Bitcoin: A Natural Oligopoly

Nick Arnosti* Matt Weinberg†

January 26, 2020

Abstract

Although Bitcoin was intended to be a decentralized digital currency, in practice the production and ownership of Bitcoin mining hardware is highly concentrated, and the largest miners sell hardware to their competitors.

We argue that this arises naturally from the economic incentives of Bitcoin mining. We model Bitcoin mining as a rent seeking competition: each miner i chooses a mining quantity q_i, incurs cost $c_i q_i$, and receives a reward proportional to q_i.

We first consider a model where the costs c_i are exogenous and non-transferable, and show that even small cost asymmetries result in highly concentrated *ownership* of mining equipment. We then consider a model where cost advantages can be sold to competitors (i.e. superior hardware). In equilibrium, the most efficient miner sells at a price that all other miners are willing to pay. This results in concentrated *production*: the largest miner sells to all of its competitors.

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 $150 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. The concentration is frequently cited as a major concern, and has motivated several new cryptocurrencies. For example, the Nxt whitepaper[2] 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 Gold[3] 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 Satoshi Nakamoto's idealistic vision of "one CPU one vote" has been superseded by a reality where the manufacture and distribution of mining equipment has become dominated by a very small number of entities.*

*Columbia Business School

†Princeton University, Department of Computer Science.

[1] Source: www.coinmarketcap.com. Accessed January 20, 2020.

[2] Source: https://whitepaperdatabase.com/nxt-nxt-whitepaper/

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

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 owns and operates its own mining equipment, while selling to its competitors.

The key questions motivating this work are as follows.

Q1 Why is the ownership of mining equipment so concentrated?

Q2 Does it make sense for a dominant miner to sell its technology to competitors?

We address these questions by modeling Bitcoin mining as a rent-seeking contest. Regarding the first question, our model with exogenous costs predicts that seemingly small cost advantages are amplified in equilibrium, resulting in highly concentrated *ownership* of mining power. Regarding the second, our model predicts that when cost advantages can be sold to competitors, the lowest-cost miner sells its technology to all competitors. This holds no matter the quality of competitors' technology, and results in highly concentrated *production* of mining power. Below, we provide more detail about the relevant aspects of Bitcoin mining, our model, and our results.

1.1 What is Bitcoin Mining?

Bitcoin was designed to replace centralized digital currencies like Paypal or Visa. While many people (including the authors) are perfectly happy to use these systems, others are repulsed at the thought of their transaction fees going to a corporate CEO, or are concerned that the US government might pressure an organization to freeze a user's funds or reveal their transaction history. The promise of Bitcoin was that no single entity would be in control, but instead millions of miners in a peer-to-peer network would each contribute to maintaining a record of past transactions, also known as a "ledger."

There are two fundamental challenges facing Bitcoin and other cryptocurrencies. First, how to ensure consistency among the many copies of the transaction history? Second, how to incentivize users to store and update this information? Achieving consensus is a notoriously difficult problem studied by computer scientists for several decades (Shostak et al., 1980; Lamport et al., 1982; Fischer et al., 1985), and is especially challenging in permissionless environments where users can anonymously join and leave as they please.

In order to ensure consistency, the Bitcoin protocol makes it difficult to amend the ledger. In particular, each "block" of new transactions must be accompanied by the solution to a cryptopuzzle. The puzzle is designed in such a way that the fastest way to find a

[4] Source: https://www.coindesk.com/bitmain-by-the-numbers-an-inside-look-at-a-bitcoin-mining-empire

solution is simply to guess at random.[5] The process of randomly guessing solutions to a cryptopuzzle in search of rewards is referred to as "mining," and agents who participate in this process are called "miners."

Because miners are guessing solutions at random, the rate at which they succeed is proportional to their rate of guessing. Whenever a miner solves a puzzle, it receives a "block reward" (currently 12.5 bitcoin), as well as fees offered by those whose transactions are being processed. Importantly, the Bitcoin protocol has a hard-coded "difficulty adjustment" for the cryptopuzzles so that the rate at which puzzles are solved (and block rewards are granted) is independent of the total computational power ("hash rate") in the network.

Although all miners are guessing solutions randomly, there are several ways that a miner can gain an advantage over others:

- **Electricity.** Cheap electricity lowers the cost of powering mining equipment.
- **Cooling.** Cold locations lower the cost of cooling mining equipment.
- **Hardware.** Specialized hardware ("ASICs") can be tailored to use far less energy when solving bitcoin cryptopuzzles. For example, state-of-the-art ASICs are estimated to be well over 1000x as efficient (hashes per joule) as state-of-the-art general-purpose hardware (GPUs).[6]

These factors play a significant role in determining who enters the market: only those with access to specialized mining hardware and cheap electricity can mine profitably.

The end result is that Bitcoin mining is concentrated among a small number of entities. This concentration can be measured in several ways:

- **Mining Pools.** "Mining pools" are groups of miners that share rewards with each other in order to decrease the variance of short-term returns. The most frequently-cited evidence for concentration of Bitcoin mining is the fact that a small number of large mining pools account for most blocks.[7] In fact, this data may understate pool centralization: Bitmain owns two of the largest mining pools (Antpool and BTC.com), and is the sole investor in another (viaBTC).[8] Furthermore, even blocks of unknown origin might be associated with large pools.[9]

[5] At least this is widely believed to be the case, and is true under the assumption that SHA-256 is an ideal hash function.

[6] Source: https://en.bitcoin.it/wiki/Mining_hardware_comparison, https://en.bitcoin.it/wiki/Non-specialized_hardware_comparison.

[7] Current shares from each pool are available at https://blockchain.info/pools?timespan=4days. As of January 20, 2020, four large pools account for a majority of blocks.

[8] Source: https://bitcoinmagazine.com/articles/bitmain-nears-51-network-hash-rate-why-matters-and-why-it-doesnt.

[9] Source: https://diar.co/volume-3-issue-1/. This possibility seems likely given that this increase in "unknown" blocks occurred shortly after Bitmain's stake in multiple large pools attracted attention.

It is unclear whether large mining pools are a significant concern. On the one hand, the pool operator typically determines the content of blocks mined by the pool, and therefore has significant influence over the ledger. On the flip side, if each pool consists of many small miners who can easily switch pools, then this influence is mitigated.[10] Cong et al. (2018) conclude that in this case, the pool size distribution is irrelevant. One might even argue that mining pools encourage decentralization by allowing small miners to collect rewards regularly. In part because of the unclear relationship between mining pools and mining centralization, we do not model mining pools in this paper.

- **Ownership of Mining Equipment.** Although the anonymity of bitcoin mining makes it difficult to determine ownership, several recent studies have estimated that somewhere between eight and eleven large miners control a majority of global mining power (Hileman and Rauchs, 2017; Gencer et al., 2018). This is concerning because these miners could freeze any user's funds, wipe past transactions from the record, or launch other attacks (for more details, see Narayanan et al. (2016)).

- **Production of Mining Equipment.** As mentioned above, a 2018 prospectus estimates that 75% 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.[11]

This paper focuses on concentration of ownership and production of mining equipment. As noted above, both forms of concentration are prevalent, and they pose serious threats to the Bitcoin network. We seek to understand the reasons for this concentration, and whether it can be expected to persist.

1.2 Overview of Results

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

- The total value of rewards available to miners is fixed.[12]

- Miners make costly investments in computational power.

[10] Romiti et al. (2019) present evidence suggesting that in fact, rewards in several mining pools are concentrated among a small set of miners. However, this relates most closely to the concentration of ownership of mining equipment.

[11] See https://www.rudebaguette.com/en/2017/05/antbleed-bitmain-shut-half-bitcoin/.

[12] As explained above, on average one new block is created every 10 minutes, and the size of the block reward is fixed. Regarding transaction fees, the model of Huberman et al. (2017) suggests they should not depend on the hash rate, nor on how it is distributed among miners.

- Mining costs vary for several reasons. Some cost advantages (i.e. advanced hardware) can be more easily transferred than others (such as cheap electricity and cooling).
- Miners earn rewards in proportion to their computational power.

Section 3 focuses on the case where cost advantages are non-transferable using the following stylized model:

- There is a fixed reward of value 1 to be split among n miners.
- Each miner i can acquire computational power at unit cost c_i.
- Each miner i chooses computational power q_i (incurring cost $c_i q_i$).
- Miner i receives reward $\frac{q_i}{\sum_j q_j}$.

In this model, there is a unique equilibrium, described in Theorem 1. If all costs c_i are identical, the outcome is decentralized: in equilibrium, each miner possesses a $1/n$ fraction of total mining power. This outcome arises because investment in mining power exhibits diminishing marginal returns: additional investment lowers the value of earlier investments.

In light of this, one might expect that ownership to remain decentralized so long as costs are not "too different." Our (qualitative) finding is that even small cost advantages result in surprising levels of mining power centralization. More specifically, each miner's market share in equilibrium is equal to their cost advantage relative to a hypothetical marginal miner who is indifferent about whether to participate. For example, a miner with costs that are 10% lower than the marginal miner will possess 10% of all mining power. We show by example that with even moderate cost asymmetries, this results in most mining power being controlled by only a few miners.

Section 4 considers a model where cost advantages are perfectly transferable, and can be sold to competitors. This might be suitable for modeling hardware advantages, as hardware can easily be sold. We show that in equilibrium, the most efficient miner will sell to all of its competitors (this equilibrium is generically unique). It is perhaps unsurprising that, *given that sales occur*, the most efficient miner undercuts all competitors. More surprising is that this miner always chooses to sell their hardware, thereby lowering its market share. Although this miner recuperates some profit from these sales, in many cases its total profit is lower than if no sales had occurred.

We note that both of our models abstract from the question of how mining costs are divided between acquiring, powering, cooling and maintaining hardware. Instead, the important distinction is whether cost advantages of one miner can easily be sold to others. Our main findings are that

(i) even small non-transferable cost advantages lead to concentration of ownership, and

(ii) transferable cost advantages will be shared with the entire market, resulting in concentration of production.

2 Related Work

There are two classes of related work to discuss. The first consists of papers about Bitcoin. These papers touch on some of the themes explored here (and justify some of our modeling assumptions), but are for the most part mathematically unrelated to our work. The second consists of the literature on "rent-seeking contests".

2.1 Cryptocurrencies

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

Numerous sources have empirically documented high levels of concentration in Bitcoin mining. The most frequently-cited statistics come from publicly available data on large mining pools. At the time of writing, four mining pools are responsible for over 50% of mined blocks.[13] While this is potentially concerning, because pool managers do not own all of the the hardware mining on their behalf, it is unclear whether large pools are problematic.

Other measures of centralization are not publicly available, and are often intentionally obscured. Nevertheless, prior work has concluded that 56% of Bitcoin *nodes*[14] exist in large data centers (Gencer et al., 2018), and that approximately 100 nodes are responsible for the initial announcement of 75% all blocks. Even more pertinent to our work is a recent report on the distribution of mining power, which identifies eleven "large mining entities" and estimates that they collectively control a majority of global computational power dedicated to Bitcoin mining (Hileman and Rauchs, 2017).

One key feature of 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 proportional reward scheme is the unique reward scheme satisfying natural axioms (anonymity, sybil-proof, and collusion-proof), implying that our results hold for any competition with these properties.

Huberman et al. (2017) study the fee-setting game facing Bitcoin users, and observe that the revenue raised by the system is determined by users' willingness to pay to avoid delays, and by the (exogenously specified) throughput rate. These factors determine aggregate miner compensation, justifying our assumption that the total reward is exogenous to miner behavior.

Cong et al. (2018) study a game where miners choose between mining solo and joining one of several pools. They conclude that the presence of pools increases the total hashrate, but that if miners can freely switch between pools, the distribution of mining pool size is

[13] Source: https:blockchain.info/pools?timespan=4days. Accessed January 20, 2020.

[14] A node listens for transactions and blocks, and forwards them to the rest of the network (additionally checking for validity before forwarding).

irrelevant. Although they consider homogeneous miners, this conclusion partially justifies our choice to focus on concentration of ownership and production of mining equipment, rather than pool size.

Prat and Walter (2018) model the Bitcoin mining game, but focus on very different features of this game than we do. Their model is dynamic, and incorporates the purchase of gradually better equipment. However, they assume that all miners have identical electricity costs, and that machines can never be sold. By contrast, our work focuses on cost asymmetries and allows miners to sell equipment to others. Ma et al. (2018) also consider a dynamic mining game with free entry of symmetric miners.

Many papers on Bitcoin assume a large number of small and symmetric miners who earn zero profit in equilibrium (Kroll et al., 2013; Budish, 2018; Chiu and Koeppl, 2018; Ma et al., 2018; Easley et al., 2019). Our paper models markets with asymmetric miners, and concludes that in most cases, mining will be dominated by a small number of large players. This suggests that conclusions from the aforementioned papers should be examined for robustness to the presence of large profit-making miners.

There is also a line of papers which 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). These papers provide clever strategies that allow a miner with an x_i fraction of computational power to claim a fraction of mining rewards exceeding x_i. These papers treat the computational power of each miner as fixed (exogenous), although Eyal and Sirer (2014) informally discusses the idea that pools using these "attacks" may be able to recruit more participants. In contrast to these papers, we formally study miners' incentives to acquire more computational power. Moreover, our work in Appendix D complements this literature by pointing out that even the minor economies of scale induced by such strategic deviations promote centralization of mining power.

2.2 Market Participation

Our non-transferable model (Section 3) is formally identical to that proposed by Tullock (1980). That paper prompted a large literature on "imperfectly discriminating contests" (so-named to highlight the contrast with the "perfectly discriminating contest" of an all-pay auction) – see Nitan (1994) and Lockard and Tullock (2001) for surveys of this literature. Many of these papers focused on the extent of *rent dissipation* – that is, comparing expenses incurred by contestants to the value of the prize.

Two early papers explicitly address the number of entrants, as we do (Hillman and Riley, 1989; Gradstein, 1995). Proposition 5 in Hillman and Riley (1989) is similar to our Theorem 1, although it does not explicitly characterize the market share of each participant. Gradstein (1995) assumes that costs are drawn iid, and shows that as the number of potential entrants n grows, the fraction who choose to enter in equilibrium converges to zero. For obvious reasons, neither paper discusses implications of these findings for Bitcoin and other proof-of-work protocols.

We know of two other papers that use Tullock's rent-seeking model (our model of non-transferable cost advantages) to study Bitcoin. Dimitri (2017) notes that the set of participants does not depend on the size of the reward (which we normalize to one), and that at least two miners will participate, but does not study market concentration. Contemporaneous with our work, Alsabah and Capponi (2019) consider a two-stage model in which firms invest in R&D in the first stage (lowering their costs c_i), and then choose how much mining power to acquire in the second. Their analytical results consider the case where firms are ex-ante identical and play symmetric investment strategies. They also include a numerical study with cost asymmetries that suggests that the R&D phase increases market concentration.

Relative to Hillman and Riley (1989), Gradstein (1995), Dimitri (2017) and Alsabah and Capponi (2019), we provide a cleaner characterization of the equilibrium of the non-transferable game based on the quantity c^*. This makes it easier to reason about the level of concentration in equilibrium: each miner's market share is simply their cost advantage relative to c^*. Most importantly, none of these papers consider the possibility of one miner selling technology to another, which we study in Section 4.

3 Concentration of Ownership

In this section, we address our first key question: why is ownership of mining equipment so concentrated? We argue that this arises naturally if miners have different costs and cannot easily sell their cost advantages to others.

Recall our formal model: there are $n \geq 2$ miners competing for a prize of value 1. Miners simultaneously choose a quantity of computational power to acquire and operate. If miner i chooses quantity q_i, it pays cost $c_i q_i$. Miners then earn a reward proportional to q_i. In other words, the utility for miner i when computational power is given by q is

$$U_i(q) = x_i(q) - c_i q_i, \tag{1}$$

where

$$x_i(q) = q_i / \sum_{j=1}^{n} q_j, \tag{2}$$

and $x_i(q) = 0$ for all i if $q_i = 0$ for all i.

Without loss of generality, we assume $0 < c_1 \leq c_2 \leq \cdots \leq c_n$. Costs are common knowledge, and exogenously fixed. Because any advantage that miner i has over miner j is non-transferable, we refer to this as the *non-transferable game*.[15]

[15] To clarify the relationship between our model and reality, we note that c_i is intended to represent the *total* cost for miner i to acquire and operate one unit of computational power. We are agnostic about whether this cost comes primarily from acquisition or operation. The assumption that cost advantages are non-transferable means that Miner 2 cannot lower its costs by contracting with Miner 1. This might be reasonable, for example, if Miner 1's cost advantage comes from being located in a region with cheap electricity.

The main result of this section characterizes a unique equilibrium outcome. To state this outcome, it is useful to define the function $X : \mathbb{R}_+ \to \mathbb{R}_+$ by

$$X(y) = \sum_{i=1}^{n} \max\{1 - c_i/y, 0\}. \tag{3}$$

Lemma 1. *There is a unique value c^* satisfying $X(c^*) = 1$. Furthermore, $c^* \in (c_2, c_1 + c_2]$.*

Proof. X is continuous and weakly increasing in y. Furthermore, X is strictly increasing on (c_1, ∞), with $X(c_2) = 1 - c_1/c_2 < 1$ and $X(c_1 + c_2) \geq (1 - \frac{c_1}{c_1+c_2}) + (1 - \frac{c_2}{c_1+c_2}) = 1$. □

For the remainder of this paper, we let c^* denote the unique solution satisfying $X(c^*) = 1$. We note that the function $X(\cdot)$ (and thus the value c^*) is implicitly parameterized by the costs $c_1, \ldots, c_n$. The value c^* provides a clean characterization of equilibrium outcomes.

Theorem 1. *In the non-transferable game, there is a unique pure strategy equilibrium. In it, $q_i = \frac{1}{c^*} \max\{1 - c_i/c^*, 0\}$, $x_i(q) = \max\{1 - c_i/c^*, 0\}$, and $U_i(q) = x_i(q)^2$.*

The proof of Theorem 1 is in Appendix A. In the remainder of this section, we focus on its interpretation and implications, and use q^* to denote the unique equilibrium outcome. Theorem 1 immediately implies the following, confirming that diminishing marginal returns indeed implies perfect decentralization in absence of cost asymmetries.

Corollary 1. *In the non-transferable game, if miners have identical costs ($c_1 = c_2 = \cdots = c_n$), then for all i, $x_i(q^*) = 1/n$, and $U_i(q^*) = 1/n^2$.*

Though the presence of a symmetric equilibrium is not surprising, one might ask why there are no asymmetric equilibria. The answer is diminishing rewards: because additional mining lowers the return for existing machines, the more a miner invests, the less it gains from additional investment. As a result, none of the n miners benefit from increasing their investment, despite the fact that an $(n+1)^{st}$ miner with identical costs would find it profitable to enter.[16] Although this encourages a decentralized market, we will see that this force is dwarfed by the effect of asymmetric costs.

When costs differ, the quantity c^* is very helpful for describing and reasoning about equilibrium outcomes. First, it serves as a participation threshold: miner i chooses $q_i^* > 0$ if and only if $c_i < c^*$. Second, the ratio c_i/c^* is sufficient to determine the market share and utility of miner i. In particular, the market share of each miner is precisely the percentage by which their costs are lower than c^*: a miner with costs that are 10% lower than this "break-even" cost will control 10% of *all* computational power in equilibrium. This implies that any market where miners with notably different costs participate will be dominated by a few large players.

[16] For this same reason, it is never an equilibrium for a single miner to control the entire market, no matter their cost advantage. Such a miner would have an incentive to lower their investment, at no cost to their market share.

Corollary 2. *In the non-transferable game, if miner i participates at all in the unique equilibrium (that is, $q_i > 0$), then $x_j(q) \geq 1 - \frac{c_j}{c_i}$ for all j. That is, the market share of miner j is at least $1 - \frac{c_j}{c_i}$.*

Proof. By Theorem 1, $x_j(q) = 1 - c_j/c^*$ and $c_i < c^*$ for any miner who participates. □

One way to think about Corollary 2 is as follows. Suppose that we wish to know whether miner k will participate in equilibrium. For each lower-cost miner, ask "by what percentage are this miner's costs lower than those of miner k?" Miner k will participate if and only if the sum of these percentages is less than 100.[17]

We now provide two brief examples illustrating a high degree of centralization. In both cases, mining costs differ by a factor of at most two. In the first example, only 7 miners choose to participate. In the second, all miners participate, but the two largest control 75% of the market.

Example 1. *Let $c_i = i/(i+1)$. We have $c_7 < c^* \approx 0.88 < c_8$, so 7 miners participate. Moreover, $c_7 = 7/8$, so Corollary 2 implies that Miner 1 controls more than $1 - \frac{1/2}{7/8} = 3/7$ of the market and Miner 2 controls more than $1 - \frac{2/3}{7/8} = 5/21$ of the market. Jointly, they control more than $2/3$ of the market.*

Example 2. *Let $c_i = 1 - 2^{-i}$. Note that $X(1) = \sum_i (1 - c_i) < 1$, from which it follows that $c^* > 1$ (because $X(c^*) = 1$ and X is non-decreasing). Therefore, all miners participate in equilibrium. However, in equilibrium we have $x_i = 1 - c_i/c^* \geq 1 - c_i = 2^{-i}$.*

Example 2 illustrates how small cost advantages are magnified in equilibrium. Note that $c_5/c_6 \approx 0.984$, so miner 5 has a 1.6% cost advantage over miner 6. Nevertheless, in equilibrium, miner 5 owns approximately *twice* as much computational power!

Although these examples clearly exhibit high concentration of mining power despite small cost asymmetries, we also seek a measure to quantify this. One standard measure of market concentration is the Herfindahl-Hirschman Index (HHI), defined as the sum of squares of market shares:[18]

$$HHI(q) = \sum_i x_i(q)^2. \tag{4}$$

Corollary 3. *In equilibrium of the non-transferable game, total miner profit is exactly equal to the Herfindahl-Hirschman Index (HHI) of market concentration:*

$$HHI(q^*) = \sum_i U_i(q^*).$$

[17] For example, if there are three miners with costs that are 20% lower than k's, and five more with costs that are 10% lower, then k will not participate, as $3 \times 20 + 5 \times 10 = 110 > 100$.

[18] See 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.

Note that by Theorem 1, the total amount of mining power acquired is $1/c^*$. If this were all held by miners with cost c^*, then all mining profits would be dissipated. In practice, miners with lower costs earn positive utilities, implying that the fraction of rewards that are spent on mining is one minus the concentration of the mining market.

4 Concentration of Production

Section 3 addressed the case where miners' cost advantages could not be shared with others. We now consider the case of a cost advantage that is perfectly transferable. For example, if one believes that hardware costs are transferable but electricity costs are not, then this section assumes that all miners pay the same electricity cost, but have access to different hardware. The key question we wish to address is whether it makes sense for the most efficient miner to sell to its competitors. We find that it does: generically, there is a unique equilibrium, in which the most efficient miner sells at a price low enough that all competitors purchase from it.

As before, we consider a model with n potential miners. The cost to miner i of producing and operating one unit of computational power is $\tilde{c}_i$, with $\tilde{c}_1 \leq \tilde{c}_2 \leq \cdots \leq \tilde{c}_n$. The game proceeds in two rounds:

1. Each miner i chooses a price $p_i \geq 0$ at which they are willing to sell their computational power.[19]

2. Each miner i chooses the quantity q_{ij} of computational power sourced from miner j. We let $q_i = \sum_j q_{ij}$ be the total computational power of miner i.[20]

Given prices $p = \{p_j\}$ and purchase quantities $q = \{q_{ij}\}$, the utility of miner i is

$$U_i(p, q) = \frac{q_i}{\sum_j q_j} - \sum_j q_{ij}\tilde{c}_j + \sum_{j \neq i} p_i q_{ji} - \sum_{j \neq i} p_j q_{ij}.$$

That is, each miner gets its share of the rewards, pays for the direct costs of the computational power that it uses, receives profit from any computational power that it sells, and pays competitors for computational power purchased from them.

[19] The restriction to anonymous pricing simplifies exposition, but is not important to our results. In a model with personalized prices, it would remain the case that unless $\tilde{c}_1 = \tilde{c}_2 \leq \tilde{c}_3/2$, the only equilibrium is for Miner 1 to sell to all of its competitors at price $\tilde{c}_2 - \tilde{c}_1$: see Appendix C for more details.

[20] As noted in Section 3, we are agnostic about whether these costs are primarily from production or operation. That is, $\tilde{c}_i = \$10$ may mean that the hardware costs \$1 to produce and \$9 to operate, or vice versa. If the cost $\tilde{c}_i$ is primarily from production, then miner j should pay a high price to miner i but then have low costs of operation; if $\tilde{c}_i$ is mostly from operation, then miner j should pay a lower up-front price to i and face higher operation costs. Either scenario is consistent with our model, with the price p_i is interpreted as the *profit* (rather than revenue) that i receives from each unit of computational power sold, and $p_i + \tilde{c}_i$ the cost to j of using computational power from i.

We seek a subgame perfect equilibrium of this two-stage game. Taking prices p as given, we note that the per-unit cost to miner j of sourcing from miner i is $\tilde{c}_i + \mathbf{1}(i \neq j)p_i$. Thus, miners are effectively playing the game from Section 3 where miner i has a per-unit cost of

$$c_i(p) = \min_j \{\tilde{c}_j + \mathbf{1}(i \neq j)p_j\}. \tag{5}$$

Given these costs, miners will acquire computational power from the cheapest source, in quantities given by Theorem 1.

Definition 1. *A function Q mapping prices to purchase quantities is a* **second-stage equilibrium** *if for any miner j, and any q' such that $q'_{i\ell} = Q_{i\ell}(p)$ for $i \neq j$,*

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

It is **lexicographic** *if in addition, each miner j uses exclusively the lowest-index source offering unit cost $c_j(p)$.*

Observation 1. *There is a unique lexicographic second-stage equilibrium, Q^*. Q^* sets (total) purchase quantities according to Theorem 1, and has each miner j purchase exclusively from lowest-index source offering unit cost $c_j(p)$.*

We focus on lexicographic equilibria (assume that miners break ties by purchasing from the lowest-indexed source) in order to make the statement of Theorem 2 clean. This restriction has bite only in the case where $\tilde{c}_1 = \tilde{c}_2$, as in that case there will be equilibria where miners buy exclusively from Miner 1, exclusively from Miner 2, or from a combination of the two. However, in all equilibria, the sales price will be $\tilde{c}_2 - \tilde{c}_1 = 0$, implying that all miners obtain the same utility in every equilibrium. Appendix B includes more details.

Having determined second-stage outcomes for any prices p, we now seek an equilibrium of the pricing game, defined below.

Definition 2. *Prices p are* **a first-stage equilibrium** *for Q if for any miner j, and any p' such that $p'_i = p_i$ for $i \neq j$,*

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

They are **lexicographic** *if in addition $p_j \leq p'_j$ when the above holds with equality.*

Definition 3. *The pair (p, Q) is an equilibrium if Q is a lexicographic second-stage equilibrium and p is a lexicographic first-stage equilibrium for Q.*

We focus on lexicographic p (assume miners break ties by setting lower prices) in order to avoid many equivalent equilibria where miners whose equipment will never be purchased set their prices arbitrarily. This restriction does not drive the main result of this section, which states that in equilibrium, all sales go to the lowest-cost miner, at price $\tilde{c}_2 - \tilde{c}_1$.

Theorem 2. *$(\langle \tilde{c}_2 - \tilde{c}_1, 0, \ldots, 0\rangle, Q^*)$ is an equilibrium in which all miners acquire computational power from miner 1. This is the only equilibrium unless $\tilde{c}_1 = \tilde{c}_2 \leq \tilde{c}_3/2$.*

Theorem 2 establishes that in equilibrium two properties hold: (i) only the most efficient miner sells computational power, (ii) the price set is just low enough so that all hardware is purchased from the most efficient miner, and (iii) the most efficient miner indeed sells computational power. Property (i) is perhaps not surprising: if any other miner were selling successfully, Miner 1 would have an incentive to undercut them. To get intuition for Property (ii), observe that the price set by Miner 1 cannot be any higher than $\tilde{c}_2 - \tilde{c}_1$, because in that case Miner 2 would have the ability and desire to undercut these sales. Although Miner 1 could lower prices further (giving up market share in return for higher sales), it turns out that this is never profitable. Together, these observations imply the following Proposition, the first two properties above (see Appendix B for a rigorous proof).

Proposition 1. *In any equilibrium (p, Q^*) where sales occur,*[21] *$p_1 = \tilde{c}_2 - \tilde{c}_1$, $p_j = 0$ for all $j > 1$, and $Q^*_{ij}(p) = 0$ for all $j > 1$.*

Property (iii) is the most surprising aspect of Theorem 2: sales occur in *all* equilibria (outside of one corner case). A priori, it seems plausible that the most efficient miners might all set high prices, preventing new competition from entering the mining industry. This scenario seems especially likely given that a buyer will never pay more than the mining rewards that the buyer will earn: the seller could earn this amount by simply keeping this computational power for itself!

What drives Proposition 2 below is that the second-stage investments $q_2, \ldots, q_{n-1}$ are not fixed as a function of $c_1, \ldots, c_{n-1}$, but *they vary in response to* c_n. Therefore, it could possibly be profitable for Miner 1 to lower c_n, only because this might lower the power acquired in second-stage by Miners 2 through $n-1$. This of course does not imply that it *will* be profitable, but it is indeed profitable (except in one degenerate case).

Proposition 2. *Unless $\tilde{c}_1 = \tilde{c}_2 \leq \tilde{c}_3/2$, sales occur in all equilibria. When $\tilde{c}_1 = \tilde{c}_2 \leq \tilde{c}_3/2$, the unique equilibrium where no sales occur is $(\langle \tilde{c}_1, \tilde{c}_1, 0, \ldots, 0\rangle, Q^*)$. Miners 1 and 2 split the market evenly, and nobody else enters.*

This establishes that generically, a situation where no sales occur cannot be an equilibrium: Miner 1 would prefer to sell some small quantity (at a high price) to a new entrant. Of course, as soon as this occurs, Miner 2 would undercut this sale, resulting in a price war that leads to the unique equilibrium identified in Theorem 2.

This final outcome may be better or worse for Miner 1 than the outcome of the non-transferable game. If $\tilde{c}_1 << \tilde{c}_2$, then Miner 1 can charge a high price for its equipment without losing much market share in the transferable game. In this case, revenue from sales more than offsets the decline in market share, and Miner 1 profits strictly more than

[21] We say that "sales occur" for (p, Q) if $Q_{ij}(p) > 0$ for some $i \neq j$.

in the non-transferable game. Conversely, if $\tilde{c}_1 \approx \tilde{c}_2$, then Miner 1 must sell its equipment for very little. Furthermore, these sales may dramatically lower costs for potential competitors, thereby significantly reducing Miner 1's market share. Although Miner 1 remains the largest miner and sells to the entire market, it would earn greater profits in the non-transferable game. We also note that Miner 2 always prefers the equilibrium of the non-transferable game to the equilibrium outcome of the transferable game: in either case, Miner 2 has cost $\tilde{c}_2$ and does not make any sales, but competition is fiercer in the transferable game, as other miners' costs have been lowered thanks to sales from Miner 1.

5 Conclusions and Discussion

The concentration of Bitcoin mining among a few large entities threatens the promise of a truly decentralized digital currency. We show that this concentration can be explained using a simple model with two key features:

(i) miners share a fixed reward in proportion to their investment in computational power,

(ii) miners have different costs.

When miners can sell their cost advantage, they will, leading to concentrated production of mining hardware. Advantages that are less easy to sell lead to concentrated ownership.

We conclude that mining centralization arises from core aspects of the Bitcoin mining protocol. Furthermore, our conclusions apply not only to Bitcoin, but to any competition with these two features. In practice, there are reasons to believe that mining features economies of scale,[22] which intuitively should only exacerbate concentration (Appendix D formalizes this for a model where rewards are proportional to q_i^α for $\alpha > 1$). All together, our findings suggest that without significant changes, the vision of a large competitive market among miners is unlikely to be fulfilled. Therefore, models which assume a competitive market with many small miners should be carefully examined to understand which conclusions are sensitive to this assumption.

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 two features identified above, our model suggests that centralization will persist. For example, innovations such as BetterHash and Stratum V2,[23] which are designed to allow miners (rather than mining pools) to dictate the contents of a block, do not address

[22] For 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); Sapirshtein et al. (2016); Kiayias et al. (2016); Carlsten et al. (2016)).

[23] For 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.

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 electricity costs play an important role, mining activity is likely to remain concentrated among the few areas of the world where electricity is cheapest.

What changes might have an impact? One thought is to change the reward structure, so that rewards scale sublinearly with mining power. In Appendix D we show that if rewards are proportional to q_i^α for some $\alpha < 1$, then all miners choose $q_i > 0$ in equilibrium, implying greater decentralization. Unfortunately, sublinear rewards seem impossible so long as mining is permissionless: several recent papers 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. In this paper, we focused on cases where cost advantages are either non-transferable or perfectly transferable. In practice, advantages likely arise from a mixture of transferable and non-transferable sources. One could consider a hybrid of the models from Sections 3 and 4 in which even miners with identical equipment (purchased at identical prices) have different operating costs (i.e. due to different electricity prices). Such a model could predict an outcome with both a dominant seller (i.e. that with the best hardware) and multiple large miners (i.e. those with the cheapest electricity).

Another abstraction in this paper is to treat the acquisition of mining power as a one-time investment. Our static model conflates the cost of acquiring equipment and the cost of powering it. In practice, fluctuations in Bitcoin 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. 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. Pittfalls of bitcoin's proof of work: R&d arms race and mining centralization. 2019.

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.

J. Chiu and T. Koeppl. The economics of cryptocurrencies – bitcoin and beyond. 2018.

L. W. Cong, Z. He, and J. Li. Decentralized mining in centralized pools. 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.

M. J. Fischer, N. Lynch, and M. S. Paterson. Impossibility of distributed consensus with one faulty process. *Journal of the ACM*, 32(2):374–382, 1985.

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.

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.

A. Kiayias, E. Koutsoupias, M. Kyropoulou, and Y. Tselekounis. Blockchain mining games. In *ACM Conference on Economics and Computation (EC)*, 2016.

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.

L. Lamport, R. Shostak, and M. Pease. The byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3), 1982.

J. Leshno and P. Strack. Bitcoin: An impossibility theorem for proof-of-work based protocols. 2019.

A. Lockard and G. Tullock, editors. *Efficient rent-seeking: Chronicle of an Intellectual Quagmire*. Springer Science+Business Media, 2001.

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.

S. Nitan. Modelling rent-seeking contests. *European Journal of Political Economy*, 10: 41–60, 1994.

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.

R. Shostak, M. Pease, and L. Lamport. Reaching agreement in the presence of faults. *Journal of the ACM*, 27(2), 1980.

G. Tullock. *Efficient rent-seeking*. 1980.

A Proofs: Non-Transferable Model

First, we state a helpful lemma showing that first-order conditions are necessary and sufficient for identifying pure-strategy equilibria.

Lemma 2. *Suppose that q^* is an equilibrium in non-transferable game. Then for each i, $\sum_{j\neq i} q_j^* > 0$, and*

$$x_i(q^*) = \max\{1 - c_i \sum_i q_i^*, 0\}. \tag{6}$$

Furthermore, any q^* satisfying (6) for all i is an equilibrium.

Proof. We first prove that if q^* is an equilibrium, then for each i we have $\sum_{j\neq i} q_j^* > 0$. If this were not the case, then no $q_i^* > 0$ can be a best response, as i could do better by choosing $q_i = q_i^*/2$. However, $q_i^* = 0$ also cannot be a best response, as i could do better with any $q_i \in (0, 1/c_i)$.

Note that $U_i(q) = x_i(q) - c_i q_i$ is continuous and differentiable in q_i on $\{q : \sum_{j\neq i} q_j > 0\}$. Slightly abusing notation, for the remainder of this proof we define

$$x_i'(q) = \frac{\partial x_i(q)}{\partial q_i} = \frac{\partial}{\partial q_i}\frac{q_i}{\sum_j q_j} = \frac{1}{\sum_j q_j} - \frac{q_i}{(\sum_j q_j)^2} = \frac{1 - x_i(q)}{\sum_j q_j} \geq 0, \tag{7}$$

and

$$U_i'(q) = \frac{\partial U_i(q)}{\partial q_i} = x_i'(q) - c_i. \tag{8}$$

Differentiability of U_i implies that if q_i^* is a best response to q_{-i}^*, then

$$U_i'(q^*) \leq 0, \quad \text{with equality if} \quad q_i^* > 0. \tag{9}$$

In fact, these first-order conditions are sufficient to identify an equilibrium – that is, any q^* for which (9) holds for all i is an equilibrium. This is because $U_i(q)$ is concave in q_i. To see this, note that (7) and (8) imply:

$$\frac{\partial}{\partial q_i} U_i'(q) = \frac{\partial}{\partial q_i} x_i'(q) = \frac{-x_i'(q)\sum_j q_j - (1 - x_i(q))}{\left(\sum_j q_j\right)^2} = -\frac{2x_i'(q)}{\left(\sum_j q_j\right)^2} \leq 0.$$

It remains to show that (6) and (9) are equivalent. First, suppose that (9) holds. Then by (7) and (8), for all i we have

$$U_i'(q^*) = x_i'(q^*) - c_i = (1 - x_i(q^*))/\sum_j q_j^* - c_i \leq 0, \text{ with equality if } q_i > 0.$$

Therefore if $q_i^* > 0$, we must have $x_i(q^*) = 1 - c_i \sum_j q_j^*$, and $q_i^* = x_i(q^*) = 0$ if and only if $c_i \sum_j q_j^* \geq 1$. In other words, (6) holds. Conversely, if (6) holds, then substituting into (7) yields

$$x_i'(q^*) = \frac{1 - x_i(q^*)}{\sum_j q_j^*} = \frac{1 - \max\{1 - c_i \sum_j q_j^*, 0\}}{\sum_j q_j^*}.$$

Therefore,

i. if $q_i^* > 0$, then $x_i(q^*) > 0$ and $x_i(q^*) = c_i$, so $U_i'(q^*) = x_i'(q^*) - c_i = 0$.

ii. If $q_i^* = 0$, then $x_i'(q^*) = 1/\sum_j q_j^* \leq c_i$, so $U_i'(q^*) = x_i'(q^*) - c_i \leq 0$.

That is to say, (9) holds.

□

Next, we use Lemma 2 to claim that in any equilibrium, miners with lower cost have greater market share.

Corollary 4. *If q^* is an equilibrium of the non-transferable game and $c_i \leq c_j$, then $q_i^* \geq q_j^*$. That is, if miner j's costs are no higher than those of miner i, then miner j acquires at least as much computational power as miner i in equilibrium.*

Proof. By Lemma 2, $x_i(q_i^*) = \max\{1 - c_i \sum_j q_j^*, 0\}$, which is weakly decreasing in c. □

Proof of Theorem 1. By Lemma 2, we know that q^* is an equilibrium if and only if $x_i(q^*) = \max\{1 - c_i \sum_j q_j^*, 0\}$ for all i. But by definition of x_i, we have that

$$1 = \sum_i x_i(q^*) = \sum_i \max(1 - c_i \sum_j q_j^*, 0) = X(1/\sum_j q_j^*),$$

where X is as defined in (3). Therefore, Lemma 1 implies that $\sum_j q_j^* = 1/c^*$, so that $x_i(q^*) = \max(1 - c_i/c^*, 0)$ and $q_i^* = x_i(q^*) \sum_j q_j^* = \frac{1}{c^*} \max(1 - c_i/c^*, 0)$. From the definition of U in (1), it follows that $U_i(q) = (\max\{1 - c_i/c^*, 0\})^2 = x_i(q)^2$. In other words, there is exactly one equilibrium, and it is as described in Theorem 1.

□

We now make two observations that will prove useful in the analysis of the transferable game.

Observation 2. *The (unique) equilibrium quantities q_i^* are continuous functions of $c_1 \ldots c_n$. As a result, so are $x_i(q_i^*)$ and $U_i(q_i^*)$.*

Lemma 3. *For any $c_1, \ldots, c_n$, if we define c^* as in Lemma 1 and let $k = \sum_i \mathbf{1}(c_i < c^*)$ be the number of miners participating in equilibrium, then $(k-1)c^* = \sum_{i=1}^k c_i$.*

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

B Proofs: Transferable Model

Throughout this section, we will reuse notation from Section 3. In particular, when it is clear from context, we will just write $c_j := c_j(p)$ for the effective cost of miner j in the second-stage game (still using $\tilde{c}_j$ to denote the cost of miner j's own hardware). We again let c^* denote the unique solution to $X_{\vec{c}}(c^*) = 1$ (when costs are $c_j(p)$). We will also use shorthand notation $q_{ji} := Q_{ji}(p)$, with $q_i := \sum_j q_{ij}$, whenever p, Q is clear from context. Unsurprisingly, our analysis leverages Theorem 1.

We begin by observing a simpler form for miner utilities in any equilibrium of the transferable game.

Lemma 4. *If (p, Q) is an equilibrium of the transferable game, then:*

$$U_i(p, Q) = (\max\{1 - c_i(p)/c^*(p), 0\})^2 + \sum_j Q_{ji}(p) \cdot p_i.$$

Proof. Because $Q(p)$ is an equilibrium of the non-transferable game induced by p, this means that the utility that i derives *counting all mining rewards and cost of owned computational power* (regardless of where this power comes from) is exactly $(\max\{1 - c_i/c^*, 0\})^2$. It remains to count the profit from computational power sold to other miners, which is simply $\sum_j Q_{ji}(p) \cdot p_i$ in total. □

Lemma 4 allows us to treat the second stage implicitly, reasoning about miner utilities as a function of p. Our next observation begins reasoning about potential best responses in the first-stage game. Below and throughout this section, we will define:

$$c = c(p) := \min_j \{\tilde{c}_j + p_j\}.$$

Observe that this implies that $c_j(p) = \min\{c(p), \tilde{c}_j\}$.

Lemma 5. *Consider any $Q(\cdot)$ which is a second-stage equilibrium. Then if there exists a miner i such that $c > \tilde{c}_i$, and any miner purchases computational power from a miner $\neq i$ in $Q(p)$, then p is not a first-stage equilibrium for Q.*

Proof. We claim that miner i can better-respond by lowering p_i to $c - \tilde{c}_i - \varepsilon$, for sufficiently small ε. The point is that the loss in market share (and potential lost revenue from selling at a lower price) both move continuously with ε, whereas quantity sold jumps discretely when i lowers its price.

To see this, recall the format of $U_i(p, q)$ as given in Lemma 4, and consider the two terms separately.

Observe first that $(\max\{1 - c_i/c^*, 0\})^2$ is a continous function of $(c_1, \ldots, c_n)$. Furthermore, lowering p_i to $c - \tilde{c}_i - \varepsilon$ changes each c_j by at most ε. Therefore, for all δ, there is a sufficiently small ε such that updating p_i to $c - \tilde{c}_i - \varepsilon$ changes $(\max\{1 - c_i/c^*, 0\})^2$ by at most δ.

Similarly, q_j is also a continuous function of $(c_1, \ldots, c_n)$. Therefore, for all δ, there is a sufficiently small ε such that lowering p_i to $c - \tilde{c}_i - \varepsilon$ changes each q_j by at most δ. Importantly, however, lowering p_i to $c - \tilde{c}_i - \varepsilon$ will now cause *all miners who purchase computational power from another miner to purchase it exclusively from miner* i, because Q is a second-stage equilibrium (this is because any miner previously purchasing computational power must have had cost $c_j = c$, and now miner i is the unique miner offering computational power at cost $< c$). Therefore, if a total quantity of $z > 0$ was previously purchased from miners $\neq i$, then for all δ, there exists a sufficiently small ε such that lowering p_i to $c - \tilde{c}_i - \varepsilon$ increases the total quantity sold by miner i by at least $z - \delta$. Finally, observe that if any miner previously purchased any computational power from miner i, then we must have had $p_i = c - \tilde{c}_i$ (otherwise, Q would not be a second-stage equilibrium).

Taking all analyses together, we see that for all $\delta > 0$, there exists a sufficiently small ε such that lowering p_i to $c - \tilde{c}_i - \varepsilon$ changes the second-stage profit by at most δ, lowers profit from current sales by at most δ (because if there were any current sales to begin with, they can only have been at profit $c - \tilde{c}_i$), and also generates $(x - \delta) \cdot (c - \tilde{c}_i - \delta)$ additional profits from new sales. As $c - \tilde{c}_i > 0$ and $z > 0$ by hypothesis, there is clearly a sufficiently small δ such that this results in a strictly positive change in utility, and p cannot be a first-stage equilibrium for Q. □

We now establish Lemmas 6 and 7, which jointly imply Proposition 1. Lemma 6 shows that when sales occur, all computational power must be purchased from Miner 1, and at a price $p_1 \leq \tilde{c}_2 - \tilde{c}_1$. Lemma 7 shows that Miner 1 never wishes to lower p_1 below $\tilde{c}_2 - \tilde{c}_1$.

Lemma 6. *If (p, Q) is an equilibrium of the transferable game where sales occur, then $p_1 \leq \tilde{c}_2 - \tilde{c}_1$, and $Q_{\ell j}(p) = 0$ for all $j \neq 1$ and all ℓ.*

Proof. Clearly, if Q is not a second-stage equilibrium, then (p, Q) is not an equilibrium. So assume for the rest of this proof that Q is a second-stage equilibrium.

Assume for contradiction that sales occur, but $p_1 > \tilde{c}_2 - \tilde{c}_1$. We claim that either Miner 1 or Miner 2 must satisfy the hypotheses of Lemma 5, and therefore p is not a first-stage equilibrium for Q. Indeed, because $p_1 > \tilde{c}_2 - \tilde{c}_1$, we must either have $\tilde{c}_1 + p_1 > c$ (in which case no one purchases from Miner 1, because Q is a second-stage equilibrium), or we must have $\tilde{c}_2 < \tilde{c}_1 + p_1 = c$, in which case Lemma 5 asserts that Miner 2 is not setting price p_2 optimally (because Q is lexicographic, all sales go to Miner 1, and some sales occur by assumption). Therefore, unless $p_1 = \tilde{c}_2 - \tilde{c}_1$, p cannot be a lexicographic first-stage equilibrium.

Finally, observe that if $p_1 \leq \tilde{c}_2 - \tilde{c}_1$, then $c_j(p) = c(p) = p_1 + \tilde{c}_1$ for all j. Therefore, because Q is lexicographic, all computational power is purchased from Miner 1.

The proof of Lemma 6 is complete. We quickly note that if we remove the lexicographic requirements from equilibria, the lemma statement would instead read that in any equilibrium where sales occur, there exists a miner i for which $\tilde{c}_i = \tilde{c}_1$ and $p_i = \tilde{c}_2 - \tilde{c}_1$, another miner j for which $\tilde{c}_j = \tilde{c}_2$ and $p_j = 0$, and all miners purchase hardware from a miner i

for which $\tilde{c}_i = \tilde{c}_1$ and $p_i = \tilde{c}_2 - \tilde{c}_1$ (the proof outline is identical, but requires more corner cases). □

Next, we argue that Miner 1 will never choose to set $p_1 < \tilde{c}_2 - \tilde{c}_1$. This part of the proof appeals to the particular reward function. The proof of Lemma 7 (along with later proofs) appeal to the change in $U_1(p, Q)$ with respect to p_1 (keeping Q as a lexicographic second-stage equilibrium, and $p_2, \ldots, p_n$ fixed). We will only be interested in changing p_1 within the range that this also changes c^*. That is, when p_1 is sufficiently large, no one purchases hardware from Miner 1, so their utility doesn't change. When p_1 is small enough, however, that some miner i might purchase from Miner 1, changing p_1 changes c_i and therefore also changes c^*. In this range, c^* is strictly increasing in p_1, and therefore there is a well-defined inverse function so that we can write p_1 as a function of c^* (think of this as "what price p_1 would Miner 1 have to set in order to induce c^*). It turns out that all calculations are drastically simpler when we look at the change in $U_1(p, Q)$ *as a function of* c^*. So formally, these proofs will consider $\frac{\partial U_1(p,Q)}{\partial c^*}$, and we will think of $U_1(\cdot, Q)$ as being a function of $c^*, p_2, \ldots, p_n$ (knowing that we can write p_1 as a function of these variables to induce c^*, and then write $U_1(p, Q)$ as a function of p, Q).

Lemma 7. *Let Q be a second-stage equilibrium, and p be a first-stage equilibrium for Q. Then $p_1 \geq \tilde{c}_2 - \tilde{c}_1$.*

Proof. To prove this, we write the utility enjoyed by Miner 1 using Lemma 4, and reason about this as a function of c^*. In the range where $p_1 \in [0, \tilde{c}_2 - \tilde{c}_1)$, observe that $c_i = p_1 + \tilde{c}_1$ for all $i > 1$, and $c_1 = \tilde{c}_1$. Therefore (see Lemma 3), $c^* = c_2 + \tilde{c}_1/(n-1) = \tilde{c}_1 + p_1 + \tilde{c}_1/(n-1)$, and when $p_1 \in [0, \tilde{c}_2 - \tilde{c}_1)$, $c^* \in [\tilde{c}_1 + \tilde{c}_1/(n-1), \tilde{c}_2 + \tilde{c}_1/(n-1))$ (and c^* is monotone increasing in p_1). We will argue that in this range, the derivative of Miner 1's utility with respect to c^* is strictly positive, and therefore no $p_1 < \tilde{c}_2 - \tilde{c}_1$ can be a best response (because Miner 1 would prefer to strictly increase p_1 to strictly increase c^*).

Recall by Lemma 4 that Miner 1's utility at (p, Q) (if it is an equilibrium) is $(\max\{1 - c_1/c^*, 0\})^2 + \sum_j q_{j1} p_i$. Observe next that $c^* \geq c_1$, so $\max\{1 - c_1/c^*, 0\} = 1 - c_1/c^*$. Moreover, recall that when $p_1 \leq \tilde{c}_2 - \tilde{c}_1$, $c_i = p_1 + \tilde{c}_1$ for all $i > 1$, and *all computational power acquired by any miner $\neq 1$ is purchased from Miner* 1. Therefore, because a total of $1/c^*$ power is acquired, and Miner 1 acquires a total of $(1 - c_1/c^*)/c^*$ computational for themselves, it must be that $c_1/(c^*)^2$ units are purchased in total from Miner 1. This means that:

$$\begin{aligned} U_1(p, Q) &= (1 - c_1/c^*)^2 + p_1 \cdot c_1/(c^*)^2 \\ &= (1 - c_1/c^*)^2 + (c^* - c_1) \cdot c_1/(c^*)^2 - (c^* - c_1 - p_1) \cdot c_1/(c^*)^2 \\ &= (1 - c_1/c^*) - (c_1/(n-1)) \cdot c_1/(c^*)^2. \end{aligned}$$

Note that the final equality follows from recalling that $c^* = c_1 + p_1 + c_1/(n-1)$. Taking now a derivative with respect to c^*, we see that:

$$\frac{\partial U_1(p,Q)}{\partial c^*} = c_1/(c^*)^2 + \frac{2c_1^2}{(n-1)\cdot(c^*)^3} > 0.$$

This means that it is strictly in Miner 1's interest to increase c^* at least until $c^* = \tilde{c}_2 + \tilde{c}_1/(n-1)$, which corresponds to increasing p_1 at least until $\tilde{c}_2 - \tilde{c}_1$. □

Having established Proposition 1, we now establish Proposition 2, which states that sales must occur in equilibrium (except for one degenerate case). We establish in Lemma 8 that it is impossible for no sales to occur if Miner n acquires non-zero computational power.

Lemma 8. *If (p,Q) is an equilibrium of the transferable game where no sales occur, then $Q_n(p) = 0$.*

Proof. Assume for contradiction that (p,Q) is an equilibrium in which no sales occur, yet $q_n > 0$. Observe first that because no sales occur, then $c_i = \tilde{c}_i$ for all i. In particular, this means that Miner 1 could set $p_1 = \tilde{c}_n - \tilde{c}_1$. This maintains $c_i = \tilde{c}_i$ for all i, and thus does not change c^*. However, it does cause Miner n to now purchase $q_n > 0$ units from Miner 1 (whereas previously no sales occurred). If $\tilde{c}_n > \tilde{c}_1$, then this strictly improves Miner 1's utility (because c^*, c_i are the same, but they get increased profits from sale). If $\tilde{c}_n = \tilde{c}_1$, then all p_1 yield equal payoff for Miner 1, so for p to be lexicographic we must have $p_1 = 0$, and for Q to be lexicographic we must have all miners purchase from Miner 1.

The proof of Lemma 8 is complete. We quickly note that if we remove the lexicographic requirements from equilibria, the lemma statement would instead read that in any equilibrium where $c_i = \tilde{c}_i$ for all i, either $q_n = 0$ or $\tilde{c}_1 = \tilde{c}_n$ (and the proof is identical). □

Finally, we establish that it is impossible for no sales to occur, yet Miner n acquires no computational power. This is perhaps the most surprising lemma of the paper, as it establishes that Miner 1 *always* wishes to set a price low enough for a new competitor to enter the market, even though this will result in a lower market share for Miner 1 in the second-stage. Specifically, observe that after fixing all quantities acquired by miners 2 through $n-1$, Miner n will choose to pay profit at most $c^* - c_1$ to Miner 1 per unit of power. Alternatively, Miner 1 could themselves acquire computational power at marginal reward c^* and cost c_1, for marginal profit $c^* - c_1$. So in particular, once the quantities $q_2, \ldots, q_{n-1}$ are fixed, Miner 1 has nothing to gain by selling to Miner n instead of just acquiring power for themselves. What drives Lemma 9 is that $q_2, \ldots, q_{n-1}$ are not fixed as a function of $c_1, \ldots, c_{n-1}$, but vary in response to c_n. Therefore, Miner 1 might profit by lowering c_n, because this in turn lowers $q_2, \ldots, q_{n-1}$ in response (and indeed, this is the case).

Lemma 9. *Let Q be a second-stage equilibrium, and p be a first-stage equilibrium for Q, and let no sales occur in (p,Q). Then unless $\tilde{c}_1 = \tilde{c}_2 \leq \tilde{c}_3/2$, $q_n > 0$.*

Proof. If no sales occur, then $c_i = \tilde{c}_i$ for all i. If $q_n = 0$, then $\tilde{c}_n \geq \tilde{c}^*$. Let k be the maximally-indexed miner with $q_k > 0$ (equivalently, $\tilde{c}_k < \tilde{c}^*$). Observe, then, that the payoff for Miner 1 for any strategy which sets price $p_1 \in (\tilde{c}_k - \tilde{c}_1, \tilde{c}^* - \tilde{c}_1)$ causes exactly miners $i > k$ to purchase from Miner 1, and to have $c_i = c_1 + p_1$. Therefore, Miner 1's payoff, as a function of p_1 in this range, is:

$$\begin{aligned} U_1(p, Q) &= (1 - c_1/c^*)^2 + (n-k)p_1 \cdot (1 - \frac{c_1 + p_1}{c^*})/c^* \\ &= (1 - c_1/c^*)^2 + p_1 \cdot (1 - \sum_{j=1}^{k} 1 - \frac{c_j}{c^*})/c^* \\ &= (1 - c_1/c^*)^2 + p_1 \cdot (-(k-1) + \sum_{j=1}^{k} \frac{c_j}{c^*})/c^*. \end{aligned}$$

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 1 through k are purchased from Miner 1 (and that $\sum_{j=1}^{k} 1 - c_j/c^*$ units are owned by miners 1 through k).

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 1 would strictly profit by lowering p_1 to induce a $c^* < \tilde{c}^*$.

$$\frac{\partial U_1(p, Q)}{\partial c^*} = \frac{2(1 - c_1/c^*)c_1}{(c^*)^2} + \frac{\partial p_1}{\partial c^*} \cdot (-(k-1) + \sum_{j=1}^{k} \frac{c_j}{c^*})/c^* + \frac{(k-1)p_1}{(c^*)^2} - \frac{2p_1 \sum_{j=1}^{k} c_j}{(c^*)^3}.$$

To evaluate the derivative at $c^* = \tilde{c}^*$, observe first that $(-(k-1) + \sum_{j=1}^{k} \frac{c_j}{\tilde{c}^*}) = 0$. The reason for this is because exactly miners $1, \ldots, k$ participate 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_1, because there are no sales at (p, Q)). Similarly, observe that the price which induces $c^* = \tilde{c}^*$ is exactly $p_1 = \tilde{c}^* - \tilde{c}_1$ (any lower price would strictly decrease c^*). So after these substitutions, we get that:

$$\begin{aligned}
\frac{\partial U_1(p,Q)}{\partial \tilde{c}^*}(\tilde{c}^*) &= \frac{2(1-c_1/\tilde{c}^*)c_1}{(\tilde{c}^*)^2} + \frac{(k-1)(\tilde{c}^*-c_1)}{(\tilde{c}^*)^2} - \frac{2(\tilde{c}^*-c_1)\sum_{j=1}^k c_j}{(\tilde{c}^*)^3} \\
&= \frac{2(1-c_1/\tilde{c}^*)c_1}{(\tilde{c}^*)^2} + \frac{(k-1)(\tilde{c}^*-c_1)}{(\tilde{c}^*)^2} - \frac{2(\tilde{c}^*-c_1)(k-1)\tilde{c}^*}{(\tilde{c}^*)^3} \\
&= \frac{2(\tilde{c}^*-c_1)c_1 + (k-1)(\tilde{c}^*-c_1)\tilde{c}^* - 2(\tilde{c}^*-c_1)(k-1)\tilde{c}^*}{(\tilde{c}^*)^3} \\
&= \frac{2(\tilde{c}^*-c_1)c_1 - (\tilde{c}^*-c_1)(k-1)\tilde{c}^*}{(\tilde{c}^*)^3} \\
&= \frac{(\tilde{c}^*-c_1)}{(\tilde{c}^*)^3}\cdot(2c_1-(k-1)\tilde{c}^*) < 0.
\end{aligned}$$

The first equality uses Lemma 3 to replace $\sum_{j=1}^k c_j$ with $(k-1)c^*$, and the rest follow from algebraic manipulation. The final inequality follows as $(k-1)\tilde{c}^* = \sum_{j=1}^k c_i > 2c_1$, unless $c_1 = c_2$ and $k = 2$ (which is a special case where the lemma fails to hold). The only way to have $\tilde{c}_1 = \tilde{c}_2$ and $k = 2$ is to also have $\tilde{c}_3 \geq 2\tilde{c}_1$, as this is necessary to have only miners 1 and 2 participate in equilibrium. Observe that this implies that Miner 1 would strictly profit by decreasing c^* from $\tilde{c}^*$, and therefore it is not an equilibrium.

Having completed the proof, we note that we can replace the subscript of 1 with i everywhere, so long as $\tilde{c}_i < c^*$, and arrive at the same conclusion as long as $2c_i < (k-1)\tilde{c}^*$. Because $(k-1)\tilde{c}^* = \sum_{j=1}^k c_j$ (see Lemma 3), this holds for any miner i with $q_i > 0$ unless $k = 2$. This implies that starting from a no-sales outcome of the transferable game, as long as at least three miners participate in the second stage, any of these miners would wish to deviate and sell to a new entrant. □

At this point, Proposition 2 follows from Lemma 8 and 9. For the sake of completeness, we also analyze the corner case where $\tilde{c}_1 = \tilde{c}_2 \leq \tilde{c}_3/2$.

Lemma 10. *If $\tilde{c}_1 = \tilde{c}_2 \leq \frac{1}{2}\tilde{c}_3$, then (p, Q^*) is an equilibrium of the transferable game either for when all miners set price 0, or when $p_1 = p_2 = \tilde{c}_1$ with all other $p_i = 0$. There are no other equilibria.*

Proof. It is trivial to check that $p_i = 0$ for all i is an equilibrium, as once one miner in $\{1, 2\}$ sets a price of 0, this fixes the payoff for all other miners, independent of their price.

Lemmas 6, 7, and 8 together imply that any equilibrium in which sale occurs must be with $p_i = 0$ for all i. So the only other possible equilibrium is one in which no sale occurs. When the lemma hypotheses are violated, Lemma 9 establishes that this is not an equilibrium. In the special case when the lemma hypotheses hold, we now show that this is an equilibrium.

To this end, observe that when $\tilde{c}_1 = \tilde{c}_2 \leq \tilde{c}_3/2$, that $\tilde{c}^* = 2\tilde{c}_1$, and therefore only $k = 2$ miners have $q_i > 0$ at these costs. Therefore, Miner 1's utility after inducing any $c^* \in (\tilde{c}_1, 2\tilde{c}_1]$ is (exactly as in Lemma 9):

$$\begin{aligned} U_1(p, Q) &= (1 - c_1/c^*)^2 + p_1 \cdot (-(k-1) + \sum_{j=1}^{k} \frac{c_j}{c^*})/c^* \\ &= (1 - c_1/c^*)^2 + p_1 \cdot (-1 + 2c_1/c^*)/c^*. \end{aligned}$$

We again are interested in taking the derivative with respect to c^*, but this time we wish to evaluate it at all c^* in the relevant range (rather than just $\tilde{c}^*$).

$$\begin{aligned} \frac{\partial U_1(p, Q)}{\partial c^*} &= \frac{2(1 - c_1/c^*)c_1}{(c^*)^2} + \frac{\partial p_1}{\partial c^*} \cdot (-1 + 2c_1/c^*)/c^* + \frac{p_1}{(c^*)^2} - \frac{4p_1 c_1}{(c^*)^3} \\ &= \frac{2c_1 c^* - 2c_1^2 + p_1 c^* - 4p_1 c_1}{(c^*)^3} + \frac{\partial p_1}{\partial c^*} \cdot (-1 + 2c_1/c^*)/c^*. \end{aligned}$$

To continue, we now must relate c^* to p_1, as we care about $\frac{\partial p_1}{\partial c^*}$. Observe that there are $n - 2$ miners with cost $\tilde{c}_1 + p_1$, and two with cost $\tilde{c}_1$. So for an equilibrium where all miners participate, we have $(n-1)c^* = 2\tilde{c}_1 + (n-2)(\tilde{c}_1 + p_1)$, and:

$$p_1 = (1 + \frac{1}{n-2})c^* - (1 + \frac{2}{n-2})\tilde{c}_1.$$

This therefore implies that:

$$\frac{\partial p_1}{\partial c^*} = (1 + \frac{1}{n-2}).$$

We can now continue evaluating $\frac{\partial U_1(p,Q)}{\partial c^*}$:

$$
\begin{aligned}
\frac{\partial U_1(p,Q)}{\partial c^*} &= \frac{2c_1c^* - 2c_1^2 + p_1c^* - 4p_1c_1}{(c^*)^3} + \frac{\partial p_1}{\partial c^*} \cdot (-1 + 2c_1/c^*)/c^* \\
&= \frac{2c_1c^* - 2c_1^2 + (1+\frac{1}{n-2})(c^*)^2 - (1+\frac{2}{n-2})c_1c^* - 4(1+\frac{1}{n-2})c^*c_1 + 4(1+\frac{2}{n-2})c_1^2}{(c^*)^3} \\
&\quad + (1+\frac{1}{n-2}) \cdot (-1 + 2c_1/c^*)/c^* \\
&= \frac{-(3+\frac{6}{n-2})c_1c^* + (2+\frac{8}{n-2})c_1^2 + (1+\frac{1}{n-2})(c^*)^2 - (1+\frac{1}{n-2})(c^*)^2 + 2(1+\frac{1}{n-2})c_1c^*}{(c^*)^3} \\
&= \frac{-(1+\frac{4}{n-2})c_1c^* + (2+\frac{8}{n-2})c_1^2}{(c^*)^3} \\
&= \frac{c_1}{(c^*)^3} \cdot \left(-(1+\frac{4}{n-2})c^* + (2+\frac{8}{n-2})c_1 \right) \\
&= \frac{c_1}{(c^*)^3} \cdot \left(1 + \frac{4}{n-2} \right) \cdot (2c_1 - c^*) \\
&> 0, \text{ when } c^* < \tilde{c}^* = 2c_1.
\end{aligned}
$$

This implies that Miner 1 is strictly better off inducing a $c^* = \tilde{c}^*$ than any other $c^* < \tilde{c}^*$, and therefore it is indeed an equilibrium for Miner 1 (and also Miner 2) to set price c_1. This is the only lexicographic p which induces no sales. Without the lexicographic requirement, any prices which induce no sales would be an equilibrium (which are any prices $\geq \tilde{c}_1$). □

C Transferable Game with Personalized Pricing

We now consider the transferable game with personalized pricing, and conclude that Theorem 2 holds in this model as well. All definitions are identical to the transferable game, except that prices set in the first-stage may be personalized. We'll refer to this as the transferable game with personalized pricing. Specifically the two rounds proceed as:

1. Each miner i chooses price for Miner j $p_{ji} \geq 0$ at which they are willing to sell their computational power.

2. Each miner i chooses the quantity q_{ij} of computational power sourced from miner j. We let $q_i = \sum_j q_{ij}$ be the total computational power of miner j.

Given prices $p = \{p_{ij}\}$ and purchase quantities $q = \{q_{ij}\}$, the utility of miner i is

$$
U_i(p,q) = \frac{q_i}{\sum_j q_j} - \sum_j q_{ij}\tilde{c}_j + \sum_{j \neq i} p_{ji}q_{ji} - \sum_{j \neq i} p_{ij}q_{ij}.
$$

Again, each miner gets its share of the pool of rewards, pays for the direct costs of the computational power that it uses, receives profit from any computational power that it sells, and pays competitors for computational power purchased from them. The definition of a second-stage equilibrium remains identical (including lexicographic). The definition of a first-stage equilibrium also remains identical (after modifying "$p_j \leq p'_j$" in the definition of lexicographic to read "$p_{ij} \leq p'_{ij}$, $\forall i$").

The main result of this section is the following theorem, extending Theorem 2 to the personalized pricing game as well.

Theorem 3. *In the transferable game with personalized pricing, it is an equilibrium for Miner 1 to set $p_{j1} = \tilde{c}_2 - \tilde{c}_1$ for all j, and all other miners i to set $p_{ji} = 0$ for all j. This is the only equilibrium unless $\tilde{c}_1 = \tilde{c}_2 \leq \tilde{c}_3/2$.*

Many of the same lemmas towards Theorem 2 also hold, verbatim (or nearly verbatim), towards Theorem 3. Specifically, we repeat statements of these modified lemmas below, and note that their proofs are identical to a corresponding lemma in Appendix B.

Lemma 11. *If (p, Q) is an equilibrium of the transferable game with personalized pricing, then:*

$$U_i(p, Q) = \left(\max\{1 - c_i(p)/c^*(p), 0\}\right)^2 + \sum_j Q_{ji}(p) \cdot p_{ji}.$$

The proof of Lemma 11 is identical to that of Lemma 4. The only difference in the statement/proof is updating notation of p_i to p_{ji}.

Lemma 12. *Consider any $Q(\cdot)$ which is a second-stage equilibrium. Then if there exists a miner i and miner j such that $Q_{j\ell}(p) > 0$ for some $\ell \neq i$, but $c_j(p) > \tilde{c}_i$ (i.e. some miner j purchases from a miner $\ell \neq i$ at a price strictly larger than $\tilde{c}_i$), then p is not a first-stage equilibrium for Q.*

The proof of Lemma 12 is even simpler than that of Lemma 5, as Miner i may now simply lower their price p_{ji}.

Lemma 13. *If (p, Q) is an equilibrium of the transferable game with personalized pricing where $Q_j(p) > 0$, then $p_{j1} \leq \tilde{c}_2 - \tilde{c}_1$, and $Q_{\ell j}(p) = 0$ for all $j \neq 1$ and all ℓ.*

The proof of Lemma 13 is identical to that of Lemma 6. The only difference in the statement/proof is updating notation from "sales occur" to "$Q_j(p) > 0$" (for a particular miner j).

Lemma 14. *Let Q be a second-stage equilibrium, and p be a first-stage equilibrium for Q. Then $p_{j1} \geq \tilde{c}_2 - \tilde{c}_1$ for all j.*

The proof of Lemma 14 follows identical calculations to Lemma 7. The only difference in the statement/proof is replacing p_1 with p_{j1}.

Lemma 15. *If (p, Q) is an equilibrium of the transferable game with personalized pricing where no sales occur, then $Q_n(p) = 0$.*

Th proof of Lemma 15 is identical to Lemma 8, and the statement is identical.

The above lemmas rule out equilibria aside from the one in Theorem 3 in which sales occur, and in which no sales occur but $Q_n(p) > 0$. Lemma 9, which was the most surprising step of Theorem 2, is the only lemma missing from the above list. Our last step is to prove a generalization of Lemma 9 for the personalized pricing model.

Lemma 16 below highlights the following surprising fact: *even if* all participating miners $\neq 1$ are currently paying per-unit price (which must be between 0 and $\tilde{c}_i - \tilde{c}_1$) to Miner 1, it is *still* in Miner 1's interest to lower c^* to recruit a new miner to the market. Lemma 16 is a strict generalization of Lemma 9, but the proof is significantly more involved.

Lemma 16. *In the transferable game with personalized pricing, let Q be a second-stage equilibrium, and p be a first-stage equilibrium for Q. Then unless $\tilde{c}_1 = \tilde{c}_2 \leq \tilde{c}_3/2$, $q_i > 0$ for all i.*

In particular, for any set of miners M with $q_i = 0$ for all $i \in M$, there exists an $\varepsilon > 0$ such that Miner 1 strictly prefers to lower p_i by ε for all $i \in M$ (and keep all other prices the same).

Proof. The first part of the proof follows similarly to Lemma 9. To simplify notation, we relabel the miners by effective cost c_i, instead of $\tilde{c}_i$ (so that again $c_1 \leq \ldots \leq c_n$), and let k be the maximally-indexed miner with $q_k > 0$. We consider now a deviation where Miner 1 sells hardware to any set of $m > 0$ miners outside these k at price $p = c^* - c_1 - \varepsilon$, for sufficiently small ε, while keeping all p_{ji} the same for all $i \leq k$. In this range, we can write the utility of Miner 1 as follows (because the only impact this has on miners 2 through k is via c^*, and the m other miners also have their c_i changing with p):

$$\begin{aligned} U_1(p, Q) &= (1 - c_1/c^*)^2 + \sum_{j=2}^{k} p_j \cdot (1 - c_j/c^*)/c^* + mp \cdot (1 - \frac{c_1 + p}{c^*})/c^* \\ &= (1 - c_1/c^*)^2 + \sum_{j=2}^{k} p_j \cdot (1 - c_j/c^*)/c^* + p \cdot (1 - \sum_{j=1}^{k} 1 - \frac{c_j}{c^*})/c^* \\ &= (1 - c_1/c^*)^2 + \sum_{j=2}^{k} p_j \cdot (1 - c_j/c^*)/c^* + mp \cdot (-(k-1) + \sum_{j=1}^{k} \frac{c_j}{c^*})/c^*. \end{aligned}$$

The first equality again follows by observing that all computational power not purchased by miners 1 through k is purchased by the remaining m miners, and profit p per unit goes to Miner 1. We are again interested in first taking the derivative with respect to c^*, and

then evaluating the derivative at our starting c_0^*, where c_0^* denotes c^* for the proposed (p, Q).

$$\frac{\partial U_1(p,Q)}{\partial c^*} = \frac{2(1-c_1/c^*)c_1}{(c^*)^2} + \left(\sum_{j=2}^{k} 2p_j c_j/(c^*)^3 - p_j/(c^*)^2\right) + \frac{\partial p}{\partial c^*} \cdot (-(k-1) + \sum_{j=1}^{k} \frac{c_j}{c^*})/c^* + \frac{(k-1)p}{(c^*)^2} - \frac{2p\sum_{j=1}^{k} c_j}{(c^*)^3}.$$

To evaluate the derivative at $c^* = c_0^*$, observe first that $(-(k-1) + \sum_{j=1}^{k} \frac{c_j}{\tilde{c}_0^*}) = 0$. The reason for this is because exactly miners $1, \ldots, k$ participate in equilibrium at the initial costs, and therefore they must be responsible for the entire market share. Similarly, observe that the price which induces $c^* = c_0^*$ is exactly $p = c_0^* - \tilde{c}_1$ (any lower price would strictly decrease c^*). So after these substitutions, we get the following chain of equalities. Below, we first simplify the term common to the term above and the non-personalized-pricing case (taking exactly the same algebraic manipulations), and then simplify the new term here.

$$\begin{aligned}
\frac{\partial U_1(p,Q)}{\partial c^*}(c_0^*) &= \frac{2(1-c_1/c_0^*)c_1}{(c_0^*)^2} + \frac{(k-1)(c_0^*-c_1)}{(c_0^*)^2} - \frac{2(c_0^*-c_1)\sum_{j=1}^{k} c_j}{(c_0^*)^3} \\
&\quad + \left(\sum_{j=2}^{k} 2p_j c_j/(c_0^*)^3 - p_j/(c_0^*)^2\right) \\
&= \frac{2(1-c_1/c_0^*)c_1}{(c_0^*)^2} + \frac{(k-1)(c_0^*-c_1)}{(c_0^*)^2} - \frac{2(c_0^*-c_1)(k-1)c_0^*}{(c_0^*)^3} \\
&\quad + \left(\sum_{j=2}^{k} 2p_j c_j/(c_0^*)^3 - p_j/(c_0^*)^2\right) \\
&= \frac{2(c_0^*-c_1)c_1 + (k-1)(c_0^*-c_1)c_0^* - 2(c_0^*-c_1)(k-1)c_0^*}{(c_0^*)^3} \\
&\quad + \left(\sum_{j=2}^{k} 2p_j c_j/(c_0^*)^3 - p_j/(c_0^*)^2\right) \\
&= \frac{2(c_0^*-c_1)c_1 - (c_0^*-c_1)(k-1)c_0^*}{(c_0^*)^3} \\
&\quad + \left(\sum_{j=2}^{k} 2p_j c_j/(c_0^*)^3 - p_j/(c_0^*)^2\right)
\end{aligned}$$

$$= \frac{(c_0^* - c_1)}{(c_0^*)^3} \cdot (2c_1 - (k-1)c_0^*) + \left(\sum_{j=2}^{k} 2p_j c_j/(c_0^*)^3 - p_j/(c_0^*)^2 \right)$$

$$= \frac{(c_0^* - c_1)}{(c_0^*)^3} \cdot (2c_1 - (k-1)c_0^*) + \left(\sum_{j=2}^{k} \frac{p_j}{(c_0^*)^3} \cdot (2c_j - c_0^*) \right)$$

$$\leq \frac{(c_0^* - c_1)}{(c_0^*)^3} \cdot (2c_1 - (k-1)c_0^*) + \max \left\{ \sum_{j=2}^{k} \frac{c_j - c_1}{(c_0^*)^3} \cdot (2c_j - c_0^*), 0 \right\}$$

The final inequality follows by simply observing that if $2c_j < c_0^*$, then the the term is maximized at $p_j = 0$. Otherwise, the term is maximized when p_j is as large as possible, which is $c_j - c_1$. From here, we continue with an observation.

Observe that for any $j > 1$, $2c_j \geq c_0^*$ (Lemma 3). Therefore, the max is superfluous, and we can continue:

$$\frac{\partial U_1(p,q)}{\partial c^*}(c_0^*) \leq \frac{(c_0^* - c_1)}{(c_0^*)^3} \cdot (2c_1 - (k-1)c_0^*) + \max \left\{ \sum_{j=2}^{k} \frac{c_j - c_1}{(c_0^*)^3} \cdot (2c_j - c_0^*), 0 \right\}$$

$$= \frac{(c_0^* - c_1)}{(c_0^*)^3} \cdot (2c_1 - (k-1)c_0^*) + \left(\sum_{j=2}^{k} \frac{c_j - c_1}{(c_0^*)^3} \cdot (2c_j - c_0^*) \right)$$

Next, we seek to understand how large the term $\left(\sum_{j=2}^{k} \frac{c_j - c_1}{(c_0^*)^3} \cdot (2c_j - c_0^*) \right)$, can possibly be. To this end, we first treat c_0^* and c_1 as fixed, and ask what $c_2, \ldots, c_k$ maximize this term subject to the constraints implied by c_1 and c_0^* (i.e. that each $c_j \in [c_1, c_0^*]$, and that $\sum_{j=2}^{k} c_j = (k-1)c_0^* - c_1$). Observe next that $\left(\sum_{j=2}^{k} \frac{c_j - c_1}{(c_0^*)^3} \cdot (2c_j - c_0^*) \right)$ is strictly convex in each of $c_2, \ldots, c_k$ (treating c_0^*, c_1 as fixed). In particular, it is a sum of single-variate convex functions in each variable separately. This means that the term is maximized at an extreme point (in particular, it means that the term is strictly increased by increasing c_j by ε, and decreasing c_i by ε, for sufficiently small ε, whenever $c_j < c_0^*$ and $c_i > c_1$). This means that we can upper bound the term by evaluating it only at extreme points which for which this adjustment is infeasible. Depending on the ratio of c_1 to c_0^*, this extreme point will have two possibilities:

- If $c_0^* < 2c_1$, then the only point which cannot further undergo this operation has $c_2 = c_1$, and $c_3 < c_0^*$.
- If $c_0^* \geq 2c_1$, then the point which cannot further undergo this operation has $c_2 = c_0^* - c_1$, and $c_j = c_0^*$ for all $j > 2$.

We first consider the case $c_0^* < 2c_1$. Recalling that we must have $\sum_{j=2}^{k} c_j = (k-1)c_0^* - c_1$, this means that the term is maximized when all terms are either c_1 or c_0^* (if this is even

feasible, due to the fact that an integer number of terms must take each value). If exactly ℓ terms are c_1 (note that this upper bounds the term in question, even if $\ell \notin \mathbb{N}$), and the remaining $k-\ell-1$ are c_0^*, we get:

$$\begin{aligned} &\ell c_1 + (k-1-\ell)c_0^* = (k-1)c_0^* - c_1 \\ \Rightarrow &\ell = \frac{c_1}{c_0^* - c_1}. \end{aligned}$$

$$\begin{aligned} \Rightarrow \left(\sum_{j \le k} \frac{c_j - c_1}{(c_0^*)^3} \cdot (2c_j - c_0^*)\right) &\le \ell \cdot 0 + (k-\ell-1) \cdot (c_0^* - c_1) \cdot c_0^*/(c_0^*)^3 \\ &= \frac{(k-1)c_0^* - kc_1}{c_0^* - c_1} \cdot (c_0^* - c_1) \cdot c_0^*/(c_0^*)^3 \\ &= \left((k-1)(c_0^*)^2 - kc_1c_0^*\right)/(c_0^*)^3. \end{aligned}$$

Combining this with our previous bound, we then get:

$$\begin{aligned} \frac{\partial U_1(p,q)}{\partial c^*}(c_0^*) &\le \frac{(c_0^* - c_1)}{(c_0^*)^3} \cdot (2c_1 - (k-1)c_0^*) + \left(\sum_{j=2}^{k} \frac{c_j - c_1}{(c_0^*)^3} \cdot (2c_j - c_0^*)\right) \\ &\le \frac{(c_0^* - c_1)}{(c_0^*)^3} \cdot (2c_1 - (k-1)c_0^*) + \left((k-1)(c_0^*)^2 - kc_1c_0^*\right)/(c_0^*)^3 \\ &= \frac{2c_1c_0^* - 2c_1^2 - (k-1)(c_0^*)^2 + (k-1)c_1c_0^* + (k-1)(c_0^*)^2 - kc_1c_0^*}{(c_0^*)^3} \\ &= \frac{c_1c_0^* - 2c_1^2}{(c_0^*)^3} \\ &= \frac{c_1(c_0^* - 2c_1)}{(c_0^*)^3} \\ &< 0, \text{ when } c_0^* < 2c_1. \end{aligned}$$

Therefore, we may conclude that whenever $c_0^* < 2c_1$, Miner 1 would strictly prefer to lower their price to any number of miners who currently are not participating. We next need to consider the case where $c_0^* \ge 2c_1$. Here, we get (note that some inequalities hold only when $c_0^* \ge 2c_1$, which is the case we're considering):

$$\frac{\partial U_1(p,q)}{\partial c^*}(c_0^*) \le \frac{(c_0^* - c_1)}{(c_0^*)^3} \cdot (2c_1 - (k-1)c_0^*) + \left(\sum_{j=2}^{k} \frac{c_j - c_1}{(c_0^*)^3} \cdot (2c_j - c_0^*)\right)$$

$$
\begin{aligned}
&\leq \frac{(c_0^* - c_1)}{(c_0^*)^3} \cdot (2c_1 - (k-1)c_0^*) + \frac{(c_0^* - 2c_1) \cdot (c_0^* - 2c_1) + (k-2) \cdot (c_0^* - c_1) \cdot c_0^*}{(c_0^*)^3} \\
&= \frac{(k+1)c_1 c_0^* - 2c_1^2 - (k-1)(c_0^*)^2 + (k-1)(c_0^*)^2 + 4c_1^2 - (k+2)c_1 c_0^*}{(c_0^*)^3} \\
&= \frac{2c_1^2 - c_1 c_0^*}{(c_0^*)^3} \\
&= \frac{c_1(2c_1 - c_0^*)}{(c_0^*)^3}.
\end{aligned}
$$

Above, the second inequality follows from evaluating at the extreme point with $c_2 = c_0^* - c_1$ and $c_j = c_0^*$ for all $j > 2$. The subsequent inequalities follow from algebraic manipulation. Let us now evaluate the final term. First, it's clear that htis is < 0 whenever $c_0^* > 2c_1$. It is equal to 0 whenever $2c_1 = c_0^*$. But recall that $\left(\sum_{j=2}^{k} \frac{c_j - c_1}{(c_0^*)^3} \cdot (2c_j - c_0^*)\right)$ is *strictly* convex in $c_2, \ldots, c_k$. This means that the *unique* maximum is the extreme point in question, and any choice that is not this extreme point is strictly negative. But observe finally that if $c_0^* = 2c_1$, $c_2 = c_0^* - c_1$, and $c_3 = c_0^*$, then we have exactly our corner case: $c_1 = c_2 \leq c_3/2$.

In conclusion, we have now established that, unless $c_1 = c_2 \leq c_3/2$, Miner 1 would strictly profit by lowering their price to any set of miners who currently have $q_i = 0$. The proof followed mostly the same outline as Lemma 9, but required extra optimization to show that the same conclusion holds even when Miner 1 may be receiving payments from miners with $q_i > 0$.

□

These lemmas together now constite a proof of Theorem 3. Lemmas 13 and 14 establish that there are no equilibria aside from that posed in Theorem 3 where sales occur. Lemmas 15 and 16 establish that there are no equilibria where no sales occur.

D Economies of Scale

In this section, we consider a model with Economies of Scale (EoS). The goal of this section is to focus on concentration of ownership, so we assume that all costs are non-transferable. We call this the *EoS model*.

The setup is almost entirely the same: there are still $n \geq 2$ miners competing for a prize of value 1. Each miner chooses an investment q_i, paying $c_i q_i$ to do so. However, miners now share rewards proportionally to q_i^α, for some parameter $\alpha \geq 1$. That is, Miner i's reward is $q_i^\alpha / \sum_j q_j^\alpha$. We'll use the same notation as the previous section and denote by $x_i(q) = q_i^\alpha / \sum_j q_j^\alpha$ miner i's market share. The miner's utility is still $x_i(q) - c_i q_i$. We proceed with the main theorem statement of this section.

Theorem 4. *Let q be any equilibrium in the EoS model. Then for all i, j (including $i = j$) such that $q_i, q_j > 0$, $x_i(q) \geq 1 - \frac{1}{\alpha}\frac{c_i}{c_j}$.*

Theorem 4 is almost a strict generalization of Corollary 2, except for the assumption that $q_i > 0$. This assumption is necessary, as the conclusion otherwise does not hold (Example 3). Note that the lower bound on $x_i(q)$ is higher than that in Corollary 2, implying greater market concentration. Furthermore, Theorem 4 has bite even with perfectly symmetric costs: taking $i = j$ implies the following.

Corollary 5. *Suppose q is an equilibrium in the EoS model. Then for all i, $q_i > 0 \Rightarrow x_i(q) \geq 1 - 1/\alpha$, and the number of miners with $q_i > 0$ is at most $\frac{\alpha}{\alpha-1}$.*

This implies, for example, that if $\alpha = 1.05$, then at most 21 miners will participate (even if costs are identical). If $\alpha = 1.1$, at most 11 will. By contrast, in the non-transferable model in Section 3, if costs are identical, then all n miners participate.

D.1 Proof of Theorem 4

The proof of Theorem 4 is similar to that of Theorem 1, but does require some new ideas (most notably, Proposition 3). We begin with a generalization of Lemma 2 to the EoS model.

Lemma 17. *Suppose that q is an equilibrium in the EoS model. Then for all i,*

$$c_i q_i = \alpha \cdot x_i(q)(1 - x_i(q)). \tag{10}$$

Moreover, for all i, $c_i \geq \alpha \cdot q_i^{\alpha-1} \cdot \frac{1-x_i(q)}{\sum_j q_j^\alpha}$,[24] and if $q_i > 0$, then $x_i(q) \geq 1 - 1/\alpha$.

Proof. Consider the derivative of $x_i(q)$ with respect to q_i. We have that:

$$\frac{\partial}{\partial q_i} x_i(q) = \frac{\partial}{\partial q_i} \frac{q_i^\alpha}{\sum_j q_j^\alpha} = \frac{\alpha q_i^{\alpha-1}}{\sum_j q_j^\alpha} - \frac{\alpha q_i^{2\alpha-1}}{(\sum_j q_j^\alpha)^2} = \alpha x_i(q)/q_i - \alpha x_i(q))^2/q_i = \alpha \frac{x_i(q) \cdot (1 - x_i(q))}{q_i}.$$

The rest of the lower bound on c_i follows identically as in Lemma 2 and is omitted.

For the final claim, $q_i > 0$ implies that the miner gets non-negative utility from participating, meaning that $x_i(q) \geq c_i q_i = \alpha x_i(q)(1 - x_i(q))$. Rearranging yields $x_i(q) \geq 1 - 1/\alpha$. □

One might hope for an extension of Corollary 4, claiming that $q_i \geq q_j$ in equilibrium implies $c_i \leq c_j$. Interestingly, this is not true in the EoS model (see Example 3). In particular, equilibria exist where less efficient miners participate and more efficient miners *don't participate at all*. However, we show that this is the *only* non-monotonicity that can occur: if both miner i and miner j participate, then $q_i \geq q_j$ does in fact imply $c_i \leq c_j$.

[24] This is implied by Equation (10) when $q_i > 0$, but needs a separate statement when $q_i = 0$.

Proposition 3. *Let q be any equilibrium of the EoS model, and let $q_i \geq q_j > 0$. Then $c_i \leq c_j$, and equality holds if and only if $q_i = q_j$.*

Proof. Lemma 17 states that $c_i q_i = \alpha x_i(q)(1 - x_i(q))$, and analogously for j. If $q_i, q_j > 0$, we can note that $q_i = x_i(q)^{1/\alpha} \left(\sum_j q_i^\alpha\right)^{1/\alpha}$ and divide by $x_i(q)^{1/\alpha}$ to get

$$c_i \left(\sum_j q_j^\alpha\right)^{1/\alpha} = \alpha x_i(q)^{1-1/\alpha}(1 - x_i(q)) = \alpha f(x_i(q)),$$

where we define $f(x) = x^{1-1/\alpha}(1 - x)$.

We know by Lemma 17 that $x_i(q)$ and $x_j(q)$ are both at least $1 - 1/\alpha$. Because $x_i(q) \geq x_j(q) \Leftrightarrow q_i \geq q_j$, to prove the Proposition, it suffices to show that f is decreasing on $(1 - 1/\alpha, 1)$. But

$$\begin{aligned} f'(x) &= (1 - 1/\alpha)x^{-1/\alpha} - (2 - 1/\alpha)x^{1-1/\alpha} \\ &= \frac{2\alpha - 1}{\alpha} x^{-1/\alpha} \left(\frac{\alpha - 1}{2\alpha - 1} - x\right). \end{aligned}$$

Thus, $f'(x) < 0$ for all $x > \frac{\alpha-1}{2\alpha-1}$, and in particular for all $x > \frac{\alpha-1}{\alpha}$, since $\frac{\alpha-1}{\alpha} \geq \frac{\alpha-1}{2\alpha-1}$. □

The remainder proof of the proof of Theorem 4 now follows easily.

Proof of Theorem 4. Let Miner i and Miner j both participate in q, and let $c_i \leq c_j$. Therefore, by Proposition 3, $q_j \leq q_i$. We have from (10) that

$$\frac{c_i}{c_j}\frac{q_i}{q_j} = \frac{x_i(1 - x_i)}{x_j(1 - x_j)}.$$

Solving for $1 - x_i$ and using the fact that $x_i/x_j = q_i^\alpha/q_j^\alpha$, we get:

$$\begin{aligned} (1 - x_i) &= \frac{c_i}{c_j} \cdot \frac{x_j}{x_i} \cdot \frac{q_i}{q_j} \cdot (1 - x_j) = \frac{c_i}{c_j} \cdot \left(\frac{q_j}{q_i}\right)^{\alpha-1} \cdot (1 - x_j) \leq \frac{1}{\alpha} \cdot \frac{c_i}{c_j} \\ &\Rightarrow x_i \geq 1 - \frac{1}{\alpha}\frac{c_i}{c_j} \end{aligned}$$

The final step follows because $1 - x_j \leq \frac{1}{\alpha}$ (Corollary 5) and $q_j/q_i \leq 1$ (hypothesis + Proposition 3). This takes care of the case where $c_i \leq c_j$. When $c_i \geq c_j$, the theorem immediately follows from Lemma 17. □

We now give an example of equilibria in which the most efficient miner chooses not to participate. Intuitively, one can think of this as inefficient incumbents deterring a more efficient potential entrant.

Example 3. *Fix an integer $m \geq 2$, and consider the case with $n > m$ miners, and $\alpha = m/(m-1)$, where $c_1 = f(m) = (1-1/m)^{1-1/m} < 1 = c_2 = c_3 = \cdots = c_n$. Then for any $M \subseteq \{2, 3, \ldots, n\}$ with $|M| = m$, there exists an equilibrium where $q_i = 1/m$ for $i \in M$, and $q_i = 0$ for $i \notin M$. When $m = \alpha = 2$, these are in fact the only equilibria. In all of these equilibria, $q_1 = 0$, despite the fact that c_1 is the lowest cost.*

We also note that if $\alpha > 2$, there is no pure-strategy equilibrium. Pure-strategy equilibria may fail to exist for a similar reason that they fail to exist in an all-pay auction: given any investment profile by others, a miner's optimal choice is either to invest just slightly more than the highest competitor, or nothing at all.

Finally, as mentioned in Section 5, we consider the case where $\alpha < 1$ (diseconomies of scale). In this case, all miners choose $q_i > 0$, regardless of the costs c_i.

Theorem 5. *For $\alpha \leq 1$, there exists a unique pure strategy equilibrium. If $\alpha < 1$, then $q_i > 0$ for all i.*

Proof. Letting S satisfy $S^\alpha = \sum_i q_i^\alpha$, Lemma 17 states that at a pure strategy equilibrium, the following must hold:

$$x_i^{1-1/\alpha}(1 - x_i) = c_i S/\alpha. \tag{11}$$

$$\sum x_i = 1. \tag{12}$$

Thus, we can search for pure strategy equilibria by searching for possible values of S. The quantity $x_i^{1-1/\alpha}(1-x_i)$ is strictly decreasing for $x_i \in (0, 1]$, taking the value zero at $x_i = 1$ and increasing without bound as $x_i \to 0$. To see this, note that the derivative is

$$(1 - 1/\alpha)x^{-1/\alpha}(1-x) - x^{1-1/\alpha} = x^{-1/\alpha}(1 - 1/\alpha - (2 - 1/\alpha)x).$$

If $\alpha \in [1/2, 1]$, then the quantity in parentheses is clearly negative for any $x > 0$. If $\alpha < 1/2$, then $(2-1/\alpha)$ is negative, so the quantity in parentheses is at most $1-1/\alpha-(2-1/\alpha) = -1$.

It follows that for any S, there is a unique x_i satisfying (11), and furthermore that this value is decreasing in S. It follows that there is a unique value of S for which both of the above equations hold. □