a short chain from the longest main honest chain. Therefore, none of these transactions are added up as the final balance of the adversary. The adversary forfeits the block revenue R on their private chain when they fail to perform their double-spend attack. A small adversary ($\theta \ll 0.5$) fails more often than not. The equation below captures the profits for the adversary when they have just launched the attack. They earn $(v + R)$ on the first block if it is eventually included in the main chain with a probability of $[1 - (1 - \theta)^2]$. Subsequently, they earn the usual honest mining rewards in the future.

$$\pi_{\text{double spend}} = [1 - (1 - \theta)^2](v + R) - c + \delta_m(\theta R - c) + \delta_m^2(\theta R - c) + \dots \quad (35)$$

As an alternative to this attack, an adversary could have simply engaged in the usual mining, i.e., not mining private blocks. The adversary would avoid losing the fee on their abandoned private block.

$$\pi_{\text{honest}} = (R - c) + \delta_m(\theta R - c) + \delta_m^2(\theta R - c) + \dots \quad (36)$$

This attack is worthwhile if the likely loss of block fees on the adversary's abandoned private chain is compensated by the double-spend value v . Equation (37) represents the additional payoff in launching a double-spend attack.

$$\pi_{\text{double spend}} - \pi_{\text{honest}} = \Delta\pi = v[1 - (1 - \theta)^2] - R(1 - \theta)^2, \quad (37)$$

¹² The adversary has other options – (1) release the private chain in state 1 and let the honest miners decide on a toss up between two equal-height chains or (2) continue mining privately at state 3 in hopes of once again obtaining their advantage. These alternatives are preferable in certain parameter ranges but do not change the core result.

To dissuade an adversary from performing a double-spend attack, we need $\Delta\pi \leq 0$. From (37), this condition simplifies to

$$v \leq \hat{v},$$

where

$$\hat{v} = \frac{R(1-\theta)^2}{1-(1-\theta)^2}. \quad (38)$$

The above condition shows that only the transactions with value $v \leq \hat{v}$ are safe from a double-spend attack. Thus, in equilibrium, users with transactions above \hat{v} will not use Bitcoin and rather will use traditional banks. From the definition of \hat{v} , it is clear that the value of \hat{v} decreases with R . In other words, the demand for Bitcoin decreases with the fee-revenue R . Next, we analyze if both collusion and double-spend attacks can be avoided by the Bitcoin community.

Assume that the Bitcoin community eliminates collusion by banning large miners. When only transactions with value below \hat{v} are using Bitcoin, the revenue of the Bitcoin system can be obtained similar to (6) (by replacing V_{max} with \hat{v} , N by $\frac{N\hat{v}}{V_{max}}$ and γ with $n_F \frac{V_{max}}{N\hat{v}}$). Specifically, we have

$$R = \hat{v} \left(1 - \frac{n_F V_{max}}{N\hat{v}}\right) \rho \alpha_h \frac{n_F V_{max}}{N\hat{v}} \frac{N\hat{v}}{V_{max}} \quad (39)$$

Substituting the value of R in (38) and simplifying for \hat{v} , we obtain

$$\hat{v} - \frac{n_F}{N} V_{max} \geq \hat{v} \eta \frac{N}{n_F} \quad (40)$$

where

$$\eta = \frac{1-(1-\theta)^2}{(1-\theta)^2} \times \frac{1}{\rho \alpha_h N} \quad ; \quad \frac{\partial \eta(\theta)}{\partial \theta} > 0 \quad \forall \quad \theta \in [0, 1] \quad (41)$$

Note that the inequality can only be satisfied if $\eta < \frac{n_F}{N}$. Both the block revenue (LHS) and the adversary's double-spend payoff (RHS) in (40) grow linearly with \hat{v} . This is true for $\hat{v} \in [0, V_{max}]$; there is no demand for payment values greater than V_{max} . Let a $\hat{v} \in [0, V_{max}]$ exist such that the condition above is satisfied for at least one value of $n_F \in [0, N]$. This implies that the condition must also be satisfied for $\hat{v} = V_{max}$ for at least that same value of n_F . A marginal increase in \hat{v} increases the block revenue more than the adversary's double-spend payoff. This condition results in a quadratic inequality in n_F at $\hat{v} = V_{max}$:

$$\left(\frac{n_F}{N}\right)^2 - \left(\frac{n_F}{N}\right) + \eta \leq 0 \quad (42)$$

satisfied for

$$n_F \in [1 - (n_F)_{max}, (n_F)_{max}] \quad ; \quad (n_F)_{max} = \frac{N}{2} + \frac{N}{2} \sqrt{1 - 4\eta(\theta)} \quad (43)$$

where $(n_F)_{max}$ is the larger of the two solutions of the quadratic when this inequality is tight.

At $n_F = N$, the double-spend protection collapses ($\eta(\theta) = 0$ only at $\theta = 0$). This represents an extreme setting whereby the block capacity is sufficient to serve the user demand. In the absence of collusion, the block-fee revenue decreases to near zero. This makes a double-spend attack worthwhile for even a minuscule adversary ($\theta = 0 + \epsilon$). Strategic miners do not want to deploy any mining hardware. Thus, an attempt to eliminate collusion and serve the full user demand is catastrophic.

If the adversary is extremely powerful, no real-valued solution for γ exists that satisfies $\hat{v} \leq V_{max}$. This occurs when $\eta(\theta) < 1/4$. This condition simplifies to $\theta > \theta_{UB}$, where

$$\theta_{UB} = 1 - \sqrt{\frac{4}{4 + \rho\alpha_h N}} \quad (44)$$

In this case, even the revenue-maximizing capacity ($n_F = N/2$) may not be sufficient to avert such a powerful adversary ($\theta > \theta_{UB}$).

More generally, the existence of such an adversary (say, $\theta < \theta_{UB}$) places a bound on the maximum block capacity. The unique solution for n_F decreases in θ . 18 shows the tradeoff between block capacity and minimum adversary power required for a worthwhile double-spend attack. Security and collusion are a double-edged sword. Collusion reduces the artificial capacity (n_P). If collusion is eliminated, the actual block capacity (n_F) must be reduced to guarantee security.

Proposition 5 *If collusion is eliminated, the Bitcoin block size remains bounded by the threat of double-spend attacks:*

$$(n_F)_{max} = \frac{N}{2} + \frac{N}{2} \sqrt{1 - 4\eta(\theta)} \quad (45)$$

- *An adversary greater than θ_{UB} would perform an attack, even if the block size is set at the block-revenue-maximizing level of $\frac{N}{2}$.*
- *The adversary threat decreases with the demand N . A larger exogenous demand for payments provides greater security.*

In July 2014, Bitcoin faced the first and perhaps only realistic risk of a 51% double-spend attack. A single mining pool, Ghash.IO, accumulated close to 51% of the total computing power. While mining pools are composed of a large number of contributing miners, a pool owner typically coordinates block addition. The pools total power would have allowed Ghash.IO to easily launch a double-spend attack. Ghash.IO voluntarily promised to reduce its power below 40% to minimize the potential risk of double spend (ArsTechnica 2014). Ghash.IO acted rationally since they most likely expected to earn more from continued honest block revenues than a double-spend attack. The average transaction

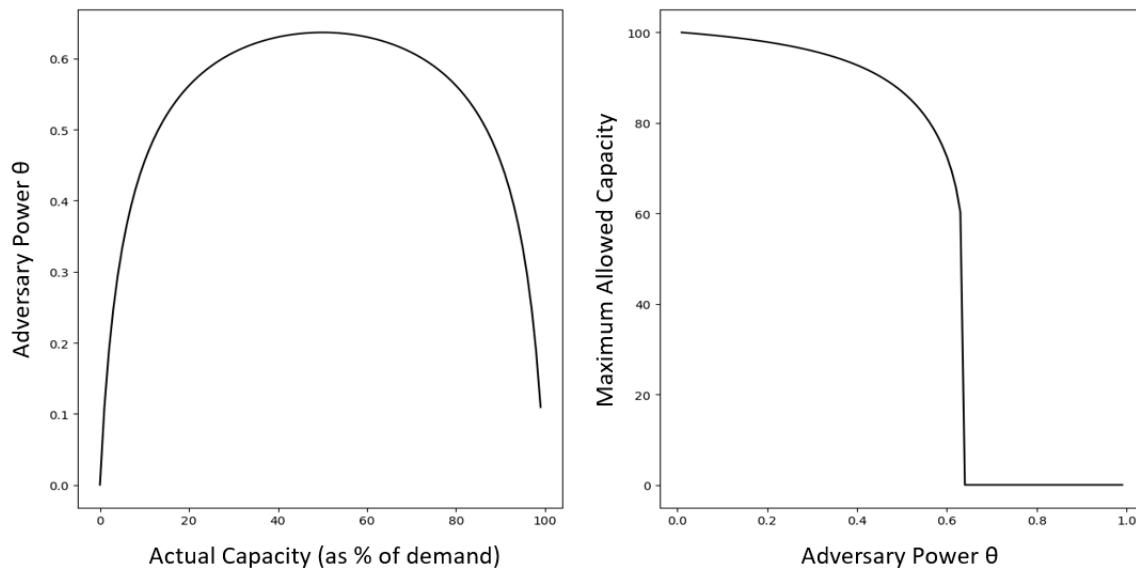


Figure 18 (Left) The adversary power required to perform double-spend attacks decreases if the capacity is close to 100% or 0% of demand. (Right) The maximum allowed capacity n_F decreases with adversary power. Note that even a $\theta > 0.5$ adversary may not perform an attack if there are sufficient revenues to be acquired.

value on Bitcoin was \$5000, while a single-block reward was worth \$100,000 (bitinfocharts 2018). Ethereum is the 2nd-largest public blockchain by usage and market capitalization after Bitcoin. Ethereum has also not yet faced a double-spend attack. It is not surprising that miners preferred to continue earning the reward instead of launching a double-spend attack.

Bitcoin Gold is a fork of Bitcoin launched to eliminate large ASIC miners. Bitcoin Gold has a substantially smaller demand from users and is valued 100-times less than Bitcoin. In May 2018, Bitcoin Gold experienced double-spend attacks. A payment worth \$17.5 Mn (Osborne 2018) was stolen at a time when Bitcoin Gold blocks paid less than \$1000 per block (bitinfocharts 2018). Ethereum classic is a fork of Ethereum valued at 30-times less than Ether (Ethereum coin) and pays \$20 per block (40 blocks ever 10 minutes). In January 2019, Ethereum classic suffered a theft of \$1.1 Mn from a double-spend attack (Casey 2019). Crypto51 (2019) estimates the cost to rent sufficient hashing power to match the current network hashing power for an hour. This cost is estimated at \$239,263 for Bitcoin, \$79,699 for Ethereum, \$3,969 for Ethereum Classic and \$737 for Bitcoin Gold. This cost is a reasonable proxy for comparing the cost of a double-spend attack. Both Ethereum Classic and Bitcoin Gold have low mining power due to a combination of smaller revenues and their ASIC-resistant mining puzzles.

The Bitcoin community wants to upgrade the block size and mitigate collusion by banning ASIC mining. Removing collusion has two effects (1) lower block revenues for the remaining miners and

(2) zero revenues for the banned large miners. Both of these factors invite double-spend attacks and make the platform insecure for large transactions. The block capacity cannot arbitrarily be raised without an adverse impact on security and thus the utility of the platform for users.

5. Conclusion

Bitcoin is a decentralized platform that provides P2P payments. It is a promising idea that attempts to provide trust without an intermediary. Naturally, Bitcoin's ability to mimic roles typically played by intermediaries has come under scrutiny. Two major questions have received significant attention. First, researchers have developed numerous attack strategies to highlight its security limitations. Second, practitioners have questioned its ability to scale with user demand. Our work places economic bounds on the scalability and security of Bitcoin.

Bitcoin allows 3 transaction per second, compared to 5000 per second on VISA and 400 per second on SWIFT. Contrary to popular belief, merely upgrading the block size does not scale Bitcoin. We show why miners will under-utilize block capacities to force users into competition on fees. This tacit collusion among miners can be sustained by an increasingly smaller group if the block capacity is increased. Artificial capacity constraints created by collusion discriminate among users. Users with large-valued payments stand to earn the most by avoiding proportional fees levied by traditional banks. These users will be willing to pay higher fees for preferential treatment on the small throughput. This reinforces Bitcoin as a currency for the rich.

Bitcoin mining puzzles could be re-designed to eliminate large miners. This promises to remove collusion-driven artificial constraints and serve greater user demand at low fees. Unfortunately, this results in too little mining revenue. As miners exit, the platform becomes susceptible to double-spend attacks. The fees raised by competition among users are not an idiosyncratic feature of Bitcoin but rather are a critical feature to thwart hacks. The scale of Bitcoin is economically bounded by threats to security.

A few limitations of our research are worth highlighting. First, we model an equilibrium setting (a few years down the line) whereby Bitcoin block revenue is completely composed of fees. Currently, transaction fees make up less than 5% of miner revenues. Nevertheless, miners who function on small margins (between revenue and cost) still care about this small component. As the Bitcoin supply via block rewards goes to zero, the mining network dependence on fee competition will only be exacerbated. While our model implications remain true today, we expect strategic under-filling to become more aggressive in the future. Second, we largely ignore Bitcoin features, such as privacy, and traditional bank services (e.g., liquidity and customer service) in the user's utility. Researchers in

the future could model user choice more thoroughly. Third, investors and speculators have played a major role in Bitcoin valuation and demand. We model users and miners and discuss some options for Bitcoin designers. However, we refrain from modeling investors. Finally, we only model the specific consensus mechanism of Bitcoin. Bitcoin is now one of hundreds of public blockchains. We expect the core intuition in the paper to be broadly applicable to major proof-of-work public blockchains.

Bitcoin makes the promise of avoiding profit-maximizing intermediaries. However, it suffers various drawbacks due to its inability to respond to issues such as scale and security. These limitations have led to calls for cryptocurrency regulation. Increasingly, new private blockchain are confining P2P interactions under the oversight of a central authority. We highlight a decentralized platform's vulnerability in serving user demand efficiently. We also show how a large number of anonymous miners start to resemble a traditionally organized intermediary. It is critical that policy makers not consider Bitcoin a wild west of irrational enthusiasts. Instead, they must closely consider the economic incentives that determine the strategic actions of Bitcoin participants.

References

- Arora, A., A. Greenwald, K. Kannan, and R. Krishnan. 2007. Effects of information-revelation policies under market-structure uncertainty. *Management Science* 53 (8): 1234–1248.
- ArsTechnica 2014. Bitcoin pool GHash.io commits to 40% hashrate limit after its 51% breach. *arstechnica.com*.
- Asvanund, A., K. Clay, R. Krishnan, and M. D. Smith. 2004. An empirical analysis of network externalities in peer-to-peer music-sharing networks. *Information Systems Research* 15 (2): 155–174.
- Athey, S., and G. Ellison. 2008. mPosition Auctions with Consumer Search, nmimeo.
- August, T., M. F. Niculescu, and H. Shin. 2014. Cloud implications on software network structure and security risks. *Information Systems Research* 25 (3): 489–510.
- Biais, B., C. Bisiere, M. Bouvard, and C. Casamatta. 2018. The Blockchain Folk Theorem.
- Bitcoin Wiki 2017. SegWit2x. Available at:
<https://en.bitcoin.it/wiki/SegWit2x>.
- bitinfocharts 2018. Bitcoin Avg. Transaction Value historical chart. <https://bitinfocharts.com>.
- Blockchain.com 2018. <https://www.blockchain.com/en/btc/tx/261d69b25896034325d8ad3e0668f963346fd79baefb6a>
- BlockchainExplorer 2018. *Blockchain.com*.
- Buntinx, J. 2017. Bitcoins Transaction fee exceeds \$50 as Network Issues Remain. *newsbtc.com*.

- Casey, M. J. 2019. The Ethereum Classic Attacker Has Sent a Bigger Message. *coindesk.com*.
- Cezar, A., H. Cavusoglu, and S. Raghunathan. 2013. Outsourcing information security: Contracting issues and security implications. *Management Science* 60 (3): 638–657.
- Chang, M.-H. et al. 1991. The effects of product differentiation on collusive pricing. *International Journal of Industrial Organization* 9 (3): 453–469.
- Chase, J. 2018.
- Cong, L. W., Z. He, and J. Li. 2018. Decentralized mining in centralized pools.
- Cong, L. W., Y. Li, and N. Wang. 2018. Tokenomics: Dynamic adoption and valuation.
- Crypto51 2019. PoW 51% Attack Cost. *crypto51.app*.
- Dey, D., A. Kim, and A. Lahiri. 2018. Online Piracy and the Longer Arm of Enforcement. *Management Science*.
- Dwyer, B. 2018. Credit card processing fees seem complicated. But take away the sales jargon and it makes more sense. *CardFellow.com*.
- Eyal, I., A. E. Gencer, E. G. Sirer, and R. Van Renesse. 2016. Bitcoin-NG: A Scalable Blockchain Protocol. In *NSDI*, 45–59.
- Gao, X., W. Zhong, and S. Mei. 2013. Information security investment when hackers disseminate knowledge. *Decision Analysis* 10 (4): 352–368.
- Gervais, A., G. O. Karame, K. Wüst, V. Glykantzis, H. Ritzdorf, and S. Capkun. 2016. On the security and performance of proof of work blockchains. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 3–16. ACM.
- Glazer, E. 2018. Justice Department Probing Wells Fargos Wholesale Banking Unit. *Wall Street Journal*.
- Glazer, E., and M. Farrell. 2018. Big U.S. Banks Face Increase in Attempted Cyberattacks. *Wall Street Journal*.
- Häckner, J. 1994. Collusive pricing in markets for vertically differentiated products. *International Journal of Industrial Organization* 12 (2): 155–177.
- Hsu, C., J.-N. Lee, and D. W. Straub. 2012. Institutional influences on information systems security innovations. *Information systems research* 23 (3-part-2): 918–939.
- Huberman, G., J. D. Leshno, and C. C. Moallemi. 2017. Monopoly Without a Monopolist: An Economic Analysis of the Bitcoin Payment System.

- Johar, M., S. Menon, and V. Mookerjee. 2011. Analyzing sharing in peer-to-peer networks under various congestion measures. *Information Systems Research* 22 (2): 325–345.
- Jordan, D., and S. S. Kerr. 2018. Baffled by Bitcoin? How Cryptocurrency Works. *Wall Street Journal*.
- Kannan, K., and R. Telang. 2005. Market for software vulnerabilities? Think again. *Management science* 51 (5): 726–740.
- Li, Z., and A. Agarwal. 2016. Platform integration and demand spillovers in complementary markets: evidence from Facebooks integration of Instagram. *Management Science* 63 (10): 3438–3458.
- Malinova, K., and A. Park. 2017. Market Design with Blockchain Technology.
- McMillan, R. 2018. Thieves Can Now Nab Your Data in a Few Minutes for a Few Bucks. *Wall Street Journal*.
- Nakamoto, S. 2008. Bitcoin: A Peer-To-Peer Electronic Cash System.
- Natoli, C., and V. Gramoli. 2017. The Balance Attack or Why Forkable Blockchains Are Ill-Suited for Consortium. In *Dependable Systems and Networks (DSN), 2017 47th Annual IEEE/IFIP International Conference on*, 579–590. IEEE.
- Nayak, K., S. Kumar, A. Miller, and E. Shi. 2016. Stubborn Mining: Generalizing Selfish Mining and Combining with an Eclipse Attack. In *Security and Privacy (EuroS&P), 2016 IEEE European Symposium on*, 305–320. IEEE.
- Osborne, C. 2018. Bitcoin Gold suffers double spend attacks, \$17.5 million lost. *zdnet.com*.
- Popper, N. 2017. Bitcoin Hasnt Replaced Cash, but Investors Dont Care. *The New York Times*.
- Rosenfeld, M. 2014. Analysis of hashrate-based double spending. *arXiv preprint arXiv:1402.2009*.
- Ross, T. W. 1992. Cartel stability and product differentiation. *International Journal of Industrial Organization* 10 (1): 1–13.
- Shy, O., and Z. Wang. 2011. Why do payment card networks charge proportional fees? *American Economic Review* 101 (4): 1575–90.
- Sompolinsky, Y., and A. Zohar. 2015. Secure High-Rate Transaction Processing in Bitcoin. In *International Conference on Financial Cryptography and Data Security*, 507–527. Springer.
- Tasca, P. 2018. The Hope and Betrayal of Blockchain. *The New York Times*.
- Thomadsen, R. 2007. Product positioning and competition: The role of location in the fast food industry. *Marketing Science* 26 (6): 792–804.
- TransferWise 2018. How long does an international wire transfer take? *transferwise.com*.

- Wei, Z., and M. Lin. 2016. Market mechanisms in online peer-to-peer lending. *Management Science* 63 (12): 4236–4257.
- Wilmoth, J. 2018. The First 8MB Bitcoin Cash Block Was Just Mined.
- Yli-Huumo, J., D. Ko, S. Choi, S. Park, and K. Smolander. 2016. Where is current research on blockchain technology? a systematic review. *PloS one* 11 (10): e0163477.
- Zaitsev, D. 2018. Quarterly Cryptocurrency and ICO Market Analysis. *Medium.com*.
- Zhu, Y., and K. C. Wilbur. 2011. Hybrid advertising auctions. *Marketing Science* 30 (2): 249–273.

Appendices

A. Readings on Bitcoin

Bitcoin’s ecosystem is composed of four major components: (1) users, (2) miners, (3) the platform protocol, and (4) the cryptocurrency. We suggest Huberman et al. (2017) for a deeper understanding of user waiting queues and transaction fee decisions. Cong et al. (2018) provides an in-depth discussion of the arms race by miners for computing hardware and their organization into mining pools. The protocol itself is most accurately described by Satoshi Nakamoto – the creator of Bitcoin (Nakamoto 2008). Finally, Cong et al. (2018) can be used as a resource for a primer on cryptocurrencies and their adoption.