

Bitcoin: A Natural Oligopoly

Nick Arnosti*

Matt Weinberg†

Abstract

We argue that the concentrated production and ownership of Bitcoin mining hardware arises naturally from the economic incentives of Bitcoin mining. We model Bitcoin mining as a two-stage competition: miners compete in prices to sell hardware, while competing in quantities for mining rewards. We characterize equilibria in our model, and show that small asymmetries in operational costs result in highly concentrated *ownership* of mining equipment. We further show that *production* of mining equipment will be dominated by the miner with the most efficient hardware, who will sell hardware to competitors while possibly also using it to mine.

1 Introduction

In the years since Bitcoin was introduced by Nakamoto (2008), cryptocurrencies have attracted a great deal of funding and media coverage. Bitcoin’s market capitalization now exceeds \$900 billion.¹ The protocol underlying Bitcoin is unquestionably clever, and has largely succeeded in creating a public yet anonymized record of transactions. In principle, these transactions can be authorized by anyone who wishes to become a Bitcoin “miner.”

In practice, it is widely acknowledged that Bitcoin mining is controlled by a small number of large entities. This concentration is frequently cited as a major concern, and has motivated several new cryptocurrencies. For example, the Nxt whitepaper² states,

Bitcoin’s creator, Satoshi Nakamoto, intended for the bitcoin network to be fully decentralized, but nobody could have predicted that the incentives provided by Proof of Work systems would result in the centralization of the mining process.

Meanwhile, the abstract of the whitepaper introducing Bitcoin Gold³ states,

The purpose [of Bitcoin Gold] is to make Bitcoin mining decentralized again. Satoshi Nakamoto’s idealistic vision of “one CPU one vote” has been superseded

*Columbia Business School

†Princeton University, Department of Computer Science.

¹Source: www.coinmarketcap.com. Accessed March 3, 2021.

²Source: <https://whitepaperdatabase.com/nxt-nxt-whitepaper/>

³Source: <https://bitcoingold.org/wp-content/uploads/2017/10/BitcoinGold-Roadmap.pdf>

by a reality where the manufacture and distribution of mining equipment has become dominated by a very small number of entities.

These quotes highlight different ways in which Bitcoin mining is centralized. The first quote addresses centralized *ownership*: most of the hardware used to mine Bitcoin is controlled by a few large organizations (Hileman and Rauchs, 2017; Gencer et al., 2018). The second quote addresses centralized *production*: a 2018 Bitmain prospectus estimates that mining equipment manufactured by Bitmain accounts for 75% of the market.⁴ Interestingly, Bitmain also owns and operates its own mining equipment: the prospectus claims that Bitmain earned mining revenue of \$53 million in 2016 and \$199 million in 2017, accounting for 9.6% and 6.2% of the global total, respectively.⁵

Two key questions motivate this work. First, why is the ownership of mining equipment so concentrated? Second, why is the production of mining equipment so concentrated?

We address these questions by modeling Bitcoin mining as a two-stage contest, where miners compete over a fixed reward. Each miner has access to different hardware, and pays different operational costs for electricity, land, wages, and other expenses. In the first stage, miners set prices at which they are willing to sell their hardware to competitors. In the second stage, miners choose what hardware to acquire and power. Miners split the available mining rewards proportionally to their computational power.

Regarding our first question, our model predicts that even small differences in operational costs lead to highly concentrated *ownership* of mining equipment. Regarding the second, our model predicts that the miner with the best hardware will sell it at a price low enough to capture the whole market, regardless of whether it also chooses to mine. This implies highly concentrated *production* of mining equipment.

Below, we provide more detail about Bitcoin mining, our model, and our results.

1.1 What is Bitcoin Mining?

Bitcoin is a digital currency, with the innovative feature that the record of transactions is not maintained by a single entity, but rather by a collection of self-appointed “miners.” Whenever a miner verifies a block of transactions, it receives a “block reward” (currently 6.25 bitcoin), in addition to fees offered by those whose transactions are included. Each block must be accompanied by the solution to a cryptopuzzle.⁶ The fastest way to solve

⁴Source: <http://enterprise.press/wp-content/uploads/2018/09/BitmainProspectus.pdf>

⁵Source: <https://www.statista.com/statistics/731383/bitcoin-mining-revenue/>.

⁶Narayanan et al. (2016) has a deeper explanation of these puzzles, and why they are used.

these puzzles is to guess solutions at random, so the rate at which miners are rewarded is proportional to their rate of guessing (“hash rate”). Importantly, the difficulty of these puzzles adjusts so that on average, a new block is created every ten minutes, regardless of the total hash rate in the network.

Although all miners guess solutions randomly, there are several ways to reduce the cost of each guess:

- **Hardware.** Specialized hardware (“ASICs”) designed to solve bitcoin cryptopuzzles uses 1000 times less energy per guess, compared to general-purpose GPUs.⁷
- **Electricity.** Cheap electricity lowers the cost of powering mining hardware.
- **Cooling.** Cold locations lower the cost of cooling mining hardware.
- **Land & Labor.** Cheap land and labor lower the cost of operation.

These factors play a significant role in determining who enters the market: only miners with low costs can mine profitably. This contributes to mining concentration, which can be measured in several ways.

Ownership of Mining Equipment. Although the anonymity of bitcoin mining makes it difficult to determine ownership, studies estimate that a majority of global mining power is controlled by eight to eleven large miners (Hileman and Rauchs, 2017; Gencer et al., 2018). Together, these miners could freeze any user’s funds, erase past transactions, or launch other attacks. Narayanan et al. (2016) describes these risks in more detail.

Production of Mining Equipment. A 2018 prospectus estimates that 75% of Bitcoin mining uses equipment manufactured by Bitmain. The risks associated with this concentration were highlighted in 2017, when a backdoor was discovered that gave Bitmain the ability to shut down any Antminers that it had produced.⁸

Mining Pools. “Mining pools” are groups of miners that share rewards with each other in order to decrease the variance of short-term returns. Most blocks come from a small number of large pools.⁹ Many of the remainder are of unknown origin, and might also come from large pools.¹⁰ This concentration is concerning because a single pool operator

⁷Sources: https://en.bitcoin.it/wiki/Mining_hardware_comparison, https://en.bitcoin.it/wiki/Non-specialized_hardware_comparison.

⁸See <https://www.rudebaguette.com/en/2017/05/antbleed-bitmain-shut-half-bitcoin/>.

⁹Current pool shares are available at <https://blockchain.info/pools?timespan=4days>.

¹⁰This possibility is discussed here <https://diar.co/volume-3-issue-1/>, and seems likely given that “unknown” blocks increased dramatically shortly after Bitmain’s stake in multiple large pools attracted attention. Bitmain owns Antpool and BTC.com, and is the sole investor in viaBTC: bitcoinmagazine.com/articles/bitmain-nears-51-network-hash-rate-why-matters-and-why-it-doesnt.

typically determines the content of blocks mined by the pool. However, Cong et al. (2018) argue that if miners can easily switch between pools then the pool size distribution is irrelevant. One could even claim that mining pools encourage decentralization by allowing small miners to collect rewards regularly.

We do not model mining pools in this paper. Instead, we focus on concentration of ownership and production of mining equipment. As noted above, both forms of concentration are prevalent, and both pose threats to the Bitcoin network. We seek to understand the reasons for this concentration, and whether it can be expected to persist.

1.2 Model Overview

We highlight several key features of the Bitcoin protocol, which form the basis of our model:

- The total value of rewards available to miners is fixed.¹¹
- Miners make costly investments in computational power. In particular:
 - Miners produce proprietary hardware, which they can sell to competitors.
 - Miners pay different operational costs (for electricity, land, cooling, labor, etc.).
- Miners earn rewards in proportion to their computational power.

In essence, there are two competitions occurring simultaneously: hardware producers are competing in price, while miners are competing in quantities. However, the line between hardware producer and miner is a blurry one.

We capture this using a stylized two-stage model in which miners have varying hardware production costs and operational costs. In the first stage, miners post a price at which they are willing to sell their hardware. In the second stage, miners decide how much computational power to acquire, and split a fixed prize in proportion to their computational power. We model the second stage game as an asymmetric Tullock rent-seeking contest, with asymmetries arising from differences in operating costs and access to high-quality hardware. Meanwhile, the first stage game is a Bertrand competition: miners will always purchase the hardware that allows them to mine for the lowest cost.

¹¹As explained above, on average one new block is created every 10 minutes, and the size of the block reward is fixed. Furthermore, Huberman et al. (2017) conclude that transaction fees should not depend on the hash rate, nor on how it is distributed among miners.

1.3 Overview of Results: Concentration of Ownership

For any hardware prices set in the first stage, there is an essentially unique second-stage equilibrium, which we describe in Theorem 1. If all costs c_i are identical, then in equilibrium, each miner possesses an equal amount of mining power (Corollary 1).

One might expect ownership to remain decentralized so long as costs are not “too different.” Our (qualitative) finding is that even small cost advantages may cause significant concentration. More precisely, each miner’s market share in equilibrium is equal to the percentage by which their costs are lower than the market-wide “break even cost” (Corollary 2). For example, a miner with costs that are 10% lower than this break-even cost will possess 10% of all mining power.

This mining concentration has economic implications. As we discuss in Section 2, many analyses of Bitcoin mining assume the market is competitive, and that miners make little profit. This assumption is also incorporated into one frequently-cited method for estimating Bitcoin’s energy footprint. Our model makes a different prediction. Corollary 3 establishes that total mining profit is proportional to the Herfindahl-Hirschman Index of market concentration: concentrated mining implies significant miner profits. This suggests that the aforementioned estimate of Bitcoin’s energy consumption is likely too high.

In summary, the thesis of Section 4 is that the Bitcoin reward scheme naturally induces a concentrated market in which low-cost miners can earn significant profits. This finding has implications for various analyses of Bitcoin, including estimates of its carbon footprint.

1.4 Overview of Results: Concentration of Production

It is natural to wonder why a hardware producer would sell their hardware, rather than simply using it to mine. For example, the CEO of blockchain company Sia wrote¹²

At the end of the day, cryptocurrency miner manufacturers are selling money printing machines... The buyer needs to understand why the manufacturer is selling the units instead of keeping them for themselves.

One white paper proposes the following intuitive answer:¹³

¹²Source: <https://blog.sia.tech/the-state-of-cryptocurrency-mining-538004a37f9b>

¹³Source: https://research.ark-invest.com/hubfs/1_Download_Files_ARK-Invest/White_Papers/ARKInvest_031220_Whitepaper_BitcoinMining.pdf.

The expertise to operate mining facilities is very different from that to manufacture hardware... Moreover, without a competitive advantage in sourcing electricity, mining machine manufacturers could be uncompetitive despite their hardware cost advantages.

In other words, the manufacturer may be unable to use their equipment to mine profitably.

Our model proposes a different answer: Theorem 2 establishes that there is always an equilibrium in which the lowest-cost hardware is *sold to the entire market*. This equilibrium arises even if the hardware manufacturer has access to cheap electricity and chooses to mine. Moreover, Theorem 4 establishes that in many cases, this outcome is the *only* equilibrium.

The reason for this is that the competition among mining manufacturers resembles a Bertrand competition: miners will purchase whichever equipment enables them to mine most cheaply. Competition from other manufacturers drives prices to the point that *all* active miners end up using the most efficient hardware. Although the miner who manufactures this hardware might be tempted to raise its price, it does not truly control delivery: a price increase opens the door for another manufacturer to step in and capture the market.

To summarize, our main conclusions regarding the sale of hardware are that (a) production of Bitcoin mining hardware naturally concentrates, and (b) a miner with good hardware will choose to sell it to the entire market, even while simultaneously using it to mine. Section 5 presents these findings in more detail.

2 Related Work

We focus on three areas of related literature. First, we discuss papers on Bitcoin which justify some of our modeling assumptions, but are for the most part mathematically unrelated to our work. Second, we discuss rent-seeking contests, which are closely related to our second-stage game. Finally, we discuss the literature on price discrimination in input markets, which uses a closely related model in which one agent can sell hardware (but cannot mine), and the remaining agents mine (but cannot sell hardware to each other).

2.1 Cryptocurrencies

Though cryptocurrencies are relatively new, the literature on them is growing quickly. Halaburda and Haeringer (2019) provide a helpful survey of papers studying economic and game-theoretic questions related to cryptocurrencies.

Several sources document concentration in Bitcoin mining. The most readily available statistics relate to mining pools: currently, four mining pools are responsible for over 50% of mined blocks.¹⁴ Other measures of centralization are not publicly available, but Romiti et al. (2019) present evidence that rewards in several mining pools are concentrated among a few miners, Hileman and Rauchs (2017) estimate that eleven “large mining entities” collectively control a majority of Bitcoin’s global hashrate, and Gencer et al. (2018) conclude that 75% of all blocks are initially announced by one of 100 Bitcoin nodes.¹⁵

Other papers provide justification for our modeling assumptions. Our assumption that miners compete for a fixed prize is consistent with work by Huberman et al. (2017), who conclude that the fees paid by Bitcoin users (and total mining rewards) are not affected by miner behavior. Our focus on concentration of ownership and production of mining equipment, rather than pool size, is consistent with the conclusion of Cong et al. (2018) that the distribution of pool size is irrelevant if miners can freely switch between pools.

Another key assumption in our model is that miners are rewarded proportionally to their mining power. Recent work by Chen et al. (2019) and Leshno and Strack (2019) proves that this is the only approach that is anonymous, sybil-proof, and collusion-proof, implying that our results extend to any competition with these properties.

Many papers study strategic deviations from the Bitcoin protocol (Babaioff et al., 2012; Eyal and Sirer, 2014; Eyal, 2015; Carlsten et al., 2016; Kiayias et al., 2016; Sapirshtein et al., 2016). Whereas these papers treat miners’ computational power as fixed, we study miners’ incentives to *acquire* computational power. Prat and Walter (2018) permit miners to decide how much hardware to acquire, but assume that they have access to the same hardware, which is available at an exogenous price. By contrast, we assume that miners have different hardware, and study how they price it.

2.2 Rent-Seeking Contests

There is a large literature on “rent-seeking” activities, which seek to influence the allocation of a prize. Hillman and Riley (1989) summarize one important idea from this literature:

Often influence-seeking activities are not directly observable, but the value of the prize secured by the successful contender is known. One may then seek to infer the value of the resources allocated to influencing the allocation of a prize from the value of the prize itself.

¹⁴Source: <https://blockchain.info/pools?timespan=4days>. Accessed March 3, 2021.

¹⁵A node listens for transactions and blocks, and forwards them to the rest of the network.

Several papers apply this idea to Bitcoin mining. Kroll et al. (2013) assume that “the total mining reward... is equal to the total global cost of mining,” and Budish (2018) assumes that “the prize...is dissipated by expenditures aimed at winning the prize.” This assumption is also made by Hayes (2015), Chiu and Koepl (2018), Thum (2018), Ma et al. (2018) and Easley et al. (2019). To justify this assumption, Chiu and Koepl (2018) and Thum (2018) appeal to a class of games introduced by Tullock (1980), in which contestants share a prize in proportion to their expenditures. When contestants are symmetric, total expenditures approach the value of the prize as the number of contestants grows.

The second stage of our model is a Tullock contest with *asymmetric* contestants, which has also been studied by Hillman and Riley (1989), Gradstein (1995), Dimitri (2017) and Alsabah and Capponi (2020). The first two papers note that the number of active participants may be small, but (for obvious reasons) do not discuss implications for Bitcoin. Dimitri (2017) does not explicitly address mining concentration, except to note that at least two miners must be active in equilibrium. Alsabah and Capponi (2020) do study mining concentration, and are closest to our work. They focus on miners’ R&D investments, and consider a model with ex-ante symmetric miners.¹⁶ By contrast, our miners are inherently asymmetric, and one of our key questions is whether a miner with good hardware will want to sell it. The sale of hardware is not permitted in any of the aforementioned papers.

2.3 Selling Inputs to Competing Firms

There is a literature on price discrimination in input markets, in which a supplier sells an input to a set of downstream firms that compete in quantities. The foundational papers of Katz (1987) and DeGraba (1990) assume that if the supplier offers price p_i to downstream firm i , then the marginal cost of production for i is $p_i + \beta_i$. This resembles a version of our model in which the first miner has prohibitively high operational costs (and thus will never mine), and the remaining miners have prohibitively high hardware costs.

Yehezkel (2004) considers a model where where a manufacturer sells a necessary input to a retailer, and also competes against that retailer in a final goods market. He notes that by charging a high price to the retailer, the manufacturer softens competition in the final goods market. This effect is also present in our model, although the nature of the second-stage competition is quite different.

¹⁶The analysis of asymmetric contests is to assess off-equilibrium behavior.

3 Model

There are $n \geq 2$ miners competing for a prize of value R . Each miner i has an *operational cost* $\text{OC}_i > 0$, and a *hardware cost* $\text{HC}_i > 0$. Without loss of generality, we sort the miners in increasing order of hardware cost. The game proceeds in two stages:

1. Each miner i chooses a price $p_i \geq \text{HC}_i$ at which they will sell their hardware.¹⁷
2. Each miner i chooses the quantity q_{ij} of computational power sourced from miner j .

Given quantities q_{ij} , miner i 's **market share** is

$$x_i(q) = \frac{\sum_j q_{ij}}{\sum_{k,j} q_{kj}}. \quad (1)$$

We define $x_i(q) = 0$ if both the numerator and denominator are zero (although this will never occur in equilibrium).

At prices p , the per-unit cost to miner i of sourcing from miner j is

$$c_{ij}(p) := \text{OC}_i + \text{HC}_j I(j = i) + p_j I(j \neq i). \quad (2)$$

Miner i earns **mining profit** of

$$\Pi_i^M(p, q) = R \cdot x_i(q) - \sum_j q_{ij} c_{ij}(p), \quad (3)$$

and additional **sales profit** of

$$\Pi_i^S(p, q) = (p_i - \text{HC}_i) \sum_{j \neq i} q_{ji}, \quad (4)$$

for a total **utility** of

$$U_i(p, q) = \Pi_i^M(p, q) + \Pi_i^S(p, q). \quad (5)$$

¹⁷Our model considers anonymous rather than personalized pricing. However, most of our conclusions hold for personalized pricing as well.

Definition 1.

A function Q mapping prices to purchase quantities is a **second-stage equilibrium** if for all price vectors p , miners j , and any q' such that $q'_{i\ell} = Q_{i\ell}(p)$ for $i \neq j$,

$$U_j(p, Q(p)) \geq U_j(p, q').$$

Prices p are a **first-stage equilibrium** for Q if for any miner j , and any p' such that $p'_i = p_i$ for $i \neq j$,

$$U_j(p, Q(p)) \geq U_j(p', Q(p')).$$

The pair (p, Q) is an **equilibrium** if Q is a second-stage equilibrium and p is a first-stage equilibrium for Q .

4 The Second-Stage Game: Concentration of Ownership

Recall our motivating question: *Why is the ownership of Bitcoin mining hardware so concentrated?* This section provides an answer by analyzing the second-stage game in our model. We start by characterizing second-stage equilibria. Our characterization relies heavily on a “break-even cost” c^* : in Section 4.1, we show how this cost can be estimated from publicly available data. In Section 4.2, we answer our main question: ownership concentrates because small asymmetries in cost lead to large asymmetries in equilibrium market share. In Section 4.3, we discuss economic implications of this concentration.

We now analyze the second-stage game. For a fixed price vector p , the **effective cost** for miner i is

$$c_i(p) = \min_j c_{ij}(p). \tag{6}$$

Lemma 1. *If Q is a second-stage equilibrium, then $c_{ij}(p) > c_i(p)$ implies $Q_{ij}(p) = 0$.*

In other words, miners only acquire hardware from sources offering the lowest cost. Thus, for the purposes of calculating mining profit, we can ignore *who* each miner purchases from, and focus on *how much* computational power each miner acquires.

The resulting game is a rent-seeking contest, as described by Tullock (1980). Although several previous papers have characterized the equilibrium of this game, we give a new

characterization in Theorem 1 below. We begin by defining the $X : \mathbb{R}_+ \times \mathbb{R}_+^n \rightarrow \mathbb{R}_+$ by

$$X(y, \vec{c}) = \sum_{i=1}^n \max\{1 - c_i/y, 0\}. \quad (7)$$

Lemma 2. *For any \vec{c} , there is a unique value c^* satisfying $X(c^*, \vec{c}) = 1$. Furthermore, increasing any $c_i < c^*$ strictly increases c^* , while increasing any $c_i \geq c^*$ leaves c^* unchanged.*

For the remainder of this paper, we let $c^* = c^*(p)$ denote the unique solution to $X(c^*(p), \vec{c}(p)) = 1$. The following result establishes that $c^*(p)$ and $c_i(p)$ are sufficient for determining equilibrium outcomes for miner i .

Theorem 1. *For any prices p , and any second-stage equilibrium Q , the total hash rate is*

$$\sum_{i,j} Q_{ij}(p) = \frac{R}{c^*(p)}, \quad (8)$$

the market share of miner i is

$$x_i(Q(p)) = \max\left\{1 - \frac{c_i(p)}{c^*(p)}, 0\right\}, \quad (9)$$

and the mining profit of miner i is

$$\Pi_i^M(p, Q(p)) = x_i(Q(p))^2 R. \quad (10)$$

The proof of Theorem 1 is in Appendix B.

4.1 Using Theorem 1 to Calculate a Break-Even Cost

Before proceeding, we show how Theorem 1 can be applied to Bitcoin mining. There are several other characterizations of equilibria in Tullock’s rent-seeking model (see e.g. Hillman and Riley (1989), Gradstein (1995), Dimitri (2017), Alsabab and Capponi (2020)), which express outcomes in terms of the number of active miners and each of their costs. Unfortunately, these factors are not observable in practice. By contrast, the expressions in Theorem 1 rely heavily on the quantity c^* , which can be easily estimated from available data. By (8), c^* is equal to the ratio $R/\sum Q_{ij}$. The quantity R corresponds to the rate of mining revenue, which is easily tracked. Meanwhile $\sum Q_{ij}$ corresponds to the total hash rate, which can be accurately estimated from the difficulty and rate of block discovery.

Let’s apply this idea. On October 13, 2020, total daily mining revenue was approximately 11.6 million USD, and total hash rate was approximately 144.3 ExaHash/second.¹⁸ This implies that

$$c^* \approx \frac{\$11.6 \text{ million/day}}{144.3 \text{ ExaHash/second}} \times \frac{1 \text{ day}}{60 \cdot 60 \cdot 24 \text{ seconds}} = \frac{\$0.93}{\text{ExaHash}}.$$

This value serves as a break-even point: given the current reward and difficulty of block discovery, any miners with average cost above c^* cannot mine profitably, while those with average cost below c^* will earn profits from mining.¹⁹

4.2 A Natural Oligopoly: Linking Cost, Market Share, and Profit

The claim that only miners with costs below c^* can make a profit is almost model-free, and should not be controversial. What is less clear is, for a miner with cost $c_i < c^*$, how much hardware will they control, and what profit will they make? We now use Theorem 1 to answer this question, and study concentration of ownership.

If we assume that there are many miners with identical costs, the following corollary establishes that all of these costs must be close to c^* , and that each miner will control a small share of the mining market.

Corollary 1. *If all miners have identical effective costs $c_1(p) = c_2(p) = \dots = c_n(p)$, then for all i , $\frac{c_i}{c^*} = \frac{n-1}{n}$ and $x_i(Q(p)) = 1/n$.*

In reality, miners have different costs, and one should expect that miners with lower costs will choose to operate more mining hardware, and will make more profit. A priori it is not clear how strong these effects should be. One might hope that small cost asymmetries (i.e. on the order of 5-10%) would lead to only small asymmetries in market share. Unfortunately, the following Corollary of Theorem 1 shows otherwise: minor asymmetries in cost can imply drastic asymmetries in market share.

¹⁸An ExaHash is 10^{18} hashes, or 1 million TeraHashes. Data reflects a 7-day average, taken from <https://www.blockchain.com/charts/miners-revenue> and <https://www.blockchain.com/charts/hash-rate>.

¹⁹Some Bitcoin blogs refer to a “break even” price of Bitcoin. This supposed to be the price at which mining would become unprofitable, assuming mining difficulty remains constant. This is an incoherent counterfactual: as Dimitri (2017) points out, if the price of Bitcoin falls, then so will the network hash rate. Our use of the phrase “break even” is quite different: for a fixed difficulty *and price* of Bitcoin, c^* is a “break even cost” because miners are able to make a profit if and only if their cost is below c^* .

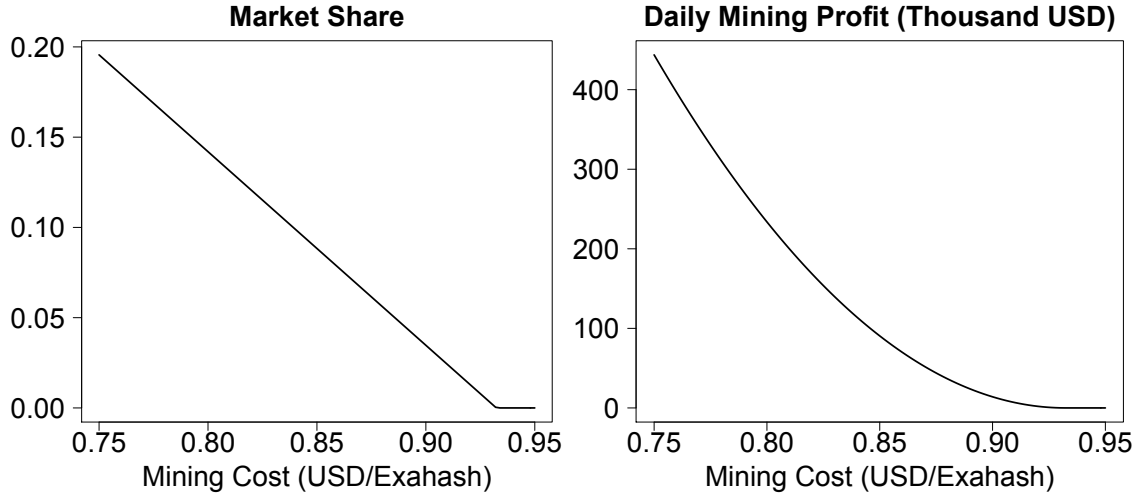


Figure 1: From the estimated break-even cost of \$0.93/ExaHash, our model predicts a miner’s market share (left) and profit (right) as a function of its marginal cost. It suggests that a miner with costs 10% below the break-even cost could control 10% of the market and be making over \$100,000 in daily profit.

Corollary 2. *In equilibrium, $c_i(p) = (1 - x_i(Q(p)))c^*(p)$ for each i . That is, miner i ’s market share is precisely the percentage by which its effective cost $c_i(p)$ is below $c^*(p)$.*

To get intuition for Corollary 2, consider a miner with marginal cost of hash power equal to $c_i < c^*$. This miner’s marginal revenue from additional hash power is equal to c^* times the fraction of mining hardware controlled by competitors. In order for marginal cost to equal marginal revenue, the fraction of hardware controlled by others must be c_i/c^* .

We now apply Corollary 2 through a simple example. If we use our previously calculated break-even cost of \$0.93/ExaHash, then a miner who pays \$0.88/ExaHash (5.4% below the break-even cost) should control 5.4% of all mining hardware in equilibrium, while a miner who pays \$0.83/ExaHash is predicted to control 10.8% of all mining hardware. The latter miner owns *twice* as much hardware, despite a seemingly minor 5.7% cost advantage. The imbalance in mining profit, as predicted by (10), is even more extreme: the latter miner earns *four times* as much profit as the former. Figure 1 plots estimated mining profit as a function of mining cost — note the sharp growth in profits as the cost drops.

We conclude that Bitcoin (and other cryptocurrencies based on proof-of-work) are natural oligopolies. Inevitable differences in electricity costs, cooling costs, and local wages will result in much larger differences in market share. A miner with costs moderately below

the break-even cost c^* should in fact control a significant share of *all* mining hardware.

For anyone who considers decentralized mining as an explicit goal of the Bitcoin protocol, this is disappointing news. It also invites the question of whether a network in which most transactions are processed by a handful of large miners is truly secure. However, the conclusion that Bitcoin mining is an oligopoly rather than a competitive market also has implications for the seemingly unrelated question of the energy consumption of the Bitcoin network, as we now discuss.

4.3 Implications For Calculating Bitcoin’s Energy Footprint

Many papers that model Bitcoin mining as a rent-seeking contest assume that the mining market is competitive (see Section 1). Section 4.2 establishes that Bitcoin is instead an oligopoly (and should be expected to remain an oligopoly). Beyond its standalone interest, this has implications for other economic analyses. This section considers one such example.

One aspect of Bitcoin that has attracted significant attention is its large carbon footprint. Many studies attempt to estimate the total energy consumption of the Bitcoin network. The challenge is that although the network hash rate can be readily estimated, it is much harder to know the average efficiency of currently-used hardware.

One way to get around this problem is to assume that miners’ electricity costs are proportional to total mining revenue (which is observable). The website Digiconomist uses the following equations to derive an estimate:

$$\text{Total Mining Cost} = \text{Total Mining Revenue}. \quad (11)$$

$$\text{Total Cost of Electricity} = (\text{Total Mining Cost}) \times (\% \text{ of Cost From Electricity}). \quad (12)$$

$$\text{Total Electricity Used} = (\text{Total Cost of Electricity}) / (\text{Average Price per KWh}). \quad (13)$$

DeVries (2018) explains the first equation with the comment “in equilibrium, not even Bitmain... should be able to generate a profit.” In fact, this relies on faulty logic: Theorem 1 implies that all active miners should make a (possibly small) profit.

In fact, there is a close relationship between the level of concentration in the mining industry, and total mining profit. We measure concentration using the Herfindahl-Hirschman

Index (HHI), defined as the sum of squares of market shares:²⁰

$$HHI(q) = \sum_i x_i(q)^2. \tag{14}$$

By (10), we have the following result.

Corollary 3. *If Q is a second-stage equilibrium, then for every p , the proportion of total mining revenue that constitutes mining profit is equal to the Herfindahl-Hirschman Index of market concentration:*

$$\sum_i \Pi_i^M(p, Q(p)) = R \cdot HHI(Q(p)).$$

By assuming that miners make zero profit, Digiconomist likely overestimates the money that miners spend on electricity. A simple first-order correction would be to modify (11) by multiplying mining revenue by one minus the market concentration.²¹ However, this correction ignores the fact that the miners with the greatest market share are also those with the lowest costs, implying that using the *average* electricity cost in (13) will overestimate electricity consumption. This point is related to the critique of Koomey (2019) that “these estimates... don’t deal with the large geographic variations in electricity prices.” Indeed, Corollary 2 suggests that miners with lower electricity costs will constitute a noticeably higher market share.

While we do not claim to offer a reliable estimate of the electricity use of the Bitcoin network, our model provides several reasons to believe that the simple calculations offered by Digiconomist are an over-estimate. A more careful economics-based analysis would take into account both mining profit and the correlation between electricity cost and market share.

²⁰See https://en.wikipedia.org/wiki/Herfindahl%E2%80%93Hirschman_Index for more information. Rhoades (1993) describes how HHI has been used by the Department of Justice and the Federal Reserve in the analysis of the competitive effects of mergers. In the industrial organization literature, it is common to treat each company’s market share as an integer percentage (i.e. 20 instead of 0.2). This results in an HHI that ranges from 0 to 10,000. In this paper, we define HHI by (14), so that it ranges from 0 to 1.

²¹One limitation is that estimating the concentration of hardware ownership is challenging. Blogs occasionally compute HHI for mining pools, as this data is more readily available. However, as discussed in Section 1, this is potentially quite different from the HHI of hardware ownership.

5 The First-Stage Game: Concentration of Production

We now fully characterize equilibria in our model, and use this characterization to answer the question, *Why is the production of Bitcoin mining hardware so concentrated?*

The answer, in some ways, is simple. In the first-stage game, miners essentially engage in Bertrand competition: the lowest price claims the entire market. Theorem 2 establishes that there is always an equilibrium in which the miner with the lowest-cost hardware sells to the rest of the market. Below, we say that *sales occur* if any miner sells hardware to another miner (i.e. there exists an $i \neq j$ such that $Q_{ij}(p) > 0$).

Theorem 2. *There exists an equilibrium (p, Q) where sales occur. If $HC_1 < HC_2$, then in all equilibria (p, Q) where sales occur, Miner 1 sets price $p_1 \in [HC_1, HC_2]$, and all hardware is produced by Miner 1 ($Q_{ij}(p) = 0$ for all i , and all $j \geq 2$).*

Although Theorem 2 relies on familiar intuition, our model is significantly richer than standard Bertrand competition, because sellers have the opportunity to mine themselves. Thus, a priori, there are three potential equilibrium outcomes.

- I. Miner 1 both mines and sells hardware to competitors.
- II. Miner 1 sells hardware to competitors, but does not mine.
- III. Miner 1 mines, but does not sell hardware to competitors (and therefore, by Theorem 2, no sales occur).

Theorem 2 establishes that at least one of the first two outcomes can be sustained in equilibrium, but which one? Will the first miner choose to mine, or make its profit entirely from selling hardware? What price will it charge in equilibrium? Theorem 2 also doesn't rule out the existence of other equilibria: can there be an equilibrium where miners with efficient hardware choose not to sell this hardware to others?

We answer these questions below. First, Section 5.1 characterizes the price set by Miner 1 in equilibria with sales. While it is possible that either I or II is an equilibrium, in both cases Miner 1 sets the highest price consistent with its participation decision and the need to undercut Miner 2. Section 5.2 establishes that it is possible for an equilibrium with no sales to exist, but only if the market primitives satisfy stringent necessary conditions.

5.1 Characterizing the Dominant Miner’s Behavior

We first characterize the price set by Miner 1 in equilibria where sales occur, and understand whether they choose to mine. Theorem 3 establishes that there are only two possible prices Miner 1 might set in equilibrium. In particular, if Miner 1 both mines and sells hardware, then it will charge the highest possible price HC_2 . Meanwhile, if Miner 1 abstains from mining, then it either charges HC_2 , or the highest price that maintains its incentive not to mine.

Theorem 3. *In all equilibria (p, Q) where sales occur, either: $p_1 = HC_2$, or: $p_1 < HC_2$, $c^*(p) = HC_1 + OC_1$, and Miner 1 does not mine ($Q_{1j}(p) = 0$ for all j).*

That is, there are only (at most) two potential prices for Miner 1 to check: HC_2 , and whatever price induces $c^*(p) = HC_1 + OC_1$. Generically, only one of these will be an equilibrium (because Miner 1 chooses their preferred price in $[HC_1, HC_2]$).²²

Theorem 3 implies that a necessary condition for any price $p_1 < HC_2$ to be an equilibrium is that Miner 1’s decision of whether to mine hinges on the price it sets. That is, if we express c^* as a function of p_1 , we must have

$$c^*(HC_1) \leq HC_1 + OC_1 < c^*(HC_2).$$

If OC_1 is too small, then no matter what price Miner 1 sets, they will choose to mine (and Theorem 3 implies that they therefore set price HC_2). If OC_1 is too large, then no matter what price Miner 1 sets (subject to guaranteeing themselves the entire market of sales), they will choose not to mine (and Theorem 3 implies they will also set price HC_2). But if the above inequalities hold, then by Lemma 2 there is a unique p_1 such that $c^*(p) = HC_1 + OC_1$, and this will be an equilibrium if and only if this price is at least as profitable as selling at HC_2 .

Interestingly, in the cases where $p_1 < HC_2$ is an equilibrium, Miner 1 would actually earn higher profit by setting a price of HC_2 and *committing not to mine*. Our model does not permit this possibility: in practice mining activity is not directly observable, and any commitment by a manufacturer not to use their own hardware would not be credible. In our model, this commitment can be made credible by lowering the sales price, thereby encouraging other miners to purchase more equipment and reducing Miner 1’s return from

²²Non-generic costs may result in the existence of multiple equilibria where sales occur – see Example 2 in Appendix A.

mining to the point where it is no longer profitable. This is precisely the mechanism that can cause Miner 1 to set a low price.

5.2 When must sales occur?

Recall our second key question: *Does it make sense for the dominant miner to both mine and sell hardware to its competitors?* Theorem 2 suggests that the answer is ‘yes’, because there is an equilibrium where sales occur (and in all such equilibria, all mining hardware is purchased from Miner 1). But this answer is incomplete, as there could a priori be another equilibrium where no sales occur. We now turn to the question of whether Miner 1 will ever choose to keep their hardware to themselves, in order to control more of the mining market. The following example illustrates that this is possible.

Example 1 (Equilibrium Without Sales). *Let $HC = (1/8, 1/4, 1)$ and $OC = (3/8, 3/4, 1)$.*

Because Miner 1 has the lowest hardware and lowest operational costs, Theorem 1 implies that it will always choose to mine. Theorem 2 states that there is an equilibrium where sales occur, and Theorem 3 implies that in this equilibrium, Miner 1 must set a price of $HC_2 = 1/4$.

When Miner 1 sets a price of $HC_2 = 1/4$, effective costs are $\vec{c} = (1/2, 1, 5/4)$ and $c^ = 11/8$. By Theorem 1, Miner 1 has a market share of $7/11$, mining profit of $(7/11)^2$, and profit from sales of $8/11 \cdot 4/11 \cdot (1/4 - 1/8) = 4/121$, for a total profit of $53/121 \approx 0.438$. Miner 2 has market share $3/11$ and profit $9/121$, and miner 3 has market share $1/11$ and profit $1/121$.*

Without sales, effective costs are $\vec{c} = (1/2, 1, 2)$, and $c^ = 3/2$. By Theorem 1, Miner 1 has a market share of $2/3$ and profit of $4/9 \approx 0.444$, Miner 2 has market share of $1/3$ and profit of $1/9$.*

The calculations above establish that Miner 1 prefers the “no sales” outcome to the equilibrium in which sales occur at price $1/4$. However, this is not sufficient to ensure that “no sales” is an equilibrium. We must instead verify that even at a high price, neither Miner 1 nor Miner 2 wishes to sell to Miner 3. The calculations for this are presented in Appendix A.2.²³

²³Although Example 1 has an equilibrium without sales, there are situations where the unique equilibrium is for Miner 1 to sell, even though all active miners prefer the “no sales” outcome to the equilibrium outcome. This suggests the possibility of collusion among miners with low mining costs, which we do not model.

So, our model is rich enough to admit the possibility of a no-sales equilibrium. However, Theorem 4 establishes fairly stringent necessary conditions on the market primitives in order for the no-sales outcome to be an equilibrium. Below, we say that miner i is **active** if $q_{ii} > 0$ in the no-sales outcome and **inactive** if $q_{ii} = 0$ (recall that Theorem 1 establishes that miner i will be active if and only if its effective cost $HC_i + OC_i$ is less than the break-even cost c^*).

Theorem 4. *Suppose that the no-sales outcome is an equilibrium. Then:*

- I. *Every inactive miner has weakly higher hardware cost than every active miner.*
- II. *If there are at least three active miners, then every inactive miner has strictly higher operational cost than every active miner.*
- III. *Unless all active miners have identical hardware costs, there is an inactive miner j who could profitably mine using hardware from **any** active miner i : $OC_j < c^* - HC_i$.*

If any of these conditions are violated, then sales occur in all equilibria. These conditions come from considering the following potential deviations from the no-sales outcome:

- I. An inactive miner could sell hardware to an active miner.
- II. An active miner could sell hardware to an inactive miner.
- III. An active miner could sell hardware to another active miner.

The first deviation will be profitable if any inactive miner has lower hardware costs than any active miner. The second deviation will be profitable if any inactive miner has low operational costs (since then, an active miner could sell them hardware at a high markup — this is the most technical bullet in the proof of Theorem 4).²⁴ The third deviation is guaranteed to be profitable if every inactive miner has high operational costs (since then, Miner 1 could sell to other active miners without attracting new entrants).

The conditions in Theorem 4 are restrictive for two reasons. First, observe that active miners are exactly those with $OC_i + HC_i < c^*$, so every active miner has lower total cost than every inactive miner. Bullets one and two of Theorem 4 imply that if no-sales is an

²⁴Theorem 4 establishes that if there are at least three active miners, then $OC_j \leq OC_i$ is sufficient for some active miner to want to sell to inactive miner j . Appendix A.3 provides an example demonstrating that this condition is not sufficient if there are only two active miners. However, it remains true that an inactive miner with sufficiently low operational costs will tempt an active miner to deviate.

equilibrium, then every active miner not only has lower total cost, but lower hardware cost *and* lower operational cost. In particular, for there to be an equilibrium without sales, it is necessary that every miner with low operational costs *can also produce hardware cheaply enough to mine profitably*. This seems unlikely.

Second, observe a quantitative tension between bullets two and three. Bullet two requires that every inactive miner have high operational cost (strictly higher than every active miner). Yet, bullet three requires the existence of an inactive miner with *low enough* operational cost to purchase hardware from every active miner. So in particular, the lowest operational cost of any inactive miner must lie in $(\max_{\text{active } i}\{\text{OC}_i\}, c^* - \max_{\text{active } i}\{\text{HC}_i\})$.²⁵

In summary, Theorem 4 establishes stringent necessary conditions for no-sales to be an equilibrium, suggesting that for practical market primitives sales occur in all equilibria. This conclusion is only strengthened if miners are permitted to sell at personalized prices. In that case, it will always be profitable for Miner 1 to deviate from the “no sales” outcome by selling to other active miners.

6 Conclusions and Discussion

The concentration of ownership of mining hardware among a few large entities, and the concentration of production of mining hardware within a single large entity, threatens the promise of a truly decentralized digital currency. We show that both aspects of concentration can be explained using a simple model with three key features. First, miners share a fixed reward in proportion to their investment in computational power. Second, miners have different costs, separated into hardware costs and operational costs. Third, hardware can be readily bought and sold.

We show that seemingly minor differences in operational costs will result in concentrated ownership of mining hardware: low-cost miners will have significant market share, and earn significant profits. Our conclusion that mining profits are very sensitive to mining cost also justifies our assumption that miners will purchase the hardware that lets them mine most cheaply. As a result, the sale of mining hardware resembles a Bertrand competition, where the lowest-cost manufacturer captures the entire hardware market.

In summary, mining centralization arises from core aspects of the Bitcoin mining pro-

²⁵In particular, observe that this range is not even necessarily non-empty. In order for the range to be non-empty, it must be that the worst hardware among active miners can be profitably used by the active miner with the highest operational costs.

tol, and is not a temporary aberration. Without significant changes, the vision of a competitive mining market is unlikely to be fulfilled. As Section 4.3 discusses, analyses based on the assumption of a competitive mining market should be revisited, to understand how the reality of a profitable mining oligopoly affects their conclusions.

There are several proposals for ways to reduce mining centralization. When evaluating these proposals, the simplicity of our model is an advantage: unless the change addresses one of the three features identified above, our model suggests that centralization will persist. For example, innovations such as BetterHash and Stratum V2,²⁶ which are designed to allow miners (rather than mining pools) to dictate the contents of a block, do not address the cost asymmetries that are the focus of this paper. While these technologies may offer some advantages, they are unlikely to result in decentralized mining. Meanwhile, a change to ASIC-resistant hash functions (such as Equihash, which is used by Bitcoin Gold) might change the supplier of mining hardware, but so long as cooling and electricity remain significant costs, mining activity is likely to remain concentrated among the few areas of the world where these costs are cheapest.

What changes might have an impact? One thought is to change the reward structure, so that rewards scale sublinearly with mining power: diseconomies of scale directly disincentivizes large miners. Unfortunately, sublinear rewards seem impossible so long as mining is permissionless: recent works point out that a proportional division is the only one that is *anonymous*, *robust to sybil attacks*, and *robust to mergers* (Chen et al., 2019; Leshno and Strack, 2019). The key challenge is that any effort to make rewards sublinear in computational power will result in miners dividing their power among multiple false identities. A more promising approach is to equalize costs across miners. Because storing and transporting electricity is difficult, cost asymmetries seem inherent to any proof-of-work protocol. However, proof-of-stake protocols reward miners in proportion to the amount of currency that they own, rather than their computational power. It seems plausible that variation in the cost of purchasing cryptocurrency should be much smaller than variation in electricity prices. If this is the case, then proof-of-stake protocols might contribute to mining decentralization.

Although the simplicity of our model is one of its key advantages, it also implies certain limitations. For example, we model both operational and hardware costs as linear in com-

²⁶For more information, see <https://medium.com/hackernoon/betterhash-decentralizing-bitcoin-mining-with-new-hashing-protocols-291de178e3e0> and <https://www.coindesk.com/a-plan-to-decentralize-bitcoin-mining-again-is-gaining-ground>.

putational power. In practice, there are reasons to believe that mining features economies of scale.²⁷ However, this fact should only serve to exacerbate concentration of ownership: it is telling that this concentration arises even when economies of scale are absent.

More importantly, we consider a static model, and combine hardware costs and operational costs using a simple additive form. This captures the fact that miners purchasing identical hardware (at identical prices) may nevertheless have different operational costs. However, our model does not clearly distinguish between the one-time cost of purchasing hardware, and the ongoing cost of powering it. In practice, fluctuations in Bitcoin and electricity prices imply the source of the cost matters: if powering equipment is the major cost, then less efficient miners should use their equipment only when Bitcoin prices are high, or electricity prices low. If most of the cost is hardware acquisition, then miners will power their hardware regardless of the Bitcoin price. While modeling and understanding these dynamics is a viable direction for future work, the basic insight that small cost asymmetries can result in highly centralized mining operations should persist.

References

- H. Alsabah and A. Capponi. Pitfalls of bitcoin’s proof of work: R&d arms race and mining centralization. 2020.
- M. Babaioff, S. Dobzinski, S. Oren, and A. Zohar. On bitcoin and red balloons. In *ACM Conference on Electronic Commerce (EC)*, 2012.
- E. Budish. The economic limits of bitcoin and the blockchain. 2018.
- M. Carlsten, H. A. Kalodner, S. M. Weinberg, and A. Narayanan. On the instability of bitcoin without the block reward. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, October 24-28, 2016*, pages 154–167, 2016.
- X. Chen, C. H. Papadimitriou, and T. Roughgarden. An axiomatic approach to block rewards. In *Proceedings of the 1st ACM Conference on Advances in Financial Technologies, AFT 2019, Zurich, Switzerland, October 21-23, 2019*, pages 124–131, 2019.

²⁷For example, there are large fixed costs associated with setting up a data center, and large miners may be able to negotiate bulk discounts from their suppliers. Furthermore, large miners can slightly increase their share of rewards by deviating from the mining protocol (see e.g. Eyal and Sirer (2014); Sapirshstein et al. (2016); Kiayias et al. (2016); Carlsten et al. (2016)).

- J. Chiu and T. Koepl. The economics of cryptocurrencies – bitcoin and beyond. 2018.
- L. W. Cong, Z. He, and J. Li. Decentralized mining in centralized pools. 2018.
- P. DeGraba. Input market price discrimination and the choice of technology. *American Economic Review*, 90(5):1246–1253, 1990.
- A. DeVries. Bitcoin’s growing energy problem. *Joule*, 2:801–809, 2018.
- N. Dimitri. Bitcoin mining as a contest. *Ledger*, 2017.
- D. Easley, M. O’Hara, and S. Basu. From mining to markets: the evolution of bitcoin transaction fees. *Journal of Financial Economics*, 134:91–109, 2019.
- I. Eyal. The miner’s dilemma. In *Security and Privacy (SP), 2015 IEEE Symposium on*, pages 89–103. IEEE, 2015.
- I. Eyal and E. G. Sirer. Majority is not enough: Bitcoin mining is vulnerable. In *Financial Cryptography and Data Security*, pages 436–454. Springer, 2014.
- A. E. Gencer, S. Basu, I. Eyal, R. V. Renesse, and E. G. Sirer. Decentralization in bitcoin and ethereum networks. In *Financial Cryptography and Data Security (FC)*, 2018.
- M. Gradstein. Intensity of competition, entry and entry deterrence in rent seeking contests. *Economics & Politics*, 7(1):79–91, 1995.
- H. Halaburda and G. Haeringer. Bitcoin and blockchain: What we know and what questions are still open. 2019.
- A. Hayes. A cost of production model for bitcoin. 2015.
- G. Hileman and M. Rauchs. Global cryptocurrency benchmarking study, 2017. https://www.jbs.cam.ac.uk/fileadmin/user_upload/research/centres/alternative-finance/downloads/2017-global-cryptocurrency-benchmarking-study.pdf.
- A. L. Hillman and J. G. Riley. Politically contestable rents and transfers. *Economics & Politics*, 1(1):17–39, 1989.
- G. Huberman, J. Leshno, and C. Moallemi. Monopoly without a monopolist: An economic analysis of the bitcoin payment system. 2017.

- M. Katz. The welfare effects of third-degree price discrimination in intermediate good markets. *American Economic Review*, 77(1):154–167, 1987.
- A. Kiayias, E. Koutsoupias, M. Kyropoulou, and Y. Tselekounis. Blockchain mining games. In *ACM Conference on Economics and Computation (EC)*, 2016.
- J. Koomey. Estimating bitcoin electricity use: A beginner’s guide. Coin Center Report, May 2019.
- J. A. Kroll, I. C. Davey, and E. W. Felten. The economics of bitcoin mining, or bitcoin in the presence of adversaries. In *Workshop on the Economics of Information Security (WEIS)*, 2013.
- J. Leshno and P. Strack. Bitcoin: An impossibility theorem for proof-of-work based protocols. 2019.
- J. Ma, J. Gans, and R. Tourky. Market structure in bitcoin mining. NBER Working Paper, 2018.
- S. Nakamoto. Bitcoin: A peer-to-peer electronic cash system. 2008.
- A. Narayanan, J. Bonneau, E. Felten, A. Miller, and S. Goldfeder. *Bitcoin and Cryptocurrency Technologies*. Princeton University Press, 2016.
- J. Prat and B. Walter. An equilibrium model of the an equilibrium model of the market for bitcoin mining. CESifo Working Paper, 2018.
- S. A. Rhoades. The herfindahl-hirschman index. *Fed. Res. Bull.*, 79, 1993.
- M. Romiti, A. Judmayer, A. Zamyatin, and B. Haslhofer. A deep dive into bitcoin mining pools: An empirical analysis of mining shares. In *Workshop on the Economics of Information Security (WEIS)*, 2019.
- A. Sapirshtein, Y. Sompolinsky, and A. Zohar. Optimal selfish mining strategies in bitcoin. In *Financial Cryptography and Data Security*, 2016.
- M. Thum. The economic cost of bitcoin mining. *CESifo Forum*, 19, March 2018.
- G. Tullock. *Efficient rent-seeking*. 1980.
- Y. Yehezkel. Downstream competition between an upstream supplier and an independent downstream firm. 2004.

A Examples

A.1 Degenerate Example with Multiple Sales Equilibria

In Section 5.1, we claimed that generically, there is a unique equilibrium where sales occur. For completeness, we give a non-generic example where the two possible equilibrium prices are both optimal. To keep calculations simple, we set operational costs equal to zero for some miners. A similar example in which Miner 1's profit is maximized at exactly two prices could be constructed by letting miners have small positive operational costs.

Example 2 (Multiple Sales Equilibria). *There are 4 miners. Miners 2, 3, 4 are identical, with $HC_i = 1$ and $OC_i = 0$. Miner 1 has $HC_1 = 0$.*

If $OC_1 < 0.6$, then there is a unique equilibrium: Miner 1 sells at price 1, $c^ = 1 + \frac{1}{3}OC_1$, and Miner 1 gets a profit of $\frac{1 - \frac{1}{3}OC_1 + \frac{4}{9}OC_1^2}{1 + \frac{2}{3}OC_1 + \frac{1}{9}OC_1^2}$.*

If $OC_1 > 0.6$, then for any price $p \leq \frac{2}{3}OC_1$, Miner 1 stays out of the market, and earns profit $2/3$. This is better than setting a price of $p = 1$. There are a continuum of equilibria, but all are payoff equivalent. However, there is no “lowest optimal price.”

If $OC_1 = 0.6$, then there are multiple equilibria that are not payoff equivalent ($p = 1, p \leq 0.4$). Miner 1 gets a utility of $2/3$ in each. When $p = 1$ we have $c^ = 6/5$ and each other miner gets market share $1/6$ and utility $1/36$. When $p \leq 0.4$ we have $c^* = \frac{3}{2}p$, and each of miners 2, 3, 4 get utility $1/9$.*

A.2 Omitted Calculations for Example 1

Recall that Example 1 provides an example where “no sales” is an equilibrium. We verify these calculations below.

Lemma 3. *In Example 1, “no sales” is an equilibrium.*

Proof. Consider the price-setting problem facing Miner 1. For $p \geq 1/2$, no sales occur. For $p \in (1/4, 1/2)$, the third miner buys hardware and we have $c^*(p) = 5/4 + p/2$. Miner 1's profits are

$$\left(\frac{c^* - 1/2}{c^*}\right)^2 + \frac{1}{c^*} \frac{c^* - (1 + p)}{c^*} \left(p - \frac{1}{8}\right) = \frac{34p - 4p^2 + 17}{2(5 + 2p)^2}.$$

For $p \in [1/8, 1/4]$, the second and third miners buy hardware and we have $c^*(p) = 9/8 + p$.

Miner 1's profits are

$$\left(\frac{c^* - 1/2}{c^*}\right)^2 + \frac{1}{c^*} \frac{1/2}{c^*} \left(p - \frac{1}{8}\right) = \frac{64p^2 + 112p + 21}{(9 + 8p)^2}.$$

We can verify that both profit functions are maximized at the right end points, and that no price is better for Miner 1 than a no sales price of $p = 1/2$.

We can also check that Miner 2 does not want to sell. As before, at a price $p \in (1/4, 1/2)$, the third miner buys hardware and we have $c^*(p) = 5/4 + p/2$. Miner 2's profits (including revenue) are

$$\left(\frac{c^* - 1}{c^*}\right)^2 + \frac{1}{c^*} \frac{c^* - (1 + p)}{c^*} \left(p - \frac{1}{4}\right) = \frac{10p - 4p^2}{(5 + 2p)^2}.$$

These are maximized at the no sales outcome $p = 1/2$. □

A.3 Necessity of Technical Conditions in Theorem 4

In Theorem 4, bullet two requires at least three active miners. We quickly confirm that this condition is necessary, by showing that it does not hold with only two active miners. As previously noted, even when only two miners are active, an inactive miner with sufficiently low operational costs will still tempt an active miner to deviate. However, “sufficiently low” will no longer have the clean definition given in Theorem 4.

Example 3 (No Sales with Non-Dominance). *Let hardware costs be $(0.25, 0.3, 1)$ and operational costs be $(0.75, 0.2, 0.6)$.*

With no sales, effective costs are $\vec{c} = (1, 0.5, 1.6)$, and $c^ = 1.5$. By Theorem 1, Miner 1 has market share $1/3$ and profit $1/9$. Miner 2 has market share $2/3$ and profit $4/9$.*

With sales at a price of 0.25 , effective costs are $\vec{c} = (1, 0.5, 0.9)$, and $c^ = 1.2$. By Theorem 1, Mining profits are $(4/144, 49/144, 9/144)$, and sales profit for Miner 1 is $5/6 \cdot 5/6 \cdot 1/20$. Total profit is $(9/144, 49/144, 9/144)$.*

Miner 1 prefers the outcome with no sales. Furthermore, one can verify that this is an equilibrium (even though Miner 3 has lower operational costs than Miner 1).

B Proofs: Section 4

We first prove Lemma 2, then a helper lemma, and then provide a proof of Theorem 1.

Proof of Lemma 2. For $y \leq \min_i c_i$, we have $X(y, \vec{c}) = 0$. Note that $X(y, \vec{c})$ is continuous and strictly increasing in y on $[\min_i c_i, \infty)$, and tends to n as $y \rightarrow \infty$. Therefore, there is a unique solution to $X(c^*, \vec{c}) = 1$. Furthermore, increasing any $c_i < y$ strictly decreases $X(y, c_i)$, while increasing any $c_i \geq y$ leaves $X(y, c_i)$ unchanged. \square

Given any second-stage response Q , we let

$$Q_i(p) = \sum_j Q_{ij}(p) \quad (15)$$

be the total computational power owned by miner i .

Lemma 4. *Suppose that Q is a second-stage equilibrium. Then for each i and p ,*

$$\sum_{j \neq i} Q_j(p) > 0. \quad (16)$$

$$x_i(Q(p)) = \max\left\{1 - \frac{c_i(p)}{R} \sum_j Q_j(p), 0\right\}. \quad (17)$$

Proof. We fix p and the values $Q_j(p)$ for $j \neq i$, and consider the optimization problem facing miner i . To simplify notation, we drop the dependence of Q_i, Q_j , and c_i on p , and we write miner i 's market share and profit as functions of q_i :

$$\begin{aligned} x_i(q_i) &= \frac{q_i}{q_i + \sum_{j \neq i} Q_j} \\ U_i(q_i) &= R \cdot x_i(q_i) - c_i q_i + \Pi_i^S \end{aligned} \quad (18)$$

(Note that Miner i 's profit from sales Π_i^S does not depend on q_i .)

We first claim that $\sum_{j \neq i} Q_j > 0$. This is because if $\sum_{j \neq i} Q_j = 0$, then $x_i = I(q_i > 0)$, and miner i does not have a best response (so in particular, Q_i can not be a best response).

Next, note that if $\sum_{j \neq i} Q_j > 0$, then x_i is continuous and differentiable in q_i , with

$$x'_i(q_i) = \frac{1}{q_i + \sum_{j \neq i} Q_j} - \frac{q_i}{(q_i + \sum_{j \neq i} Q_j)^2} = \frac{1 - x_i(q_i)}{q_i + \sum_{j \neq i} Q_j}. \quad (19)$$

Furthermore, x_i is easily shown to be concave, from which it follows that U is concave. Because i chooses $q_i \in [0, \infty)$, it follows that the maximizer of U satisfies the following

first-order optimality condition:

$$U'_i(Q_i) \leq 0, \quad \text{with equality if } Q_i > 0. \quad (20)$$

We will show that (18), (19) and (20) jointly imply (17). They imply that for all i ,

$$U'_i(Q_i) = R \cdot x'_i(Q_i) - c_i = R \cdot \frac{1 - x_i(Q_i)}{Q_i + \sum_{j \neq i} Q_j} - c_i \leq 0, \quad \text{with equality if } Q_i > 0.$$

Therefore if $Q_i > 0$, we must have $x_i(Q_i) = 1 - \frac{c_i}{R} \sum_j Q_j$. Meanwhile, if $Q_i = 0$, then $x_i(Q_i) = 0$ and we must have $R - c_i \sum_j Q_j \leq 0$. In other words, (17) holds. \square

Proof of Theorem 1. If Q is a second stage equilibrium, then for all p we have

$$1 = \sum_i x_i(Q(p)) = \sum_i \max\left(1 - \frac{c_i(p)}{R} \sum_j Q_j(p), 0\right) = X(R / \sum_j Q_j(p)).$$

The first equality follows from the definition of x_i in (17), the second from Lemma 4, and the third from the definition of X in (7). Therefore, Lemma 2 implies that

$$c^*(p) = \frac{R}{\sum_j Q_j(p)}.$$

In other words, (8) holds. From this and Lemma 4, (9) follows: $x_i(Q(p)) = \max\left(1 - \frac{c_i(p)}{c^*(p)}, 0\right)$.

Finally, the definition of Π_i^M in (10) implies that

$$\begin{aligned} \Pi_i^M(p, Q(p)) &= R \cdot x_i(Q(p)) - \sum_j Q_{ij}(p) c_{ij}(p) \\ &= R \cdot x_i(Q(p)) - Q_i(p) c_i(p) \\ &= R \cdot x_i(Q(p)) - R \cdot x_i(Q(p)) \frac{\sum_j Q_j(p)}{R} c_i(p) \\ &= R \cdot x_i(Q(p)) \left(1 - \frac{c_i(p)}{c^*(p)}\right) \\ &= R \cdot x_i(Q(p))^2. \end{aligned}$$

The second line uses Lemma 1 (miners only source from the lowest cost provider), the third uses the definition of x_i in (17), the fourth uses (8), and the fifth uses Lemma 4. \square

We now make two observations that will prove useful in subsequent analysis.

Lemma 5. *If Q is a second-stage equilibrium, then Q_i is a continuous function of p for all i . As a result, so are $x_i(Q(p))$ and $\Pi_i^M(Q(p))$.*

Lemma 6. *For any c_1, \dots, c_n , if we define c^* as in Lemma 2 and let $k = \sum_i \mathbf{1}(c_i < c^*)$ be the number of miners participating in equilibrium, then $(k - 1)c^* = \sum_{i=1}^k c_i$.*

This follows from rewriting the equation $\sum_{i=1}^k 1 - c_i/c^* = 1$.

C Proofs: Concentration of Production

Throughout this section, we use $Q_j(p)$ as shorthand for $\sum_i Q_{ji}(p)$, and define

$$\underline{p} = \underline{p}(p) := \min_j \{p_j\}.$$

Observe that this implies that $c_j(p) = \text{OC}_j + \min\{\underline{p}(p), \text{HC}_j\}$ for all j .

C.1 Proof of Theorem 2

Lemma 7. *Let Q be a second-stage equilibrium. If there exists a miner i with $\text{HC}_i < \underline{p}(p)$, and if any miner purchases from a miner other than i in $Q(p)$ (that is, $Q_{jk} > 0$ for $k \notin \{i, j\}$), then (p, Q) is not an equilibrium.*

Proof. We claim that miner i can choose a better price than p_i . We note that by the definition of U_i in (5) and Theorem 1, we have

$$U_i(p, Q(p)) = R \cdot \max \left(1 - \frac{c_i(p)}{c^*(p)}, 0 \right)^2 + (p_i - \text{HC}_i) \sum_{j \neq i} Q_{ji}(p). \quad (21)$$

If $p_i > \underline{p}$, then by Lemma 1, $Q_{ji}(p) = 0$ for all $j \neq i$, so the second term in (21) is zero. Because lowering p_i to \underline{p} does not change $c_i(p)$ or $c^*(p)$, it does not change the first term in (21), and may increase the second. Thus, utility U_i is weakly higher than before.

We now show that if the conditions of the Lemma are met, then for sufficiently small ε , it is strictly better for i to set a price of $p_i = \underline{p} - \varepsilon$ than to set a price of \underline{p} . The point is that as miner i lowers its price, the loss in miner i 's market share and profit-per-unit $p_i - \text{HC}_i$ move continuously, whereas quantity sold jumps discretely.

To see this, consider the expression for U_i in (21). Because $c_i(p)$ and $c^*(p)$ are continuous functions of p , the first term (which represents mining profit Π_i^M) changes continuously in ε . Any miner j who was previously making a purchase had $\text{HC}_j \geq \underline{p}$. It follows from Lemma 1 that if miner i were to set a price $p_i < \underline{p}$, then all of these miners would source exclusively from i . We know that at least one of these miners was previously sourcing elsewhere ($Q_{jk}(p) > 0$). Because total purchase quantities $Q_j(p)$ are continuous in p for all j by Lemma 5, for all sufficiently small ε Miner i gets a discrete increase in sales, with only a continuous loss in per-unit profit. Therefore, this deviation is profitable for sufficiently small ε . \square

Proof of Theorem 2. We first claim that if (p, Q) is an equilibrium in which sales occur then the minimum price \underline{p} must be at most HC_2 . This is because if sales occur and the minimum price is above HC_2 , then by Lemma 7, either Miner 1 or Miner 2 is not best responding (either one could set a price of $\underline{p} - \varepsilon$, and capture the entire hardware market).

We next show that if $\text{HC}_1 < \text{HC}_2$, then in every equilibrium (p, Q) in which sales occur, all hardware is sourced from Miner 1 (that is, $Q_{ij}(p) = 0$ for $j \geq 2$). If $p_1 < \text{HC}_2$, then this is immediate from Lemma 1 (which states that all miners purchase from the lowest-cost source) and the fact that Q is a second-stage equilibrium. Meanwhile, if $p_1 \geq \text{HC}_2$ and some hardware is purchased from a miner other than 1, then Lemma 7 establishes that setting a price p_1 just below HC_2 would result in higher utility for Miner 1.

Finally, we establish existence of an equilibrium where sales occur. Let $p_j = \text{HC}_j$ for $j \geq 2$, and let Q be a second-stage equilibrium in which all miners tie-break in favor of Miner 1 when purchasing hardware. Let p_1 be Miner 1's best response to this.

We first note that if Miner 1 sets any price $p_1 \leq \text{HC}_2$, then sales will occur. In particular, this implies that $\text{HC}_j \geq p_1$ and $c_{j1}(p) \geq c_j(p)$ for all $j \neq 1$. Because Q is a second-stage equilibrium in which miners tiebreak in favor of Miner 1, all hardware will be sourced from Miner 1. Furthermore, Miner 1 will have positive sales, because Lemma 4 implies that $Q_j(p) > 0$ for some $j \neq 1$.

Next, we argue that any price above HC_2 is weakly worse than a price of HC_2 . This is because $c_i(p)$ (and therefore $c^*(p)$) are the same for any price $p_1 \geq \text{HC}_2$, but when $p_1 = \text{HC}_2$, Miner 1 gets profit from selling hardware. Therefore, we can restrict attention to prices in the interval $[\text{HC}_1, \text{HC}_2]$.

On this interval, $U_1(p, Q(p))$ is continuous in p_1 : as argued above, all hardware is purchased from Miner 1, and by Lemma 5, the quantities $c^*(p)$ and $Q_{j1}(p)$ are continuous

in p_1 . Therefore, Mining profit Π_i^M and sales profit Π_i^S are both continuous in p_1 on $[\text{HC}_1, \text{HC}_2]$, so $U_1(p, Q)$ achieves its supremum over this compact interval. In other words, there is a best response p_1 . Furthermore, because $p_1 \leq \text{HC}_2$, no miner $j \geq 2$ can benefit from changing its price. This establishes the existence of an equilibrium in which sales occur. \square

C.2 Proof of Theorem 3

Proof of Theorem 3. Theorem 2 already establishes that when sales occur, $p_1 \in [\text{HC}_1, \text{HC}_2]$. Therefore, the claim of Theorem 3 which needs proving is that whenever $p_1 < \text{HC}_2$, $c^*(p) = \text{HC}_1 + \text{OC}_1$. Observe that in order to possibly have $p_1 < \text{HC}_2$, it must be that $\text{HC}_1 < \text{HC}_2$. Therefore, Theorem 2 already establishes that all hardware is purchased from Miner 1. Because all hardware is purchased from Miner 1, this means that Miner 1 is active if and only if $c^*(p) > \text{HC}_1 + \text{OC}_1$. So we consider the cases $c^*(p) > \text{HC}_1 + \text{OC}_1$ and $c^*(p) \leq \text{HC}_1 + \text{OC}_1$ separately.

We will use throughout both cases that all hardware is purchased from Miner 1 in any potential equilibrium with $p_1 < \text{HC}_2$.

We consider potential equilibria with $c^*(p) \leq \text{HC}_1 + \text{OC}_1$ first. We know from Theorem 1 that the total computational power purchased is $1/c^*(p)$. Because all hardware is purchased from Miner 1 (and Miner 1 is inactive), Miner 1's profit is $(p_1 - \text{HC}_1)/c^*(p)$.

We now expand Miner 1's profit, and establish that it is increasing in $c^*(p)$. Let $A(p)$ denote the set of miners i for which $\text{OC}_i + p_1 < c^*(p)$ (i.e., the miners who are active at p), and $N(p) := |A(p)|$. Then:

$$\begin{aligned}
& \sum_{i \in A(p)} 1 - \frac{\text{OC}_i + p_1}{c^*(p)} = 1 \\
\Rightarrow & \sum_{i \in A(p)} (\text{OC}_i + p_1) = (N(p) - 1) \cdot c^*(p) \\
\Rightarrow & \frac{N(p) \cdot p_1}{c^*(p)} = N(p) - 1 - \frac{\sum_{i \in A(p)} \text{OC}_i}{c^*(p)} \\
\Rightarrow & \frac{p_1}{c^*(p)} = 1 - \frac{1}{N(p)} - \frac{\sum_{i \in A(p)} \text{OC}_i}{N(p) \cdot c^*(p)} \\
\Rightarrow & \frac{p_1 - \text{HC}_1}{c^*(p)} = 1 - \frac{1}{N(p)} - \frac{\sum_{i \in A(p)} \text{OC}_i}{N(p) \cdot c^*(p)} - \frac{\text{HC}_1}{c^*(p)}
\end{aligned}$$

The first line follows as $1 - \frac{OC_i + p_1}{c^*(p)}$ is Miner i 's fraction of the total purchased power. The rest follow from algebraic manipulation. Fixing $N(p)$, it is easy to see that the final expression on the right-hand side is strictly increasing in $c^*(p)$. This means that Miner 1's profit is maximized with $c^*(p)$ as large as possible, *except for possibly discontinuities in Miner 1's profit when $N(p)$ changes* (e.g. if, for example, Miner 1's profit dropped radically when $c^*(p)$ increased and $N(p)$ decreased discretely). However, Miner 1's profit is continuous in $c^*(p)$. To see this, simply observe that p_1 is continuous in $c^*(p)$, and Miner 1's profit is $\frac{p_1 - HC_1}{c^*(p)}$ (a continuous function of continuous functions in $c^*(p)$).²⁸ Therefore, there are not in fact any discontinuities, and Miner 1's profit must be maximized when $c^*(p)$ is as large as possible.

It remains to consider the case where $c^*(p) > HC_1 + OC_1$. Observe in this case that Miner 1 is active, but still that all hardware is purchased from Miner 1 (as $p_1 < HC_2$). In this case, Miner 1 participates, and by Equation (5) their payoff is:

$$\left(1 - \frac{HC_1 + OC_1}{c^*(p)}\right)^2 + (p_1 - HC_1) \cdot \frac{HC_1 + OC_1}{(c^*(p))^2}$$

We again wish to see how this term behaves for a fixed $A(p)$. To this end, it will be significantly cleaner to write $\gamma := \frac{HC_1 + OC_1}{c^*(p)}$, and optimize with respect to γ . We write $p_1(\gamma)$ to denote the price Miner 1 must set to induce a particular γ (noting that p_1 determines c^* which determines γ), and $\Pi(\gamma)$ to denote Miner 1's profit for a particular γ . Observe first that we can write Miner 1's profit as:

$$\begin{aligned} \Pi(\gamma) &= (1 - \gamma)^2 + \frac{(p_1(\gamma) - HC_1) \cdot (HC_1 + OC_1)}{(c^*(p))^2} \\ &= (1 - \gamma)^2 + \frac{\gamma^2(p_1(\gamma) - HC_1)}{HC_1 + OC_1} \end{aligned}$$

Differentiating with respect to γ , we get:

$$\Pi'(\gamma) = -2(1 - \gamma) + 2\gamma \frac{p'(\gamma) - HC_1}{HC_1 + OC_1} + \gamma^2 \frac{p'(\gamma)}{HC_1 + OC_1}$$

²⁸If desired, one can also confirm that when $c^*(p)$ is fixed, but $N(p)$ changes, that the final expression on the right-hand side does not change. To set up the calculations, recall that any miners added/removed from $A(p)$ must have $OC_i + p_1 = c^*(p)$.

We now wish to substitute for $p'(\gamma)$. Observe that (again, fixing $A(p)$):

$$\begin{aligned}
& 1 - \frac{\text{HC}_1 + \text{OC}_1}{c^*(p)} + \sum_{i \in A(p)} 1 - \frac{\text{OC}_i + p_1}{c^*(p)} = 1 \\
\Rightarrow & 1 - \gamma + \sum_{i \in A(p)} 1 - \gamma \frac{\text{OC}_i + p_1}{\text{OC}_1 + \text{HC}_1} = 1 \\
\Rightarrow & -\gamma + N(p) - \gamma \sum_{i \in A(p)} \frac{\text{OC}_i}{\text{OC}_1 + \text{HC}_1} = \gamma N(p) \cdot \frac{p_1(\gamma)}{\text{HC}_1 + \text{OC}_1} \\
\Rightarrow & -\frac{1}{N(p)} + \frac{1}{\gamma} - \sum_{i \in A(p)} \frac{\text{OC}_i}{N(p) \cdot (\text{OC}_1 + \text{HC}_1)} = \frac{p_1(\gamma)}{\text{HC}_1 + \text{OC}_1} \\
\Rightarrow & \frac{p'_1(\gamma)}{\text{HC}_1 + \text{OC}_1} = -1/\gamma^2.
\end{aligned}$$

The first line follows as the sum of total mining power must be 1. The remaining lines substitute the definition of γ and follow from algebraic manipulation. After making this substitution, we get:

$$\begin{aligned}
\Pi'(\gamma) &= -2(1 - \gamma) + 2\gamma \frac{p(\gamma) - \text{HC}_1}{\text{HC}_1 + \text{OC}_1} - 1 \\
&= -3 + 2\gamma \left(\frac{p(\gamma) - \text{HC}_1}{\text{HC}_1 + \text{OC}_1} + 1 \right) \\
&= -3 + 2\gamma \left(\frac{1}{\gamma} - \frac{1}{N(p)} - \sum_{i \in A(p)} \frac{\text{OC}_i}{N(p) \cdot (\text{OC}_1 + \text{HC}_1)} - \frac{\text{HC}_1}{\text{HC}_1 + \text{OC}_1} + 1 \right) \\
&= -1 + 2\gamma \left(-\frac{1}{N(p)} - \sum_{i \in A(p)} \frac{\text{OC}_i}{N(p) \cdot (\text{OC}_1 + \text{HC}_1)} - \frac{\text{HC}_1}{\text{HC}_1 + \text{OC}_1} + 1 \right)
\end{aligned}$$

Observe first that, because $p(\gamma)$ is a continuous function of γ , that $\Pi'(\gamma)$ is also a continuous function of γ (this follows immediately from the first line). From the final line, it is now easy to see for fixed $A(p)$ that $\Pi'(\gamma)$ is either negative for all γ , positive for all γ , or initially negative and then positive.

Because $\Pi'(\gamma)$ is continuous, at the (finitely many) points where $N(p)$ changes, the derivative cannot change signs, and therefore the derivative over the entire range is (weakly)

negative and then (weakly) positive, meaning that the optimum must be achieved at the endpoints. In particular, this means that if the optimal price lies in the interval where $c^*(p) = \text{HC}_1 + \text{OC}_1$ up to HC_2 , it must be either at $c^*(p) = \text{HC}_1 + \text{OC}_1$, or at HC_2 . In particular this means that if $c^*(p) > \text{HC}_1 + \text{OC}_1$ that it must be that $p_1 = \text{HC}_2$.

This completes the proof of both cases: if $c^*(p) \leq \text{HC}_1 + \text{OC}_1$ in the optimal solution, the first half establishes that we must have $c^*(p) = \text{HC}_1 + \text{OC}_1$. If $c^*(p) > \text{HC}_1 + \text{OC}_1$, the second half establishes that we must have $p_1 = \text{HC}_2$. Therefore, if $p_1 < \text{HC}_2$ in equilibrium, we must have $c^*(p) = \text{HC}_1 + \text{OC}_1$. \square

C.3 Proof of Theorem 4

When it is clear from context, we will write $c_i := c_i(p)$ for the effective cost of miner i in the second-stage game, and write $\tilde{c}_i := \text{OC}_i + \text{HC}_i$ to denote the effective cost for player i when all prices are high (at least HC_i). We let $c^* = c^*(p)$ denote the unique solution to $X(c^*, c(p)) = 1$, and \tilde{c}^* denote the unique solution to $X(\tilde{c}^*, \tilde{c}) = 1$.

We prove Theorem 4 through several lemmas. The outline is as follows:

- One simple possible deviation from the “no sales” strategy profile for a inactive miner to sell hardware to a active miner. Under no sales, the inactive miner gets zero utility. But by selling hardware, they get strictly positive utility. Therefore, every active miner must have (weakly) better hardware cost than every inactive miner in order for “no sales” to possibly be an equilibrium (Lemma 8).
- A second simple deviation from “no sales” would be for a active miner to sell hardware to another active miner, without changing c^* . If they can do so *without causing a inactive miner to become active*, this is a strictly better response. Therefore, it must be that any price set by a active miner which is low enough to entice another active miner to buy their hardware *must also* entice a inactive miner to become active (Lemma 9).
- A final deviation from “no sales” would be for a active miner to sell hardware to a previously inactive miner (changing c^* in the process). It turns out that, under the necessary conditions imposed by Lemma 8 and 9, that it is strictly profitable for active Miner i to cause inactive Miner j to become active whenever $\text{OC}_i > \text{OC}_j$ (assuming the necessary conditions implied by the first two bullets), whenever at

least three miners are active in the “no sales” outcome (Lemma 10). This concludes the proof of Theorem 4.

- Example 3 demonstrates that Theorem 4 does not extend to the case where the “no sales” outcome has two active miners. Any necessary and sufficient conditions for this case are much more technical than the simple dominance condition.

We begin now with Lemma 8, establishing that all active miners must have weakly better hardware cost in order for “no sales” to possibly be an equilibrium. Throughout the remainder of this section, we will let $A := \{i | \tilde{c}_i < \tilde{c}^*\}$ denote the set of active miners under “no sales,” and $I := [n] \setminus A$ denote the remaining miners.

Lemma 8. *If “no sales” is an equilibrium, then either I is empty, or $\max_{i \in A} \{HC_i\} \leq \min_{i \in I} \{HC_i\}$.*

Proof. Assume for contradiction that I is non-empty, and that there exists $j \in I, i \in A$ such that $HC_j < HC_i$. Consider the deviation where Miner j sets price $(HC_i + HC_j)/2$. Under “no sales,” Miner j has zero utility. Under this deviation, Miner i ’s effective cost is strictly lowered, implying by Lemma 2 that $c^*(p)$ is strictly lower than \tilde{c}^* . This implies that some miner must purchase hardware from Miner j , at strictly positive profit to Miner j . Therefore, Miner j has a strictly profitable deviation, contradicting that “no sales” is an equilibrium. \square

We continue with Lemma 9, which considers the possibility of a active miner selling hardware to another active miner. Observe that Lemma 9 implies that there is some inactive miner who could profitably mine using *any* active miner’s hardware.

Lemma 9. *If “no sales” is an equilibrium, then at least one of the following occurs:*

- HC_i is the same for all $i \in A$.
- I is non-empty, and $\max_{i \in A} \{HC_i\} + \min_{i \in I} \{OC_i\} < \tilde{c}^*$.

Proof. Assume for contradiction that neither occurs. Then $HC_i \neq HC_j$ for some $i, j \in A$, and also either I is empty, or $\max_{i \in A} \{HC_i\} + \min_{i \in I} \{OC_i\} \geq \tilde{c}^*$.

Let’s first consider the case that $HC_i < HC_j$ for some $i, j \in A$ and I is empty. Then observe that Miner i can set price $\max_{j \in A} \{HC_j\} - \varepsilon$ as a price for arbitrarily small ε . Because c^* is a continuous function of p_i , for all δ , there exists a sufficiently small ε such

that Miner i 's mining profit changes by at most δ by setting this price. On the other hand, the miner $j \in \arg \max_{j \in A} \{HC_j\}$ must now purchase hardware exclusively from Miner i . Again, c_j is a continuous function of p_i so for all δ there is again a sufficiently small ε such that Miner j purchases at least $1 - \tilde{c}_j/\tilde{c}^* - \delta$ units of hardware from Miner i at price at least $HC_j - \delta > HC_i$. Taking all of these together, Miner i can pick a sufficiently small $\delta > 0$, set ε as a function of δ , and strictly increase their utility. Again, the main idea is that they get a discrete jump in sales that earn positive, while suffering only a continuous loss in mining revenue.

Let us next consider the case that $HC_i \neq HC_j$ for some $i, j \in A$, and $\max_{i \in A} \{HC_i\} + \min_{i \in I} \{OC_i\} \geq \tilde{c}^*$. Let again $j := \arg \max_{j \in A} \{HC_j\}$, and let i be some miner with $HC_i < HC_j$. Miner i will again try to set a price $HC_j - \varepsilon$, which is a strictly profitable deviation *unless it causes some inactive miner to enter the market*. But if $HC_j + \min_{i \in I} \{OC_i\} \geq \tilde{c}^*$, then in fact *no inactive miner will enter the market at price HC_j* . Moreover, even at price $HC_j - \varepsilon$, it is still the case that c^* changes continuously in p_i , and so again for all δ there exists a sufficiently small ε such that $c^* \geq \tilde{c}^* - \delta$ at price $HC_j - \varepsilon$ (even taking into account that this price causes additional miners to join the market, albeit purchasing tiny amounts of computational power). The analysis from the previous paragraph again holds for the new revenue achieved by Miner i , so this is again strictly profitable. Again, the main idea is that the mining revenues change continuously in p_i *even considering the new miners who may enter at price $< HC_j$* , but there is a discrete jump in revenue from hardware sales. \square

Finally, we prove Lemma 10, establishing conditions under which a active miner can strictly benefit by selling hardware to a inactive miner (under the hypotheses imposed by Lemmas 8 and 9).

Lemma 10. *If “no sales” is an equilibrium with $|A| \geq 3$, then either I is empty, or $\max_{i \in A} \{OC_i\} < \min_{i \in I} \{OC_i\}$.*

Proof. By Lemma 9, there are two possible cases for a no-sales equilibrium. First, perhaps $\max_{i \in A} \{HC_i\} + \min_{i \in I} \{OC_i\} < \tilde{c}^*$. This means that if miner $\ell \in A$ sets price $p_\ell := \tilde{c}^* - \min_{i \in I} \{OC_i\} > HC_\ell$, they do not yet set a price low enough to replace the hardware cost of any miners in A (who are producing some hardware by definition). On the other hand, this price is just low enough that now at least one miner (say it is a total of k miners) have cost *exactly* \tilde{c}^* .

Similarly, if HC_i is the same for all $i \in A$ (and I is non-empty), assume for contradiction that $OC_j \leq OC_i$ for some $i \in A, j \in I$. Then clearly, $OC_j + HC_i < \tilde{c}^*$, as $OC_j + HC_i \leq$

$\text{OC}_i + \text{HC}_i < \tilde{c}^*$. Then we draw the same conclusion: any miner $\ell \in A$ can set price $p_\ell := \tilde{c}^* - \min_{i \in I} \{\text{OC}_i\} > \text{HC}_\ell$, and this results in a price just low enough that now $k \geq 1$ miners have cost *exactly* \tilde{c}^* .

This means that if miner ℓ were to further slightly lower p_ℓ (which is feasible, as $p_\ell > \text{HC}_\ell$), these miners would choose to purchase non-zero hardware. For simplicity of notation, denote by $\text{OC} := \min_{i \in I} \{\text{OC}_i\}$, and observe that there exists a sufficiently small $\varepsilon > 0$ such that for any $p_\ell \in (\tilde{c}^* - \text{OC} - \varepsilon, \tilde{c}^* - \text{OC})$, Miner ℓ 's payoff as a function of p_ℓ in this range is (where c^* is a function of p_ℓ):

$$\begin{aligned} U_\ell(p, Q) &= (1 - \tilde{c}_\ell/c^*)^2 + k(p_\ell - \text{HC}_\ell) \cdot (1 - \frac{\text{OC} + p_\ell}{c^*})/c^* \\ &= (1 - \tilde{c}_\ell/c^*)^2 + (p_\ell - \text{HC}_\ell) \cdot (1 - \sum_{j \in A} 1 - \frac{\tilde{c}_j}{c^*})/c^* \\ &= (1 - \tilde{c}_\ell/c^*)^2 + (p_\ell - \text{HC}_\ell) \cdot (-(|A| - 1) + \sum_{j \in A} \frac{\tilde{c}_j}{c^*})/c^*. \end{aligned}$$

In the second equality, we have used the fact a total quantity of $1/c^*$ units of computational power are purchased, and therefore all units *not* owned by miners in A are purchased from Miner ℓ (and that $\sum_{j \in A} 1 - \tilde{c}_j/c^*$ units are owned by miners in A).

We are interested in first taking the derivative with respect to c^* , and then evaluating the derivative at $c^* = \tilde{c}^*$, to argue that Miner ℓ would strictly profit by lowering p_ℓ to induce a $c^* < \tilde{c}^*$.

$$\begin{aligned} \frac{\partial U_\ell(p, Q)}{\partial c^*} &= \frac{2(1 - \tilde{c}_\ell/c^*)\tilde{c}_\ell}{(c^*)^2} + \frac{\partial p_\ell}{\partial c^*} \cdot (-(|A| - 1) + \sum_{j \in A} \frac{\tilde{c}_j}{c^*})/c^* \\ &\quad + \frac{(|A| - 1)(p_\ell - \text{HC}_\ell)}{(c^*)^2} - \frac{2(p_\ell - \text{HC}_\ell) \sum_{j \in A} \tilde{c}_j}{(c^*)^3}. \end{aligned}$$

To evaluate the derivative at $c^* = \tilde{c}^*$, observe first that $(-(|A| - 1) + \sum_{j \in A} \frac{\tilde{c}_j}{\tilde{c}^*}) = 0$. The reason for this is because exactly miners in A are active in equilibrium at costs \tilde{c} , and therefore they must be responsible for the entire market share (intuitively, it also makes sense that we should not care about the change in p_ℓ , because there are no sales at (p, Q)). Similarly, observe that the price which induces $c^* = \tilde{c}^*$ is exactly $p_\ell = \tilde{c}^* - \text{OC}$. We will let $\Delta := \text{OC}_\ell - \text{OC}$, and observe that now $p_\ell = \tilde{c}^* - \text{OC}_\ell + \Delta$ (and therefore $p_\ell - \text{HC}_\ell = \tilde{c}^* - \tilde{c}_\ell + \Delta$).

So after these substitutions, we get that:

$$\begin{aligned}
\frac{\partial U_\ell(p, Q)}{\partial c^*}(\tilde{c}^*) &= \frac{2(1 - \tilde{c}_\ell/\tilde{c}^*)\tilde{c}_\ell}{(\tilde{c}^*)^2} + \frac{(|A| - 1)(\tilde{c}^* - \tilde{c}_\ell + \Delta)}{(\tilde{c}^*)^2} - \frac{2(\tilde{c}^* - \tilde{c}_\ell + \Delta) \sum_{j \in A} \tilde{c}_j}{(\tilde{c}^*)^3} \\
&= \frac{2(1 - \tilde{c}_\ell/\tilde{c}^*)\tilde{c}_\ell}{(\tilde{c}^*)^2} + \frac{(|A| - 1)(\tilde{c}^* - \tilde{c}_\ell + \Delta)}{(\tilde{c}^*)^2} - \frac{2(\tilde{c}^* - \tilde{c}_\ell + \Delta)(|A| - 1)\tilde{c}^*}{(\tilde{c}^*)^3} \\
&= \frac{2(\tilde{c}^* - \tilde{c}_\ell)\tilde{c}_\ell + (|A| - 1)(\tilde{c}^* - \tilde{c}_\ell + \Delta)\tilde{c}^* - 2(\tilde{c}^* - \tilde{c}_\ell + \Delta)(|A| - 1)\tilde{c}^*}{(\tilde{c}^*)^3} \\
&= \frac{2(\tilde{c}^* - \tilde{c}_\ell)\tilde{c}_\ell - (\tilde{c}^* - \tilde{c}_\ell + \Delta)(|A| - 1)\tilde{c}^*}{(\tilde{c}^*)^3} \\
&= \frac{(\tilde{c}^* - \tilde{c}_\ell)}{(\tilde{c}^*)^3} \cdot \left(2\tilde{c}_\ell - (|A| - 1) \left(\tilde{c}^* + \frac{\tilde{c}^* \Delta}{\tilde{c}^* - \tilde{c}_\ell} \right) \right) \\
&= \frac{(\tilde{c}^* - \tilde{c}_\ell)}{(\tilde{c}^*)^3} \cdot (2\tilde{c}_\ell - \tilde{c}^*(|A| - 1)) - \frac{\Delta(|A| - 1)}{(\tilde{c}^*)^2}.
\end{aligned}$$

The first equality uses the definition of \tilde{c}^* , which implies that $\sum_{i \in A} (1 - \tilde{c}_i/\tilde{c}^*) = 1$, and therefore we can replace $\sum_{j \in A} \tilde{c}_j$ with $(|A| - 1)\tilde{c}^*$. The remaining equalities follow from algebraic manipulation. We now want to conclude with sufficient conditions for this derivative to be negative (implying that lowering p_ℓ by a bit will strictly increase Miner ℓ 's utility). First, observe that $\tilde{c}^* > \tilde{c}_\ell$ because $\ell \in A$, so the sign depends entirely on the term $\left(2\tilde{c}_\ell - (|A| - 1) \left(\tilde{c}^* + \frac{\tilde{c}^* \Delta}{\tilde{c}^* - \tilde{c}_\ell} \right) \right)$.

We now argue that when $|A| \geq 3$ and $\Delta \geq 0$ that the term is strictly negative. Indeed, when $|A| \geq 3$, $2\tilde{c}_\ell - (|A| - 1)\tilde{c}^* \leq 2(\tilde{c}_\ell - \tilde{c}^*) < 0$. When $\delta \geq 0$, $\frac{\Delta(|A| - 1)}{(\tilde{c}^*)^2} \geq 0$. Therefore, $\frac{\partial U_\ell(p, Q)}{\partial c^*}(\tilde{c}^*) < 0$, and Miner ℓ can strictly profit by bringing an inactive miner into the market.

This completes the proof. To recap the outline: we only needed Lemma 9 to claim that in any possible instance where “no sales” is an equilibrium, it must be the case that *every* active miner ℓ can bring some inactive miner into the market by setting a sufficiently low price $p_\ell > \text{HC}_\ell$. Once we have this, we establish that this strategy is indeed strictly profitable if there is any inactive miner with weakly lower operational cost. \square

At this point, we conclude the proof of Theorem 4

Proof of Theorem 4. Lemma 8 establishes the first bullet: $\text{HC}_i \leq \text{HC}_j$ whenever Miner i is active and Miner j is not. Lemma 10 establishes the second: $\text{OC}_i < \text{OC}_j$ whenever Miner i is active and Miner j is not. Lemma 9 establishes the third bullet. \square