A Mean Field Games Model for Cryptocurrency Mining

Zongxi Li* A. Max Reppen* Ronnie Sircar*

December 1, 2019

Abstract

We propose a mean field game model to study the question of how centralization of reward and computational power occur in the Bitcoin-like cryptocurrencies. Miners compete against each other for mining rewards by increasing their computational power. This leads to a novel mean field game of jump intensity control, which we solve explicitly for miners maximizing exponential utility, and handle numerically in the case of miners with power utilities. We show that the heterogeneity of their initial wealth distribution leads to greater imbalance of the reward distribution, or a "rich get richer" effect. This concentration phenomenon is aggravated by a higher bitcoin price, and reduced by competition. Additionally, an advanced miner with cost advantages such as access to cheaper electricity, contributes a significant amount of computational power in equilibrium. Hence, cost efficiency can also result in the type of centralization seen among miners of cryptocurrencies.

1 Introduction

Blockchain technologies serve the purpose of record keeping in a decentralized way. Bitcoin is the famous realization of this idea (see Nakamoto (2008)). Since its creation in January 2009, Bitcoin has grown rapidly. The supply of bitcoins is constantly growing, but limited to 21 million, of which more than 17 million are in circulation now.

In the Bitcoin network, independent "miners" compete for the right to record the next transaction block on the blockchain. They follow proof-of-work protocol and solve math puzzles. Once a miner obtains a solution, the corresponding block is added on top of the blockchain and the miner obtains the reward. The math puzzle is designed such that there is no known better way of solving it than brute force calculation. In other words, the chance of getting the reward is proportional to the computational power or the hash rates that miners can provide. Moreover, the difficulty of the puzzle varies to maintain a consistent solving time, for example 10 minutes. To be specific, if miners can solve the problem in 8 minutes, the system will make the problem harder so that the average time goes back to 10 minutes. In summary, the two important properties are that (1) the probability of obtaining the next reward is proportional to computational efforts and (2) the block rewards appear with a fixed average frequency.

*ORFE Department, Princeton University, Princeton, USA.

Bitcoin is a payment system maintained by a peer-to-peer network. The miners are actually individuals who are dispersed all over the world. They record transactions on the blockchain and achieve the decentralization of the payment system. However, miners have incentives to maximize their own utility. To achieve this, they may increase their computational power to compete for the reward, which can lead to the imbalance of the reward distribution. The empirical analysis in Kondor et al. (2014) shows that the accumulation of bitcoins tends to occur among a small amount of miners, which suggests the centralization of the reward. This raises the following questions: What is the best strategy for miners to maximize their interests? How does the centralization of the reward happen in a decentralized mining activity. What factors have impact on this centralization?

At the end of 2010, the first mining pool "Slush pool" was announced. Miners can join the pool, which collects their computational power to do the mining. Once the pool gets the reward, miners share the profit within it. Nowadays, most computational power comes from mining pools that are controlled by a few companies (see Figure 5). For instance, AntPool and BTC.com are run by Bitmain. Meanwhile, these companies also contribute a significant proportion of hash rates in their own pools. That means a few miners account for a large amount of hash rates in the world. One may ask: What leads to this centralization of computational power? What advantages do those miners have?

1.1 Related literature

Our work is related to the growing literature on cryptocurrencies. A game-theoretic model is developed in Easley et al. (2019) to show the emergence of transaction fees in the Bitcoin payment system. Abadi and Brunnermeier (2018) point out the blockchain trilemma, and analyze when decentralized record-keeping is economically beneficial. Sockin and Xiong (2018) explore a model to study initial coin offerings for new decentralized digital platforms. Cong and He (2019) argue that the blockchain facilitates the creation of smart contracts, which can sustain market equilibria with a larger range of economic outcomes. Biais et al. (2019) use a stochastic game to show that the proof-of-work protocol results in multiple equilibria, some of which can lead to persistent divergence between chains. A revenue management problem in the context of bitcoin selling is studied in Dai et al. (2019). Our work differs from these studies in that we analyze centralization of the reward and computational power in mining activities as well as how price and competition impacts it.

Our work is most closely related to recent literature on miners' strategic behavior and the centralization of mining. Cong et al. (2019) examine mining pools, and unexpected impacts of their risk sharing, such as the concentration of the mining power. Arnosti and Weinberg (2018) consider asymmetric costs among miners and show that lower cost leads to higher market share. On the other hand, Alsabah and Capponi (2019) explore a two-stage mining game consisting of research and development and then competition. They explain how the arms race leads to asymmetric costs and mining centralization. Different from these static games, our work considers continuous mean field games, incorporating dynamic change of miners' wealth and decisions over time. We refer to Guéant et al. (2011) for an early introductory exposition on mean field games.

Our work also contributes to the literature on intensity control of jump processes. It is used in the model for exploration of natural resources. Deshmukh and Pliska (1980) and

Arrow and Chang (1982) study the optimal consumption rule of a natural resource. They use a point process to model the uncertainty of the discoveries for new sources of supply, where the control is exploration effort. Later on, Soner (1985) considered a similar model with holding cost, and established the existence and uniqueness of solution to the Bellman equation. Intensity control models are also used in revenue management and dynamic pricing. A buffer flow system with jumps is considered by Li (1988), where the cumulative production and demand are modeled by two counting processes, with intensity controlled by production capacity and price. In addition, Gallego and van Ryzin (1994, 1997) model dynamic pricing for inventories of products. The demand for those is modeled as point processes and the intensities are controlled by setting prices. In our work, the jump process is used to represent the acquisition of the reward. The miners control the jump intensity through adjusting their computational power or hash rates. This model approach is natural due to the two important properties of Bitcoin payment system mentioned before.

There has been recent work on games of intensity control. For instance, Ludkovski and Sircar (2012) consider the effects of stochastic resource exploration in dynamic Cournot game, where an exhaustible producer and a green producer set the production to affect the price. Gallego and Hu (2014) study dynamic pricing in an oligopolistic market. Each firm competes to sell its product and the equilibrium strategies and prices are resolved. In a mean field game setting, Chan and Sircar (2017) examine the impact of oil discovery, concluding that higher reserves lead to lower exploration. There the players' interaction was through producers' oil extraction rates. In this paper, different from most works in the literature, the mean field interaction is through the players' intensities, or hash rates.

1.2 Mean field game model

To study the centralization of mining rewards, we formalize a mean field game model, where each miner is characterized by its wealth, and chooses its hash rate to maximize expected utility at a fixed time horizon. Its wealth changes because of the mining rewards and expenses. The instantaneous probability of receiving the reward is

$$\frac{\text{pl. } i\text{'s hash rate}}{\text{total hash rate}} = \frac{\text{pl. } i\text{'s hash rate}}{\#\text{players} \times \text{mean hash rate}} = \frac{\text{pl. } i\text{'s hash rate}}{\text{pl. } i\text{'s hash rate} + (\#\text{players} - 1) \times \text{mean hash rate}}.$$

In order to utilize computational advantages of mean field games technology, our model replaces the second term in the denominator by ($M \times$ continuum mean hash rate), where $M + 1$ represents the actual total number of miners, and is large. This implies that the mean field interaction is strong, whereas often in the literature it is assumed to be small for computational and technical reasons. We argue that for cryptocurrency problems, interaction with the total hash rate is essential in a realistic model. Indeed, this does introduce numerical difficulties, for which we provide an effective algorithm in Section 2.3.1.

For mining without liquidity constraints, we find the equilibrium explicitly under exponential utility, whereas with liquidity constraints we solve the equilibrium numerically with power utility. It is suggested by the model that heterogeneity of the initial wealth distribution among miners results in greater concentration of wealth over time, or "the rich get richer". In other words, the miner with more wealth contributes more hash rate and thus has a higher probability of getting the next reward. It is also illustrated that the price of bitcoin can foster this centralization, while competition can reduce it.

In an extended model, we consider an *advanced miner* with cost advantages in the mean field game model. There are three main features distinguishing it from other miners. First, it is cost efficient. This could be due to having access to cheaper electricity or advanced equipment. Second, the advanced miner is risk neutral and maximizes running profit over time instead of terminal utility. Lastly, its hash rate affects the probability for the representative miner of getting the reward. But the representative miner only affects the advanced miner through the mean field. The model shows that the advanced miner can account for a significant amount of the total hash rate. Other miners become less active when the cost of the advanced miner decreases. Hence, cost efficiency is a factor leading to the centralization of mining power.

2 Competition among homogeneous miners

We begin by describing the general structure of both the individual's mining problem and the subsequent equilibrium. This structure is then used in the study of mining, first without liquidity constraints, and thereafter with.

2.1 General structure of the mining problem

We consider a continuum of miners who competitively engage in Bitcoin mining over some finite time period $[t_0, T]$. The miners have initial wealth $x \in \mathbb{R}$, distributed at time t_0 according to an initial density function m_0. The representative miner provides hash rate α_t, incurring a linear cost per unit of time $c\alpha_t$, where $c > 0$, and $t \in [t_0, T]$. This is interpreted as the cost of electricity, and is thus proportional to their hash rates. It can also be thought of as encompassing any other linear cost, but for simplicity we disregard the possibility of nonlinear costs. In particular, the cost of mining equipment is linear in equipment, but not in the mining rate.

There are two important features of the Bitcoin proof-of-work protocol: First, the system always generates a reward on an almost fixed frequency that does not depend on the total hash rate. In fact, the system will adjust the difficulty to make a reward available every 10 minutes on average. So it is reasonable to model the total number of rewards in the system as a whole as a Poisson process with a constant intensity $D > 0$. Second, a miner's probability of receiving the next mining reward is proportional to the ratio of its hash rate to that of the population. Since the math puzzle needs to be solved by brute force, the more hash rate a miner contributes, the more likely it will obtain the reward.

The number of rewards each miner can receive is modelled by a counting process N_t with jump intensity $\lambda_t > 0$.[1] Let $M + 1$ be the total number of miners and $\bar{\alpha}_t$ denote the mean hash rate across all miners. Here, our model for the reward intensity as a function of an individual's hash rate α_t and the mean hash rate is

$$\lambda_t := \frac{\alpha_t}{D(\alpha_t + M\bar{\alpha}_t)},$$

and we use $M\bar{\alpha}_t$ to approximate the total hash rate of other miners.

[1] Formally, $P[N_{t+\Delta t} - N_t = 1] = \lambda_t \Delta t + o(\Delta t)$ and $P[N_{t+\Delta t} - N_t \geq 2] = o(\Delta t)$, see Brémaud (1981).

Each miner is considered small and has negligible impact on the population's mean production. To model the behavior of each individual miner, let $\bar{\alpha} = (\bar{\alpha}_t)_{t \geq t_0}$ be a given process describing the mean production, where $t_0 \geq 0$ is the initial time. Then, the miner's wealth process X_t follows

$$dX_t = -c\alpha_t \, dt + p \, dN_t, \tag{2.1}$$

where p is the value of the mining reward.

The value of each reward is the product of the bitcoin prince p and its quantity.[2] Since our focus is on the strategic decision of miners and the centralization in the competition, we treat the price as a constant. The number of bitcoins as a reward is set to decrease geometrically with 50% reduction every 4 years approximately. Currently, a miner can be rewarded by 12.50 bitcoins if it adds the next block successfully. We do not distinguish between the value of the reward and the price of bitcoin since the quantity generated each time is fixed over four years.

2.1.1 The miner's problem

Suppose that $\alpha = (\alpha_t)_{t \geq t_0}$ is a Markovian control. The process α can then be associated with a function $(t, X_t) \mapsto \alpha(t, X_t; \bar{\alpha})$ of the current state. We call a control admissible if $\alpha(t, x; \bar{\alpha}) \in [0, A(x)]$, for a given non-decreasing function $A : \mathbb{R} \to [0, +\infty]$. The function $A(x)$ is part of the problem specification and should be thought of as encoding the liquidity constraints of each miner.[3] For instance, a liquidity constraint will later be imposed by requiring mining to cease when X drops to zero. This is encoded as $A(0) = 0$.

With such controls, the wealth process X is a Markov process. The objective of the representative miner is to maximize the expected utility at fixed terminal time T. We assume the utility function U is strictly increasing and concave. The objective function is written as

$$v(t_0, x; \bar{\alpha}) = \sup_{\alpha} \mathbb{E}[U(X_T) | X_{t_0} = x], \tag{2.2}$$

where we emphasize that X_T depends on α and $\bar{\alpha}$ through the cost and N.

Lemma 2.1. *Fix a choice of $\bar{\alpha} > 0$. For any time $t \in [t_0, T]$, the value function $v(t, x; \bar{\alpha})$ is finite and strictly increasing in the wealth x.*

This lemma is standard, because more wealth gives more flexibility to miners to choose their hash rates. It will be useful in the following derivation.

For a fixed mean hash rate $\bar{\alpha} > 0$, we first write down the HJB

$$\partial_t v + \sup_{\alpha \in [0, A(x)]} \left(-c\alpha \partial_x v + \frac{\alpha}{D(\alpha + M\bar{\alpha}_t)} \Delta v \right) = 0, \tag{2.3}$$

[2]The miners are also rewarded the transaction fees in successfully mined blocks. This is paid in units of bitcoins. The transaction fees usually account for small proportion of the total reward, so, as our focus is not on the structure of the reward, we do not consider these fees in our model.

[3]It could also encode constraints on the hardware capacity or access to electricity. However, this paper will focus on financial liquidity constraint.

with terminal condition $v(T,x) = U(x)$, and where $\Delta v = v(t, x+p; \bar{\alpha}) - v(t,x;\bar{\alpha})$. By Lemma 2.1, $\Delta v > 0$ and $\partial_x v > 0$, and so the optimal hash rate is given by

$$\alpha^*(t,x;\bar{\alpha}) = \begin{cases} \min\left\{-M\bar{\alpha}_t + \sqrt{\dfrac{M\bar{\alpha}_t \Delta v(t,x;\bar{\alpha})}{Dc\partial_x v(t,x;\bar{\alpha})}}, A(x)\right\}, & \text{if} \quad \bar{\alpha}_t < \dfrac{\Delta v(t,x;\bar{\alpha})}{MDc\partial_x v(t,x;\bar{\alpha})}, \\ 0, & \text{otherwise.} \end{cases} \tag{2.4}$$

If $A(x)$ is sufficiently large, the HJB equation can be simplified as

$$\begin{cases} \partial_t v + \left(\sqrt{Mc\bar{\alpha}_t \partial_x v} - \sqrt{\dfrac{\Delta v}{D}}\right)^2 = 0, & \text{if} \quad \bar{\alpha}_t < \dfrac{\Delta v}{MDc\partial_x v}, \\ \partial_t v = 0, & \text{otherwise.} \end{cases} \tag{2.5}$$

2.1.2 Equilibrium characterization

The continuum of miners are labelled by their wealth x. Let $\alpha^*(t,x;\bar{\alpha})$ be the optimal hash rate of miner x, and denote by $m(t,x;\bar{\alpha})$ the resulting density of the miners' wealth as a function of time and wealth. We say $\bar{\alpha}^*$ forms an *equilibrium mean hash rate* of the mining game if

$$\bar{\alpha}_t^* = \int_{\mathbb{R}} \alpha^*(t,x;\bar{\alpha}^*) m(t,x;\bar{\alpha}^*) dx, \quad \forall t \in [t_0, T].$$

Henceforth, let $\bar{\alpha}^*$ denote an equilibrium mean hash rate, and denote

$$v(t,x) = v(t,x;\bar{\alpha}^*), \quad \alpha^*(t,x) = \alpha^*(t,x;\bar{\alpha}^*), \quad m(t,x) = m(t,x;\bar{\alpha}^*).$$

We will assume that $\bar{\alpha}_t^* \neq 0$ for all t for the following reason. If $\bar{\alpha}_t^* = 0$, then each miner has an admissible control that dominates the choice of not mining, provided some non-zero control is admissible. Hence, unless the mass of miners with non-zero admissible controls is zero, some mining will always occur.

We assume the initial density $m_0(x)$ is continuously differentiable and satisfies

$$\int_{\mathbb{R}_{\geq 0}} m_0(x) dx = 1. \tag{2.6}$$

That is, each miner starts with nonnegative, finite wealth.

In the equilibrium, if $A(x)$ is sufficiently large,

$$\bar{\alpha}_t^* = \int_{E_t} \alpha^*(t,x) m(t,x) dx = -M\eta(t)\bar{\alpha}_t^* + \sqrt{\frac{M\bar{\alpha}_t^*}{Dc}} \int_{E_t} \sqrt{\frac{\Delta v(t,x)}{\partial_x v(t,x)}} m(t,x) dx,$$

where

$$E_t = \{x : \alpha^*(t,x) > 0\} \tag{2.7}$$

denotes the wealth level on which the miners are active and

$$\eta(t) = \int_{E_t} m(t,x) dx$$

denotes the fraction of active miners. Thus,

$$\bar{\alpha}_t^* = \frac{M}{(1+M\eta(t))^2}\left(\int_{E_t}\sqrt{\frac{\Delta v(t,x)}{Dc\partial_x v(t,x)}}m(t,x)dx\right)^2, \tag{2.8}$$

and the Fokker-Planck equation is given by

$$\partial_t m - \partial_x(c\alpha^*(t,x)m) - \frac{1}{D}\left(\frac{\alpha^*(t,x-p)}{\alpha^*(t,x-p)+M\bar{\alpha}_t^*}m(t,x-p) - \frac{\alpha^*(t,x)}{\alpha^*(t,x)+M\bar{\alpha}_t^*}m(t,x)\right) = 0, \tag{2.9}$$

with initial distribution $m(t_0,x) = m_0(x)$.

2.2 Exponential utility and mining without liquidity constraints

In the model of unconstrained liquidity, a miner's wealth X_t can take negative values by means of interest-free borrowing. This means that miners can continue their mining activity even if they would otherwise be ruined, i.e., $A(\cdot) \equiv \infty$. This may not be true in reality, but it shows us the fundamental structure of the mining problem on which we build more reasonable models later on.

The above analysis in Section 2.1 does not make any specific assumption on the utility function. In general, the model can only be solved numerically to find the equilibrium fixed point. However, with exponential utility, we are able to find the solution explicitly.

Proposition 2.2. *With exponential utility $U(x) = -\frac{1}{\gamma}e^{-\gamma x}$ ($\gamma > 0, x \in \mathbb{R}$), in the equilibrium, all miners are always active, with constant hash rate*

$$\alpha^*(t,x) \equiv \bar{\alpha}_t^* \equiv \frac{M}{(1+M)^2}\frac{1-e^{-\gamma p}}{Dc\gamma}, \tag{2.10}$$

and their individual reward rate is

$$\lambda_t \equiv \frac{1}{D(1+M)},$$

for any $t \in [t_0, T]$ and $x \in \mathbb{R}$. The value function is given by

$$v(t,x) = U(x)e^{-\frac{1-e^{-\gamma p}}{D(1+M)^2}(T-t)}. \tag{2.11}$$

Proof. We guess the form $v(t,x) = U(x)h(t)$ and then the HJB equation (2.5) becomes

$$\begin{cases} \partial_t h - \gamma\left(\sqrt{Mc\bar{\alpha}_t^*} - \sqrt{\dfrac{1-e^{-\gamma p}}{D\gamma}}\right)^2 h = 0, & \text{if} \quad \bar{\alpha}_t^* < \dfrac{1-e^{-\gamma p}}{MDc\gamma}, \\ \partial_t h = 0, & \text{otherwise,} \end{cases}$$

with terminal condition $h(T) = 1$. It is an equation involving t only, which validates our ansatz. On the other hand, this ansatz implies that $\Delta v/\partial_x v$ does not depend on x. Hence, by (2.8),

$$\bar{\alpha}_t^* = \frac{\eta^2 M}{(1+\eta M)^2}\frac{\Delta v}{Dc\partial_x v} \le \frac{M}{(1+M)^2}\frac{\Delta v}{Dc\partial_x v} < \frac{\Delta v}{MDc\partial_x v},$$

which means all miners are active and $\eta \equiv 1$.

The equilibrium mean hash rate is obtained from plugging in the ansatz into (2.8). Thereafter, (2.4) yields (2.10). The value function then satisfies

$$\partial_t h - \frac{1}{(1+M)^2}\frac{1-e^{-\gamma p}}{D}h = 0,$$

with terminal condition $h(T) = 1$. Thus we have (2.11). □

Risk-Reward Analysis

As α^* is constant, we drop the dependence on t and x. In the equilibrium, the wealth of the representative miner can be written as

$$X_{t_0+t} = X_{t_0} - c\alpha^* t + p(N^*_{t_0+t} - N^*_{t_0}) = X_{t_0} - \frac{M}{(1+M)^2}\frac{1-e^{-\gamma p}}{\gamma D}t + p(N^*_{t_0+t} - N^*_{t_0}),$$

where N^* has the jump rate $\lambda_t = \frac{1}{D(1+M)}$. Then we can get the expectation and variance of X_{t_0+t},

$$\begin{aligned}\mathbb{E}[X_{t_0+t}] &= \mathbb{E}[X_{t_0}] + \left(p - \frac{M}{1+M}\frac{1-e^{-\gamma p}}{\gamma}\right)\frac{t}{D(1+M)},\\ \mathrm{Var}(X_{t_0+t}) &= \mathrm{Var}(X_{t_0}) + \frac{p^2 t}{D(1+M)}.\end{aligned}$$

Using that $0 < \frac{1-e^{-\gamma p}}{\gamma} \le p$, the expected wealth change $\mathbb{E}[X_{t_0+t} - X_{t_0}]$ satisfies

$$\frac{pt}{D(1+M)} > \mathbb{E}[X_{t_0+t} - X_{t_0}] \ge \frac{pt}{D(1+M)^2} \ge 0.$$

In above expressions, we can identify two sources of risk. The first is the price of bitcoin. The expected wealth grows linearly in the price of bitcoin, but the variance increases quadratically. This indicates that a price increase adversely affects the risk relative to the reward. Although with high prices p, mining can bring considerable rewards, the potential loss is greater.

The second source of risk is competition. If γp is very small, the expected wealth increment is discounted by $(M+1)^2$. However, the variance is only discounted by a linear factor $(M+1)$. Hence, as more miners join the game, the expected gain shrinks very fast, but the potential risk decreases slowly. We notice that

$$P[N^*_{t_0+t} - N^*_{t_0} = 0] = e^{-\frac{t}{D(1+M)}}.$$

In fact, as the number of miners grows, the probability of getting no new bitcoin approaches 1. This reflects that the competition reduces the chance of getting the reward.

The optimal hash rate in (2.10) is increasing in the price and decreasing in the number of miners. That is, a larger reward inspires miners to hash faster, whereas competition diminishes the value of their efforts. For each M and p, $\alpha^*_t \le \frac{1}{c\gamma(1+M)D}$, which shows that

miners do not increase their hash rates arbitrarily. It is also worth noting that the social cost $(M+1)c\alpha^*$ has an upper bound $(M+1)c\alpha^* < \frac{1}{D\gamma}$. In other words, this is an upper bound for the total cost paid by all miners.

It is also interesting to see that the optimal hash rate does not depend on individual wealth. This is because miners can keep mining even with negative wealth. In the next section, we present a model where miners with negative wealth are instead eliminated from the mining game.

2.3 Liquidity-constrained mining and utilities on $\mathbb{R}_{\geq 0}$

In this section, we consider the mining problem where miners have constrained liquidity, and utility functions U defined on $\mathbb{R}_{\geq 0}$, satisfying $U(0) > -\infty$, $U'(0+) = \infty$, and $U'(\infty) = 0$. In particular, miners face ruin when their wealth falls to zero. This means that a miner with zero wealth does not have the resources to pay for electricity and other expenses, and is therefore not able to mine at all. We encode this restriction by letting

$$A(0) = 0 \text{ (equivalently, } \alpha^*(t,0) \equiv 0) \text{ and } A(x) = \infty \text{ for } x > 0.$$

As argued in Section 2.1.2, we assume that $\bar{\alpha}^* \neq 0$. Thus, with this choice of A, solving the problem (2.2), we reach the equation (2.5) for $x > 0$ with the boundary condition

$$v(t,0) = U(0). \tag{2.12}$$

This condition reflects that miners cease their mining if they are without wealth, i.e., they are out of the game once their wealth hits 0. Note that the wealth of a miner that starts with non-negative wealth will remain non-negative. Recall from (2.6) that the initial density has non-negative support. This means that the problem is fully characterized on $\mathbb{R}_{\geq 0}$.

In contrast to the problem of Section 2.2, with the boundary condition (2.12), the value function cannot be found explicitly, even with power utility, so we must solve (2.5) numerically. Nevertheless, the mean hash rate is still given by (2.8). However, the Fokker–Planck equation has two parts to account for the boundary condition. The density m solves (2.9) (at least in a weak sense) for $x > p$. On the other hand, if $0 < x < p$, there is no density at $x - p$ jumping to x, because no miners are active with negative wealth. Thus, for the optimal individual hash rate

$$\alpha^* = \alpha^*(t,x) = -M\bar{\alpha}_t + \sqrt{\frac{M\bar{\alpha}_t \Delta v(t,x;\bar{\alpha})}{Dc\partial_x v(t,x;\bar{\alpha})}}$$

and for $0 < x < p$, the density m solves

$$\partial_t m - \partial_x(c\alpha^* m) + \frac{\alpha^*}{D(\alpha^* + M\bar{\alpha}^*)} m = 0, \tag{2.13}$$

with initial condition $m(t_0, x) = m_0(x)$.

Remark 2.3. *The Fokker–Planck equation can be verified to preserve mass on $\mathbb{R}_{\geq 0}$, i.e., for*

all $t \in [t_0, T]$, $\int_{\mathbb{R}_{\geq 0}} m(t,x)dx = 1$. *Indeed,*

$$\begin{aligned}
\partial_t \int_{\mathbb{R}_{\geq 0}} m(t,x)dx &= \int_{\mathbb{R}_{\geq 0}} \partial_t m(t,x)dx, \\
&= c\alpha^*(t,\cdot)m(t,\cdot)|_0^p - \int_0^{+\infty} \frac{\alpha^*}{D(\alpha^* + M\bar{\alpha}_t^*)} m dx + c\alpha^*(t,\cdot)m(t,\cdot)|_p^{+\infty} \\
&\quad - \int_p^{+\infty} \frac{\alpha^*(t,x-p)}{D(\alpha^*(t,x-p) + M\bar{\alpha}_t^*)} m(t,x-p)dx, \\
&= c\alpha^*(t,\cdot)m(t,\cdot)|_0^{+\infty} = 0.
\end{aligned}$$

The last equation holds because no one can obtain infinite wealth in finite time, and the condition $\alpha^*(t,0) = 0$ *holds.*

2.3.1 Numerical method

The method we use to numerically find an equilibrium is as follows. First, we solve both the nonlinear HJB equation (2.3) with (2.12) for some mean hash rate. Then, we use the optimal control and the Fokker–Planck equation to find the evolution of the population density. With these two solutions, we calculate the corresponding mean hash rate and repeat the process until convergence. This procedure is described in greater detail below.

1. Initialize with a mean hash rate $t \mapsto \bar{\alpha}_t$, for instance a constant.

2. Solve for the value function and control:

   Given α^*, (2.3) is a linear PDE that can be solved by standard methods. We thus begin by approximating α^*, starting at time T.

   At time T, the value function is v known, so (2.4) yields $\alpha^*(T,x;\bar{\alpha})$. This value is then used as an approximation of $\alpha^*(T-dt,x;\bar{\alpha})$, which allows us to solve for v at $T-dt$, using the HJB:

   $$\partial_t v + \left(-c\alpha^*(T,x;\bar{\alpha})\partial_x v + \frac{\alpha^*(T,x;\bar{\alpha})}{D(\alpha^*(T,x;\bar{\alpha}) + M\bar{\alpha}_T)} \Delta v \right) = 0.$$

   The Δv term is calculated explicitly using $v(T,x+p;\bar{\alpha}) - v(T,x;\bar{\alpha})$, while the other part is discretized by an implicit finite difference scheme. With the value function v at $T-dt$, we can get $\alpha^*(T-dt,x;\bar{\alpha})$. Repeat such time steps backwards until $t=0$. This yields both functions v and α^*.

3. The next step is to solve the Fokker–Planck equation and get the mean field control.

   The $\alpha^*(t,x;\bar{\alpha})$ is obtained from the previous step allows us to solve for $m(t,x)$ using (2.9) and (2.13). In doing so, the following parts are discretized by an implicit finite difference scheme

   $$\partial_t m - \partial_x(c\alpha^*(t,x)m) + \frac{\alpha^*(t,x)}{D(\alpha^*(t,x) + M\bar{\alpha}_t)} m,$$

while m in

$$-\frac{1}{D}\frac{\alpha^*(t,x-p)}{\alpha^*(t,x-p)+M\bar{\alpha}_t}m(t,x-p)$$

is evaluated in the previous time step.

4. To reduce oscillations in searching for the equilibrium, we introduce a parameter of inertia, $w \in [0,1)$. At each time t, we update the mean hash rate according to

$$\bar{\alpha}_t^{\text{new}} = w\bar{\alpha}_t + (1-w)\int_{\mathbb{R}} \alpha^*(t,x;\bar{\alpha})m(t,x;\bar{\alpha})dx.$$

The choice of w has no impact on the equilibrium fixed point.

Finally, repeat from the first step with $\bar{\alpha} = \bar{\alpha}^{\text{new}}$ until convergence.

Because of the destabilizing effect of $M \gg 1$, we use $w = 1 - \frac{1}{M}$. This yields stable iterations at the expense of slower convergence. The choice has been successful for all figures presented here, but experiments have shown that faster choices, i.e., smaller w sometimes also works.

2.3.2 Concentration of wealth and mining effort

In this section, we numerically solve for the equilibrium with liquidity constraints and for utility functions of constant relative risk aversion (CRRA), also known as power utility. That is,

$$U(x) = \frac{1}{1-\gamma}x^{1-\gamma} \quad \text{for } \gamma > 0, \gamma \neq 1. \tag{2.14}$$

This structure of liquidity constraints and CRRA utility leads to strategic decisions of the miners that are very different from those without liquidity constraints. When taking illiquidity into consideration, those with larger wealth tend to hash more. This results in a phenomenon called *preferential attachment*, or *the rich get richer*, which means that those who have more also receive more.

Preferential attachment happens in many situations: scientific citation networks (Barabási et al., 2002), language use (Perc, 2012), distribution of cities by population and distributions of incomes by size (Simon, 1955). A recent study points out that it also appears in the Bitcoin network (Kondor et al., 2014). It states that "we find that the wealth of already rich nodes increases faster than the wealth of nodes with low balance." Our numerical results show that the heterogeneity in miners' wealth leads to this phenomenon.

Figure 1 shows the distribution of the miners' wealth at $t = 30, 45, 60, 90$ compared with the initial distribution. As time increases, the majority of the mass moves to the left, forming a big spike gradually. At the same time, there is small part of the mass moving to the right. This indicates that most miners lose their wealth, but those who have relatively more money originally accumulate wealth over time.

Figure 2(a) shows the expected instantaneous profit, namely,

$$-c\alpha^*(t,x) + \frac{p}{D}\frac{\alpha^*(t,x)}{\alpha^*(t,x)+M\bar{\alpha}_t^*}.$$

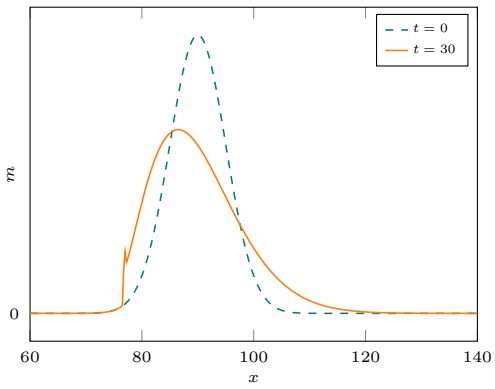

(a) Density comparison ($t = 0, 30$).

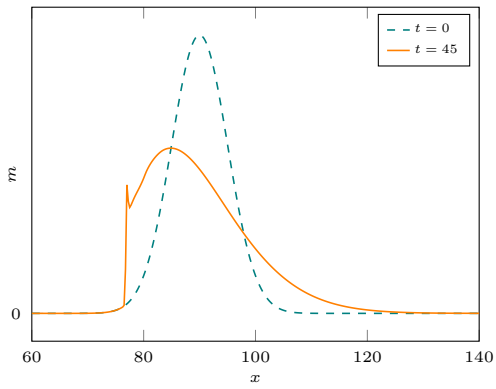

(b) Density comparison ($t = 0, 45$).

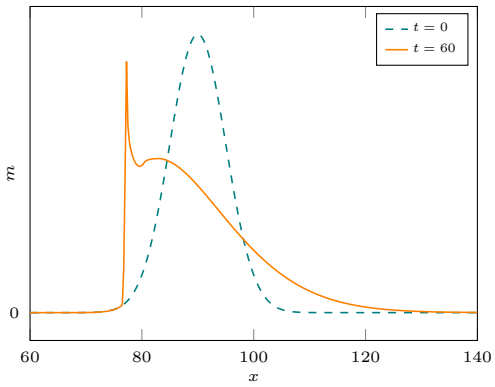

(c) Density comparison ($t = 0, 60$).

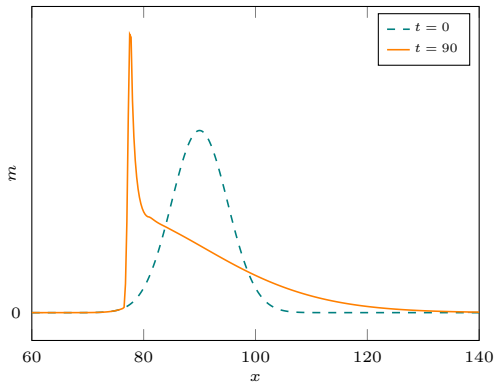

(d) Density comparison ($t = 0, 90$).

Figure 1: *The distribution of miners' wealth at different times. Parameters:* $D = 0.007$, $p = 3$, $c = 2 \times 10^{-5}$, $T = 90$, $\gamma = 0.8$, $M = 1000$. *The initial distribution* m_0 *is normal with mean* 90 *and standard deviation* 5.

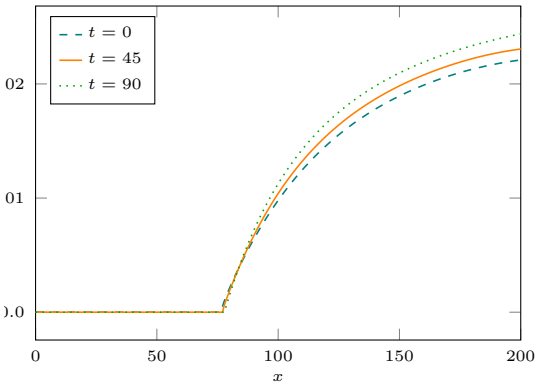

(a) Expected profit.

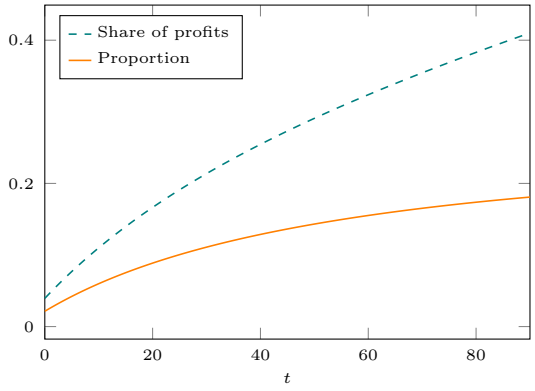

(b) Miners with wealth over 100.

Figure 2: *The left plot shows the expected instantaneous profit miners at different wealth levels and times. The right one gives the proportion of miners with wealth over 100 and their share of the total instantaneous profits. Parameters are the same as Figure 1.*

This shows that the more wealth a miner has, the higher is the reward it receives. Miners with lower wealth worry more about ruin, and thus they hash at lower rates or even zero rate. This pattern holds at all times. In addition, it is interesting to see that the miners are blocked (0 hash rate) around and below wealth level 80 (see Figure 2(a)). The risk aversion prevents miners with small wealth from participating in the mining game. This is summarized by the following lemma. Its proof is provided in the appendix.

Lemma 2.4. *For any time t and equilibrium hash rate $\bar{\alpha}_t^* > 0$, there exists $x_b(t) > 0$ such that zero rate mining is optimal, i.e., $\alpha^*(t, x) = 0$ for $x \leq x_b(t)$.*

Moreover, we calculate the proportion of miners whose wealth is over 100 and their their share of the instantaneous profits, both of which are shown in Figure 2(b). The proportion increases from around 2% to 18%, while the share of profits rises from 4% to 41%. Hence, as time goes by, the wealthy receive an increasingly large share of the profits.

We have shown "the rich get richer" phenomenon in the mining game. It is also interesting to see how the price p and competition parameter M affect this phenomenon, as these two are the sources of risk mentioned in Section 2.2. To avoid repetition, we show only the plots at $t = 30$, but a similar pattern is present also at other times.

A larger price p exacerbates the degree of preferential attachment. The density plots in Figure 3 show that those with lower wealth tend to lose money faster when the price is higher. At the same time, the density for $x \geq 110$ is clearly higher for the larger price in Figure 3 (a)(b)(c). Figure 3(d) show the expected instantaneous profit, which leads to the same conclusion.

The competition parameter M reduces the preferential attachment. In Figure 4 (a)(b)(c) this shown, as the density for $x \geq 100$ is lower for larger M. Additionally, the hash rate and the profit decrease with respect to M, as is show in Figure 4(d). Meanwhile, as the competition becomes fierce and the entry level for the game is larger. When $M = 1000$, miners need wealth around 75 to enter the game, but this increases to about 90 for $M = 10000$. Hence, the competition makes the mining less lucrative and makes it harder for miners to stay active, which reduces the preferential attachment.

3 Competition with cost advantages

In this section, we consider a model in which a miner can have cost advantages over the rest. This could be due to access to cheaper energy or more advanced equipment, and helps the miner become dominant in the mining game. Bitmain is one example of an advantageous miner. It takes advantage of the cheaper electricity in China, like the hydropower stations in Sichuan during the rainy season, and also of its expansion overseas, like Hydro Quebec in Canada, which offers some of the lowest electricity rates in North America.[4] The model studied in this section suggests that cost advantages can be a contributing factor in the centralization observed in Bitcoin mining, which is dominated by a few large entities, as is illustrated in Figure 5.

[4] `https://www.reuters.com/article/us-canada-bitcoin-china/chinese-bitcoin-miners-eye-sites-in-energy-rich-canada-idUSKBN1F10BU`

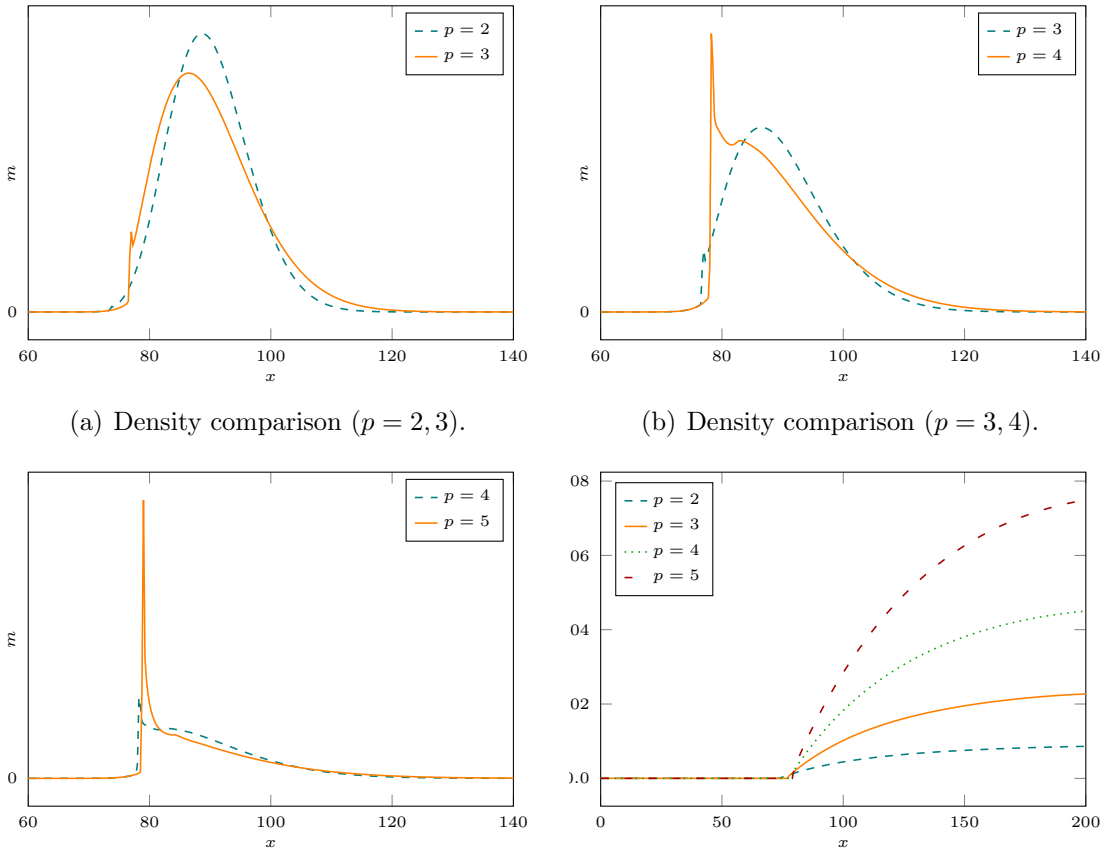

(a) Density comparison ($p = 2, 3$).

(b) Density comparison ($p = 3, 4$).

(c) Density comparison ($p = 4, 5$).

(d) Expected profit.

Figure 3: *The price effects at $t = 30$. The first three plots show the distribution of miners' wealth. The last one shows the expected instantaneous profit of miners at different wealth levels, for four different price levels. Parameters are the same as Figure 1 except that p takes multiple values.*

3.1 The cost-advantaged miner problem

We consider a cost-advantaged miner, competing with $M+1$ individual miners introduced in Section 2.1.[5] This miner chooses its hash rate β_t, with the corresponding cost $c_1\beta_t = k_c c\beta_t$, where $0 < k_c \leq 1$ is the relative cost efficiency. Given the mean hash rate $\bar{\alpha}_t$ of individual miners, let the counting process N_t^1 with intensity

$$\lambda_t^1 = \frac{\beta_t}{D(\beta_t + (M+1)\bar{\alpha}_t)}$$

denote the number of rewards received by the advanced miner.

We assume that the advanced miner is wealthy and therefore not liquidity constrained. Moreover, for simplicity, it is also risk neutral and it aims to maximize its running profit.

[5]This type of competition between an individual and a continuum of payers is related to so-called major-minor mean field games, see e.g. Huang (2010). However, the introduction of our parameter M to approximate an aggregate in terms of a mean implies that the so-called minor players are not really minor.

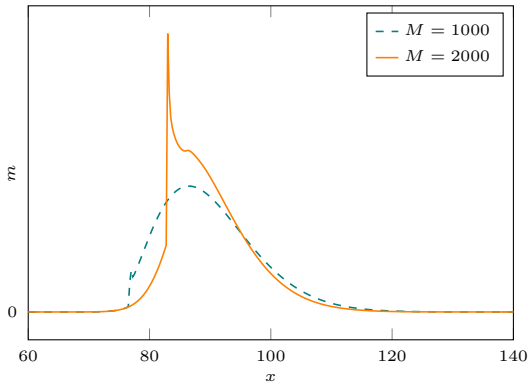

(a) Density comparison ($M = 1K, 2K$).

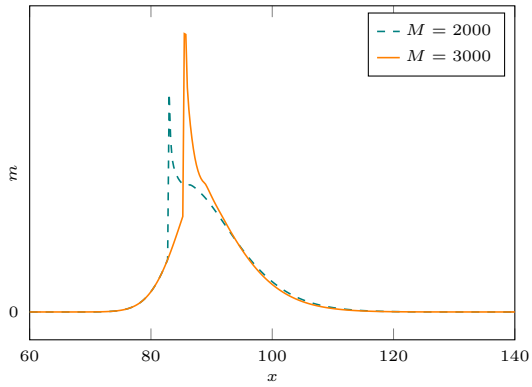

(b) Density comparison ($M = 2K, 3K$).

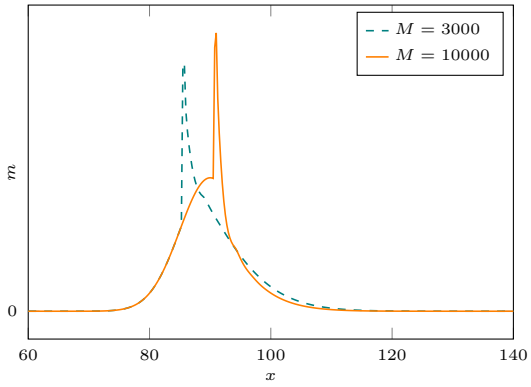

(c) Density comparison ($M = 3K, 10K$).

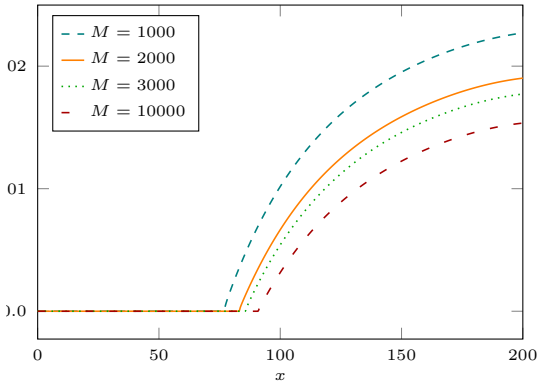

(d) Expected profit.

Figure 4: *The competition effects at $t = 30$. The first three plots show the distribution of miners' wealth. The last one shows the expected instantaneous profit for miners at different wealth levels, for four different competition levels. Parameters are the same as Figure 1 except that M takes multiple values.*

Hence, the objective for the advanced miner is

$$\sup_{\beta_t \geq 0} \mathbb{E}\left[\int_0^T -c_1 \beta_t dt + p dN_t^1\right]. \tag{3.1}$$

As the optimization problem is independent of wealth, any Markov control can be identified by a function $\beta(t; \bar{\alpha})$, i.e., the control only depends on time.

Given $\bar{\alpha} > 0$, the maximizer in (3.1) satisfies the first-order condition

$$-c_1 + \frac{p(M+1)\bar{\alpha}_t}{D(\beta_t + (M+1)\bar{\alpha}_t)^2} = 0,$$

which yields the best response

$$\beta^*(t; \bar{\alpha}) = \begin{cases} -(M+1)\bar{\alpha}_t + \sqrt{\dfrac{p(M+1)\bar{\alpha}_t}{c_1 D}}, & \text{if} \quad \bar{\alpha}_t < \dfrac{p}{c_1(M+1)D}, \\ 0, & \text{otherwise.} \end{cases} \tag{3.2}$$

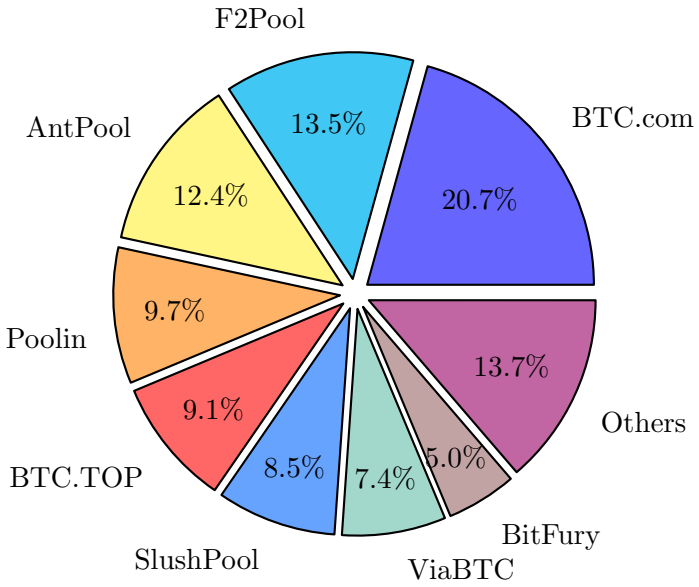

Figure 5: *Bitcoin hash rate distribution among the largest mining pools. The data is obtained on 06/30/2019 from* `https://www.blockchain.com/pools`.

3.2 The individual miner's problem

As in Section 2.1, the model for individual miners remains the same except that the intensity for N_t in (2.1) becomes

$$\lambda_t = \frac{\alpha_t}{D(\alpha_t + M\bar{\alpha}_t + \beta_t t)},$$

given the advanced miner's hash rate β_t. Here the denominator consists of both the advanced miner's hash rate and the total of individual miners. Hence, the value function defined in (2.2) depends on both $\bar{\alpha}$ and α, i.e., $v(t_0, x; \bar{\alpha}, \beta)$.

For a fixed choice of $\bar{\alpha} > 0$ and $\alpha \geq 0$, the HJB can be written as

$$\partial_t v + \sup_{\alpha \in [0, A(x)]} \left(-c\alpha \partial_x v + \frac{\alpha}{D(\alpha + M\bar{\alpha}_t + \beta_t)} \Delta v \right) = 0,$$

with terminal condition $v(T, x) = U(x)$. Like Lemma 2.1, it can be proved that v is strictly increasing in x. Hence, the maximizer is taken as

$$\alpha^*(t, x; \bar{\alpha}, \beta) = \begin{cases} -(M\bar{\alpha}_t + \beta_t) + \sqrt{\dfrac{(M\bar{\alpha}_t + \beta_t)\Delta v}{Dc\partial_x v}}, & \text{if} \quad M\bar{\alpha}_t + \beta_t < \dfrac{\Delta v}{Dc\partial_x v}, \\ 0, & \text{otherwise.} \end{cases} \tag{3.3}$$

If $A(x)$ is sufficiently large, the HJB is simplified as

$$\begin{cases} \partial_t v + \left(\sqrt{c(M\bar{\alpha}_t + \beta_t)\partial_x v} - \sqrt{\dfrac{\Delta v}{D}} \right)^2 = 0, & \text{if} \quad M\bar{\alpha}_t + \beta_t < \dfrac{\Delta v}{Dc\partial_x v}, \\ \partial_t v = 0, & \text{otherwise.} \end{cases} \tag{3.4}$$

3.3 Equilibrium characterization

Let $m(t,x;\bar{\alpha},\beta)$ be the resulting density, corresponding to the optimal hash rate $\alpha^*(t,x;\bar{\alpha},\beta)$ of individual miners. We say that $\bar{\alpha}^*$ and β^* form an *equilibrium* of the mining game with an advanced miner if

$$\bar{\alpha}_t^* = \int_{\mathbb{R}} \alpha^*(t,x;\bar{\alpha}^*,\beta^*)m(t,x;\bar{\alpha}^*,\beta^*)dx, \quad \forall t \in [t_0,T],$$

and $\beta_t^* = \beta^*(t;\bar{\alpha}^*)$, given by (3.2). Henceforth, let $\bar{\alpha}^*$ and β^* denote the equilibrium mean hash rate and equilibrium hash rate for the advanced miner, $v(t,x) = v(t,x;\bar{\alpha}^*,\beta^*)$, $\alpha^*(t,x) = \alpha^*(t,x;\bar{\alpha}^*,\beta^*)$, and $m(t,x) = m(t,x;\bar{\alpha}^*,\beta^*)$.

By the same argument presented in Section 2.1.2, it is meaningful to consider $\bar{\alpha}_t^* > 0$ for all t. And we also assume the initial density satisfies (2.6). Thus in the equilibrium, if $A(x)$ is sufficiently large, we have coupled equations $\beta_t^* = \beta^*(t;\bar{\alpha}^*)$ and

$$\bar{\alpha}_t^* = -\eta(t)(M\bar{\alpha}_t^* + \beta_t^*) + \sqrt{(M\bar{\alpha}_t^* + \beta_t^*)} \int_{E_t} \sqrt{\frac{\Delta v(t,x)}{Dc\partial_x v(t,x)}} m(t,x)dx,$$

by integrating (3.3) on x over the set (2.7). The Fokker-Planck equation is given by

$$\partial_t m - \partial_x(c\alpha^*(t,x)m) - \frac{1}{D}\left(\frac{\alpha^*(t,x-p)}{\alpha^*(t,x-p) + M\bar{\alpha}_t^* + \beta_t^*} m(t,x-p) - \frac{\alpha^*(t,x)}{\alpha^*(t,x) + M\bar{\alpha}_t^* + \beta_t^*} m(t,x)\right) = 0,$$

with initial distribution $m(t_0,x) = m_0(x)$.

3.4 Exponential utility and mining without liquidity constraints

In the absence of liquidity constraints and with exponential utility, we have the following lemma.

Proposition 3.1. *Suppose the individual miners have exponential utility $u = -\frac{1}{\gamma}e^{-\gamma x}$ and no liquidity constraints $A(\cdot) \equiv \infty$, suppose the relative cost efficiency satisfies*

$$k_c < \frac{\gamma p}{1 - e^{-\gamma p}} \frac{M+1}{M}, \tag{3.5}$$

and let

$$\kappa_1 = \frac{1 - e^{-\gamma p}}{Dc\gamma}, \quad \kappa_2 = \frac{(M+1)p}{Dc_1}.$$

Then, in equilibrium, all miners are active with

$$\alpha^*(t,x) \equiv \bar{\alpha}_t^* \equiv \frac{\kappa_1^2\kappa_2}{(\kappa_1+\kappa_2)^2} > 0, \quad \beta_t^* \equiv \frac{\kappa_1\kappa_2(\kappa_2 - M\kappa_1)}{(\kappa_1+\kappa_2)^2} > 0, \tag{3.6}$$

for all $t \in [t_0,T]$ and $x \in \mathbb{R}$.

Proof. We consider the ansatz $v(t,x) = u(x)h(t)$ and then the HJB (3.4) becomes

$$\begin{cases} \partial_t h - \gamma \left(\sqrt{c(M\bar{\alpha}_t^* + \beta_t^*)} - \sqrt{\dfrac{1-e^{-\gamma p}}{D\gamma}} \right)^2 h = 0, & \text{if} \quad M\bar{\alpha}_t^* + \beta_t^* < \dfrac{1-e^{-\gamma p}}{Dc\gamma}, \\ \partial_t h = 0, & \text{otherwise,} \end{cases}$$

with terminal condition $h(T) = 1$. Since $\bar{\alpha}^*$ and β^* are only functions of t, this validates the ansatz. In looking for an equilibrium in which $\alpha_t^* > 0$ and $\beta_t^* > 0$, we use the non-zero best response β^* in (3.2), and, using the ansatz in (3.3),

$$\alpha^*(t,x;\bar{\alpha}^*,\beta^*) = -(M\bar{\alpha}_t^* + \beta_t^*) + \sqrt{\frac{1-e^{-\gamma p}}{Dc\gamma}(M\bar{\alpha}_t^* + \beta_t^*)}.$$

Since α^* does not depend on the wealth x, all individual miners are active. Therefore, we have

$$\bar{\alpha}_t^* = \alpha_t^* \equiv -(M\bar{\alpha}_t^* + \beta_t^*) + \sqrt{\frac{1-e^{-\gamma p}}{Dc\gamma}(M\bar{\alpha}_t^* + \beta_t^*)}.$$

This, together with (3.2), yields (3.6). It is direct that α^* is positive, and β^* is positive if and only if (3.5) holds. Thus we have found the equilibrium in which everyone is active. □

Cost advantage and its effect on mining power concentration

Proposition 3.1 demonstrates that the cost-advantaged miner's efficiency leads to centralization in the following sense. It can be checked that the hash rate β_t^* in (3.6) is increasing in κ_2 and hence decreasing in c_1. Similarly, the hash rate α^* in (3.6) of the individual miners increases with respect to c_1. Thus, a smaller c_1—a bigger cost advantage—makes the advanced miner more dominant. As a consequence, individual miners with higher cost have to decrease their hash rates to regulate their risk exposure, as the advanced miner gets a larger share of the mining rewards.

To quantify this, consider the case that $\gamma \ll p$, so that also the individual miners are (almost) risk neutral. To understand the share of the reward obtained by the advanced miner, we write

$$k_c = \rho \frac{\gamma p}{1-e^{-\gamma p}} \approx \rho$$

for some constant $0 < \rho \leq 1$. Then $\kappa_2 = (M+1)\kappa_1/\rho$. The probability for the advanced miner to get the reward is

$$\frac{\beta^*}{\beta^* + (M+1)\bar{\alpha}^*} = \frac{\kappa_2 - M\kappa_1}{\kappa_1 + \kappa_2} = \frac{(1-\rho)M + 1}{M + 1 + \rho} \approx 1 - \rho$$

for sufficient large M, whereas the remaining miners have a collective probability ρ and individual probability $\rho/(M+1)$. If the advanced miner is 10% efficient ($k_c = 0.9$), then $\rho \leq 0.9$, which gives a probability around 10% for the advanced miner to get the reward.

Let Y and Y^1 denote the profits of the individual miner and the advanced miner. Then we have

$$Y_{t_0+t} = -c\alpha^* t + pN_t^*, \quad Y_{t_0+t}^1 = -c_1\alpha_1^* t + pN_t^{1*},$$

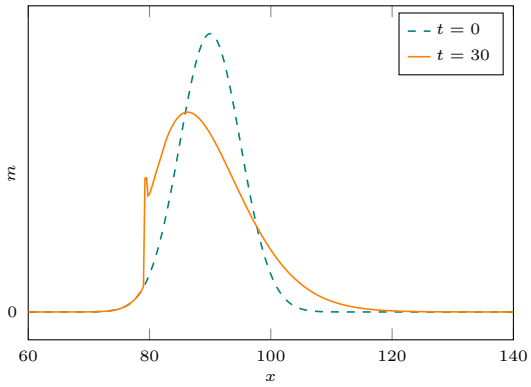

(a) Density comparison ($t = 0, 30$).

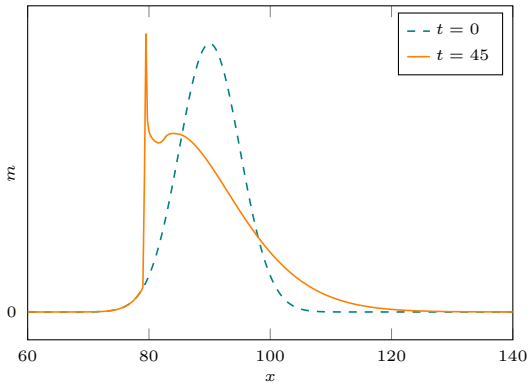

(b) Density comparison ($t = 0, 45$).

Figure 6: *The distribution of miners' wealth at $t = 0, 30, 45$. For both figures, $k_c = 0.8$. The parameters are the same as Figure 1.*

where N_t^* and N_t^{1*} have jump rates $\frac{\rho}{D(M+1+\rho)}$ and $\frac{(1-\rho)M+1}{D(M+1+\rho)}$. We can then get the expectation and variance.

$$\mathbb{E}(Y^1_{t_0+t}) = \frac{p}{D}\left(\frac{(1-\rho)M+1}{M+1+\rho}\right)^2 t, \qquad \mathbb{E}(Y_{t_0+t}) = \left(p - \frac{M+1}{M+1+\rho}\frac{1-e^{-\gamma p}}{\gamma}\right)\frac{\rho t}{D(M+1+\rho)},$$

$$\mathrm{Var}(Y^1_{t_0+t}) = \frac{p^2((1-\rho)M+1)}{D(M+1+\rho)}t, \qquad \mathrm{Var}(Y_{t_0+t}) = \frac{p^2\rho t}{D(M+1+\rho)}.$$

Notice that c_1 is increasing in ρ. Hence, $\mathbb{E}(Y_{t_0+t})$ and $\mathrm{Var}(Y_{t_0+t})$ are increasing in c_1 while $\mathbb{E}(Y^1_{t_0+t})$ and $\mathrm{Var}(Y^1_{t_0+t})$ are decreasing in c_1. Combining with the analysis before, the advanced miner hashes harder to have more expected profits and more risk at the same time. However, individual miners facing the competition from the advanced miner have to decrease their hash rates and thus receive less profits. Meanwhile, the risk they are faced with is lower.

3.5 Liquidity-constrained mining and power utility

In this section, we assume individual miners take liquidity constraints into consideration and compete with the advanced miner. The model for the individual miners is described in Section 2.3. We use power utility (2.14) in our numerical results. The HJB and the Fokker–Planck equations of Sections 3.2 and 3.3 are modified like in Section 2.3. We omit the details for brevity. The numerical method correspondingly follows the procedure in Section 2.3.1 with appropriate updates.

Figure 6 shows the wealth distribution of individual miners when $k_c = 0.8$ for $t = 0, 30, 45$. The change from Figure 1 is not immediately apparent, but there are two subtle differences. Most clearly visible is the larger portion of individuals who are in the region of optimal zero rate mining. This shift can be attributed to the upward shift of this cutoff. The other effect is the slightly slower dispersion of the distribution, which is explained by the hesitancy in the region right above the cutoff.

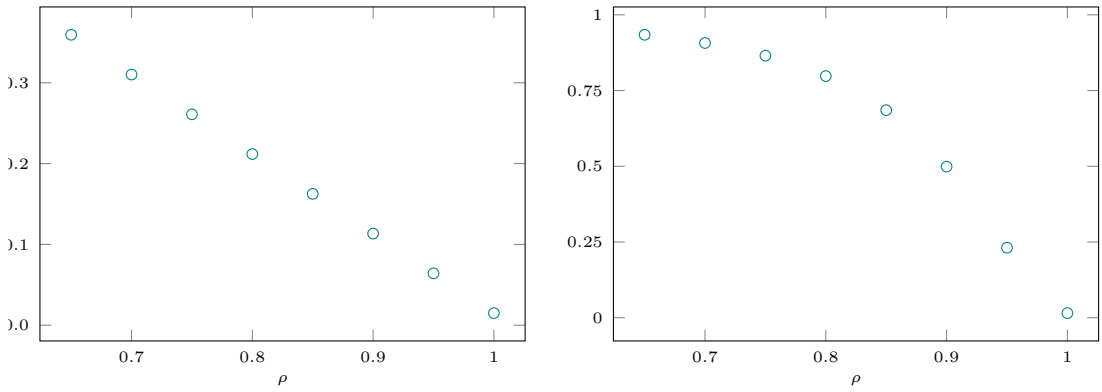

(a) The advantaged miner's share of total rewards.

(b) The advantaged miner's share of total profits.

Figure 7: *Effect of cost efficiency, k_c. The first shows the advantaged miner's probability of getting the next reward, i.e., the expected share of total rewards. The last one shows the advantaged miner's expected instantaneous profits divided by the total instantaneous profits of all miners. For both figures, $t = 30$. The parameters are the same as Figure 1.*

The effect of varying the cost efficiency k_c is plotted in Figure 7. Figure 7(a) shows the advantaged miner's probability of getting the next reward, which is also the share of the expected instantaneous reward. As it is more cost efficient from $k_c = 1.0$ to 0.65, its hash rate accounts for around 1% to 35%. Hence, the cost advantages could be one explanatory factor for the concentration of mining power. A similar idea also appears in (Arnosti and Weinberg, 2018). They suggest that if a miner's cost is (e.g.) 10% lower than those of other miners, then the miner must control at least 10% of the total mining power. Alsabah and Capponi (2019) argue that miners invest in R&D which allows them to develop more energy efficient mining equipment. Hence miners can have lower marginal cost and contribute more hash rates. As $M = 1000$, the advantaged miner contributes a somewhat higher hash rate than the rest of the population—which is on the order of $1/M$—even for $k_c = 1$. We attribute this to the difference in risk aversion, as these numbers are for $\gamma = 0.8$.

Figure 7(b) plots the share of total profits for the advantaged miner. At a 35% cost advantage, i.e., $k_c = 0.65$, the advantaged miner reaps 93% of the total profits generated, and 86% of profits for $k_c = 0.75$. This shows that most of the economic welfare in the system is received by a miner with a cost advantage.

As an example of dominant mining, Bitmain controls AntPool and BTC.com which account for around 33% of the total hash rate in the world as of June 2019 (see Figure 5). This number includes Bitmain's computational power as well as those miners who join the pools. Until Bitmain began to disclose its hash rate in 2018, it was not known how much Bitmain itself contributed. According to Bitmain,[6] it had 2339.21 PH/s on Bitcoin mining in October 2018. A rough estimate of the total hash rate at that time was 50000 PH/s on BTC.com website.[7] Thus, Bitmain accounted for about 4.5% of all computational power for

[6] `https://web.archive.org/web/20181017133438/https://blog.bitmain.com/en/hashrate-disclosure/`

[7] `https://btc.com/stats/diff`

Bitcoin mining. As already mentioned, it has access to cheaper electricity, which enables it to contribute significant large hash rate.

Moreover, Taylor (2017) points out that some mining entities develop application-specific integrated circuits (ASICs) and create related data centers with low energy cost. For instance, BitFury takes advantage of ASICs, which convert the same amount of electricity into more hash rates. "BitFury optimizes its chips for use in new immersion-cooled datacenters in the Republic of Georgia, Iceland, and Finland". Thus, more advanced equipment also helps the advanced miner acquire a big share of the mining market.

4 Conclusion

This paper develops models to study the centralization of the reward and computational power in Bitcoin mining. The mean field game among a continuum of miners is explored. As a result of the heterogeneity of the miners' wealth, more rewards tend to be collected by those who have more wealth, which is "the rich get richer". Since miners are maximizing their own utility, the rich will contribute more hash rates to compete for the reward and hence have higher probability of receiving the next reward. Moreover, the price of bitcoin fosters this phenomenon, since higher price means more reward and more incentive to hash. However, the competition will reduce this phenomenon, because the probability of getting the next reward becomes smaller for each miner. In addition, we incorporate an advanced miner into the game to study the centralization of the computational power. The result shows that if a miner is cost efficient, then it will contribute a significant amount of hash rate in the game. This explains that mining pools that have more advanced equipment or access to cheaper electricity account for most of the computational power in recent years.

Appendix A

A.1 Proof of Lemma 2.1

Proof. For the finiteness, by Jensen's inequality,

$$\mathbb{E}[u(X_T)|X_t = x] \leq u(\mathbb{E}[X_T|X_t = x]) \leq u(x + p\mathbb{E}[N_T - N_t]) \leq u\left(x + \frac{pT}{D}\right) < \infty.$$

For any $x_1 < x_2$, let α_t^i and X_t^i $(i = 1, 2)$ denote the optimal hash rate and corresponding wealth starting at time t with initial wealth x_i. Consider the case where we start with x_2. We use x_1 as the wealth in the mining and save $x_2 - x_1$ in a bank account. Then we use the hash rate α_t^1. The corresponding wealth process is denoted by $X_t^{2,1}$. Thus we have $X_t^{2,1} \geq x_2 - x_1 + X_t^1 > X_t^1$. Thus,

$$v(t, x_1) < \mathbb{E}[u(X_T^{2,1})|X_t^{2,1} = x_2] \leq v(t, x_2).$$

□

A.2 Proof of Lemma 2.4

Proof. For any $\varepsilon > 0$,

$$\frac{v(\varepsilon) - v(0)}{\varepsilon} \geq \frac{U(\varepsilon) - U(0)}{\varepsilon} \xrightarrow{\varepsilon \to 0} \infty.$$

Hence, by continuity, $\Delta v / \partial_x v$ is arbitrarily small in some neighborhood of 0. Thus, by (2.4), there exists a $x_b(t)$ such that zero rate mining is optimal for $x \leq x_b$. □

References

Abadi, J. and Brunnermeier, M. K. (2018). Blockchain economics. Working paper, Princeton University.

Alsabah, H. and Capponi, A. (2019). Pitfalls of Bitcoin's Proof-of-Work: R&D arms race and mining centralization. Available at SSRN: https://ssrn.com/abstract=3273982.

Arnosti, N. and Weinberg, S. M. (2018). Bitcoin: A natural oligopoly. In *ITCS*.

Arrow, K. J. and Chang, S. (1982). Optimal pricing, use, and exploration of uncertain natural resource stocks. *Journal of Environmental Economics and Management*, 9(1):1–10.

Barabási, A., Jeong, H., Néda, Z., Ravasz, E., Schubert, A., and Vicsek, T. (2002). Evolution of the social network of scientific collaborations. *Physica A: Statistical Mechanics and its Applications*, 311(3-4):590–614.

Biais, B., Bisière, C., Bouvard, M., and Casamatta, C. (2019). The blockchain folk theorem. *The Review of Financial Studies*, 32(5):1662–1715.

Brémaud, P. (1981). *Point processes and queues: martingale dynamics*, volume 50. Springer.

Chan, P. and Sircar, R. (2017). Fracking, renewables, and mean field games. *SIAM Review*, 59(3):588–615.

Cong, L. W. and He, Z. (2019). Blockchain disruption and smart contracts. *The Review of Financial Studies*, 32(5):1754–1797.

Cong, L. W., He, Z., and Li, J. (2019). Decentralized mining in centralized pools. Working paper 25592, National Bureau of Economic Research.

Dai, M., Jiang, W., Kou, S., and Qin, C. (2019). From Hotelling to Nakamoto: The economic meaning of Bitcoin mining. In preparation.

Deshmukh, S. D. and Pliska, S. R. (1980). Optimal consumption and exploration of nonrenewable resources under uncertainty. *Econometrica*, 48(1):177–200.

Easley, D., O'Hara, M., and Basu, S. (2019). From mining to markets: The evolution of Bitcoin transaction fees. *Journal of Financial Economics*, 134(1):91–109.

Gallego, G. and Hu, M. (2014). Dynamic pricing of perishable assets under competition. *Management Science*, 60(5):1241–1259.

Gallego, G. and van Ryzin, G. (1994). Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science*, 40(8):999–1020.

Gallego, G. and van Ryzin, G. (1997). A multiproduct dynamic pricing problem and its applications to network yield management. *Operations Research*, 45(1):24–41.

Guéant, O., Lasry, J.-M., and Lions, P.-L. (2011). Mean field games and applications. In *Paris-Princeton Lectures on Mathematical Finance 2010*, pages 205–266. Springer.

Huang, M. (2010). Large-population LQG games involving a major player: The nash certainty equivalence principle. *SIAM Journal on Control and Optimization*, 48(5):3318–3353.

Kondor, D., Pósfai, M., Csabai, I., and Vattay, G. (2014). Do the rich get richer? An empirical analysis of the Bitcoin transaction network. *PLoS ONE*, 9(2):e86197.

Li, L. (1988). A stochastic theory of the firm. *Mathematics of Operations Research*, 13(3):447–466.

Ludkovski, M. and Sircar, R. (2012). Exploration and exhaustibility in dynamic Cournot games. *European Journal of Applied Mathematics*, 23(3):343–372.

Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system. https://bitcoin.org/bitcoin.pdf.

Perc, M. (2012). Evolution of the most common English words and phrases over the centuries. *Journal of The Royal Society Interface*, 9(77):3323–3328.

Simon, H. A. (1955). On a class of skew distribution functions. *Biometrika*, 42(3-4):425–440.

Sockin, M. and Xiong, W. (2018). A model of cryptocurrencies. Working paper, Princeton University.

Soner, H. M. (1985). Optimal control of a one-dimensional storage process. *Applied Mathematics & Optimization*, 13(1):175–191.

Taylor, M. B. (2017). The evolution of Bitcoin hardware. *Computer*, 50(9):58–66.