

A Mean Field Games Model for Cryptocurrency Mining

Zongxi Li* A. Max Reppen† Ronnie Sircar*

December 2019; revised January 14, 2022

Abstract

We propose a mean field game model to study the question of how centralization of reward and computational power occur in 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, and increased wealth heterogeneity over time, or a “rich get richer” effect. This concentration phenomenon is aggravated by a higher bitcoin mining reward, and reduced by competition. Additionally, an advantaged miner with cost advantages such as access to cheaper electricity, contributes a significant amount of computational power in equilibrium, unaffected by competition from less efficient miners. Hence, cost efficiency can also result in the type of centralization seen among miners of cryptocurrencies.

1 Introduction

Blockchain technologies solve the elusive problem of creating a decentralized ledger. Proponents have made the case that a payment and banking system founded on blockchain technologies could conceivably allay data privacy concerns, and provide wider access to financial services. Various forms of digital currency solutions have been developed in the past, but Bitcoin is the famous realization that finally created a secure distributed ledger (see Nakamoto (2008)). Since its creation in January 2009, its fame and popularity have grown rapidly. The supply of bitcoins is constantly growing, but artificially limited to 21 million, of which almost 19 million are in circulation now.

In a cryptocurrency network following a so-called proof-of-work protocol, independent “miners” compete for the right to record the next transaction block on the blockchain. They do so by solving computational puzzles: once a miner obtains a solution, the corresponding block is added on top of the blockchain, and the miner obtains a reward. The computational puzzle is designed such that there is no known better way of solving it than brute

*ORFE Department, Princeton University, Princeton, USA.

†Questrom School of Business Boston University, Boston, MA, 02215, USA. Partly supported by the Swiss National Science Foundation grant SNF 181815.

force calculation. In other words, the chance of getting the reward is proportional to the computational power or the hash rates that miners contribute. The purpose of the puzzle is to disincentivize bad behavior by forcing miners to provide proof-of-work that some effort has been exerted – electricity paid. Moreover, the difficulty of the puzzle varies to maintain a consistent solving time, for example 10 minutes.

The empirical analysis in Kondor, Pósfai, Csabai, and Vattay (2014) shows that the accumulation of bitcoins tends to occur among a small number of miners, which suggests centralization in the market. This raises the following questions: What incentives drive miners’ competitive behavior? How does the centralization of the reward occur in a decentralized mining environment? What factors impact this centralization?

Another emerging phenomenon is major miners who have access to cheaper electricity or more efficient hardware and thus are cost-advantaged. Bitmain controls AntPool and BTC.com and accounts 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 that of miners who join the pools. By rough estimates, Bitmain accounted for about 4.5% of all mining power for Bitcoin mining in October 2018.^{1,2} Moreover, as Taylor (2017) points out, some mining entities also develop application-specific integrated circuits (ASICs) and create related data centers with low energy cost. One may ask: What leads to this centralization of computational power? What advantages do some miners have?

We are interested in understanding and modeling the interaction between cryptocurrency miners, the consequent evolution of wealth inequality among them, and the potential centralization of mining power. In the context of cryptocurrency mining, an economic model captures the cost-reward structure of mining: the cost – typically the marginal cost of electricity, and the potential reward – the units of currency rewarded to successful miners.

An ongoing regulatory issue is as to whether cryptocurrencies are currencies or commodities (or even share-like assets offered in initial coin offerings). The Commodity Futures Trading Commission (CFTC) in the US classifies them as commodities, and their electronic structure of production mirrors the uncertainty and language of mining exhaustible resources. The latter connects us to game theoretic models that have been developed to try and explain energy production from various sources many of which, like oil, are in finite supply (see, for instance, Ludkovski and Sircar (2015)).

Similarly to models of natural resource extraction, the miners are producers, but the product is numbers. These numbers, called hashes, come from processing transaction data. Because of the demand for processing, there is effectively also a demand for hashes. So, just like in any other market, the suppliers (miners) are competing to fill this demand. And as in other markets, the (expected) reward per hash obtained is diminishing in the aggregate supply (mining power). However, calling it a market for hashes is masking the actual mechanism, because it is not the numbers that are valuable. Instead, the value lies in the process that ‘extracts’ the numbers (mining), and the numbers themselves are just a byproduct. Therefore the basis of the models introduced in this paper is the probabilistic relation between competitive hashing and cryptocurrency discovery.

¹<https://web.archive.org/web/20181017133438/https://blog.bitmain.com/en/hashrate-disclosure/>

²<https://btc.com/stats/diff>

Our paper quantifies the competition between miners by adopting a tractable mean field games approach. The idea behind mean field games is that, with many participants, any particular player has little impact on any other player when they interact through the mean of their actions—hashing in our case. As the number of players grows, one can first view an individual’s decision making problem as being against a mean field competition, knowing that their individual contribution to the mean field is infinitesimally small. The final step is a fixed point condition in which the mean optimal control should coincide with the aforementioned mean field. This approach leads to a computationally tractable model, which is intuitively a good approximation to the finite player game, precisely because each player’s impact dissipates in the mean of many players.

1.1 Related literature

Our work is related to a growing literature on cryptocurrencies. A game-theoretic model is developed in Easley, O’Hara, and Basu (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, Bisière, Bouvard, and Casamatta (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, Jiang, Kou, and Qin (2019). Our work differs from these studies in that we analyze centralization of both the reward and computational power in mining activities as well as how the reward size and competition impact it.

Our work is most closely related to recent literature on miners’ strategic behavior and the centralization of mining. Cong, He, and Li (2021) examine mining pools’ fee-setting decisions, and their unexpected impact on risk sharing and concentration. 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 (2020) 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. Garratt and van Oordt (2020) study the importance of the cost structure in cryptocurrency mining, and they find that fixed costs improves the resilience to shocks due to its impact on exit and entry.

In contrast to the static games in these papers, our work considers continuous time mean field games, incorporating dynamically evolving wealth and mining decisions. We refer to Guéant, Lasry, and Lions (2011) for an early introductory exposition on mean field games. Bertucci, Bertucci, Lasry, and Lions (2020) provide another example of mean field game modeling of cryptocurrency mining. Their model complements ours, as they focus on the blockchain’s resilience to outside attacks, whereas ours focuses on centralization and inside attacks. Their modeling decisions and approach are therefore very different, in order to capture the phenomena of importance in that setting.

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 L. 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.

The proof-of-work system supporting Bitcoin, Ethereum,³ and the majority of alt-coins has received heavy criticism for the high energy consumption of its miners. The most prominent alternative consensus method is called proof-of-stake.⁴ In this system, instead of spending computing power as a requirement for creating valid blocks, the participants are instead "randomly" chosen in proportion to their current stake in the system. A user holding 1% of all coins will in the long run create 1% of new blocks. Because block creation is wasteless, this leads to the so-called nothing-at-stake problem in which deviation from "good" behavior is not punished. We refer to Brown-Cohen, Narayanan, Psomas, and Weinberg (2019) for a more detailed account of proof-of-stake and its drawbacks. Fanti et al. (2019) and Roşu and Saleh (2021) both study the impact of rewards on the wealth distribution of participants by using the martingale property of participants' *share* of assets in a proof-of-stake system. Roşu and Saleh (2021) show that under a constant reward scheme, the limiting distribution is stable in terms of the share of total coins, and thus fair (not exhibiting a rich get richer effect).⁵ Additionally, proof-of-stake is believed to not only have the advantage of

³Ethereum is working towards replacing proof-of-work with proof-of-stake in Ethereum 2.0 (Eth2). The first step towards this migration (Phase 0) was deployed in December 2020.

⁴Proof-of-stake was first introduced in 2011 by user QuantumMechanic on the bitcointalk forums: <https://bitcointalk.org/index.php?topic=27787.0>

⁵Although this is a very interesting result and good news for proof-of-stake systems, we perceive this as strongly connected to the martingale property of such systems, and believe that it does not translate to proof-of-work for two main reasons: First, central to the martingale property is the dilution effect (inflation tax) of introducing new coins. In a proof-of-stake system, this means that whereas a wealthier participant receives proportionally larger rewards, the inflation tax is also proportionally larger. In contrast, in a proof-of-work system, a miner does not need to hold the cryptoasset to reap the rewards, creating the possibility

reducing energy consumption and improving fairness, but also of increasing security, as any attacker must have a stake in the system. At time of writing, various proof-of-stake systems are under heavy development, both theory and in practice.

1.2 Contribution and intuition

We introduce a dynamic competitive mining game in the presence of risk aversion and liquidity constraints. The miners compete by exerting computational effort in an attempt to obtain the mining rewards, which are distributed based on the computational effort relative to the population aggregate. They are expected utility maximizers, balancing the cost of computation (e.g. electricity) and the reward associated with block creation.

We find the equilibrium explicitly under exponential (CARA) utility (Section 3.1), thereby establishing the existence and uniqueness of a solution in this case. The explicit solution allows us to understand the influence of model parameters on the equilibrium.

For miners with power (CRRA) utility, we find the equilibrium numerically (Section 3.2). In this case, heterogeneity of the initial wealth distribution among miners results in preferential attachment, i.e. increasing heterogeneity of wealth over time. In other words, a miner with greater wealth contributes a larger hash rate, and thus has a higher probability of receiving the next reward, whereas some miners with lesser initial wealth become disincentivized over time from participating entirely, leading to increased wealth inequality. Moreover, our results show that increasing centralization in mining power arises, an antithesis to the principles behind cryptocurrency security.

In an extended model (Section 4), we consider an *advanced miner* with cost advantages, for instance due to having access to cheaper electricity or advanced equipment. We show that the advanced miner accounts for a significant share of the total hash rate. Hence, cost efficiency is another factor leading to the centralization of mining power.

The main driver of our results is that, *ceteris paribus*, the individual miner’s ratio of expected reward and the standard deviation of the reward is decreasing in the hash rate (at least to a point). Coupled with risk aversion, this creates an advantage to a miner willing to take on more risk by mining at a higher rate. Although assumptions on cost structure, rewards, or preferences do not change this fundamental feature, they can modulate its strength.

The paper’s technical and methodological contributions are as follows. As each miner’s action affects the other miners equally, the problem exhibits a natural symmetry that is well suited to be modeled as a mean field game. Our adaptation of mean field games technology to this problem is novel: to our knowledge, ours is the first mean field game of control in which the control and the aggregate (rather than the mean) of other players’ control influence the intensity of a jump process. This enables us to capture dynamic features of cryptocurrency mining competition and utility optimization in a numerically tractable way. Finally, we also provide one of the few explicit mean field game equilibria outside of a linear quadratic framework.

for a miner to avoid the inflation tax if so desired. Second, even though some miners could be competitive long-term, the presence of wealth considerations and risk aversion could prematurely force them out of the game against their will. This latter effect is apparent in the model studied here.

2 Model and methodology

In this section we develop our model motivated by the specifics of proof-of-work mining and set up the equations to be tackled analytically and numerically in Section 3. Section 2.1 briefly presents the purpose and mechanics of cryptocurrency mining, followed by Section 2.2 which describes our method for representing participant miners as a continuum. Sections 2.3, 2.4, 2.5 collectively describe the mathematical model for the mining game, and finally Section 2.6 presents the numerical method we implement.

2.1 Proof-of-work mining

Under a proof-of-work protocol, mining is the process by which a block, i.e., a list of transactions, is appended to the cryptocurrency ledger, and by which the system controls who may choose the next block, and thus also the next transactions to be registered. Due to the pseudonymity inherent in the system, any real-world individual or entity can trivially pose as multiple separate users, and for this reason it is problematic to, for instance, take turns in having the right to append the next block.

A solution to this is to give users the right to append the next block in proportion to their computing power by means of a computational mining game. The underlying assumption is that computational power cannot be monopolized. Agents are incentivized to participate by rewards to successful miners. This is accomplished by mining puzzles that are unique to the data being processed. This data, among other things, includes a reference to the last mined block (thus creating the ‘chain’ of blocks) and the transactions to be processed. When the mining puzzle is solved by one miner, the miner receives the reward, the data processed by that miner is appended to the blockchain, and the game starts over for everyone, now with a reference to the new block.

The puzzle is characterized by a binary function h with range $[0, 1]$ and a **target** $\in (0, 1)$, the latter of which is changing over time. A miner successfully mines a block characterized by its **data** if they find an input y such that $h(\text{data}, y) < \text{target}$. Because the second input to the function h is arbitrary and does not carry any information about the transactions themselves, it is referred to as a nonce.

To make this game computationally difficult, the function h must have the property that there is no better way to find a nonce y satisfying $h(\text{data}, y) < \text{target}$ than brute force, i.e., haphazardly evaluating the function for different nonce values. This property is provided by so-called (cryptographic) hash functions. The function h is thus referred to as the hash function and its output as a hash. The precise hash function varies between cryptocurrencies and the target is dynamically adjusted so that the average rate at which the population as a whole solves the puzzles is stable.

All participating miners simultaneously try one nonce after the other until one miner is lucky enough to produce a small enough hash output, below the target. As a consequence, the probability of miner i appending the next block is

$$\frac{\text{miner } i\text{'s hash rate}}{\text{total hash rate}}. \quad (2.1)$$

After the successful mining of a block, the game repeats itself in the hunt for the subsequent block. We see from this expression that, provided no miner’s hash rate is large relative to

the population total, no miner will have undue power over which transactions are included in the ledger. If this is not satisfied and one or more miners run dominant shares of mining computations, the system is *centralized* and those miners have undue power over the ledger and what transactions are included. This could allow dominant miners to conduct censorship or other forms of malicious behavior. For more in-depth detail on the mining game, problems with centralization, and potential attacks, we refer the interested reader to (Narayanan, Bonneau, Felten, Miller, & Goldfeder, 2016, Chapter 5).

Finally, note that because only the data changes between blocks, the process describing the arrival of solutions to the puzzle is memoryless. This is an important property for the mathematical model in Section 2.

2.2 Continuum mean field approximation

To study the centralization of mining power and rewards, we first recognize the mean field structure inherent in mining competition and then introduce a continuum approximation, which leads to the formulation of a continuum mean field game.

From (2.1) we see that each player depends on the other players through the sum of their hash rates, which is proportional to the mean hash rate. Hence, formalizing cryptocurrency mining as a mean field game is completely natural, because each miner is affected by each other miner equally (through their aggregate or mean).⁶

Each miner is characterized by their wealth and chooses their hash rate to maximize expected utility at a fixed time horizon. Their wealth changes because of the mining rewards and expenses. The instantaneous probability of receiving the reward is given by the probability of producing the next block:

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

where the last expression is a good approximation when the number of players is large. Let $M = (\#\text{players}-1)$, which is assumed to be large. To formulate the continuum model, we replace the second term in the denominator on the right hand side by $(M \times \text{continuum mean hash rate})$. With this approximation, the continuum game we consider has significant computational advantages over finite player games, which are often infeasible to solve numerically (because of dimensionality), especially in a dynamic setting.⁷ In other words, we mathematically model a continuum of players, but their aggregate effect on each other is still of size M , and is thus *interpreted* as a game of $M + 1$ players.⁸

The purpose of using a *continuum* mean field games structure instead of a finite N -player⁹ model is that it significantly reduces the dimensionality. The observation that this

⁶This is in contrast to some other proposed applications of mean field games, such as banking, where the assumed symmetry in agent interaction does not capture the intricate graph structure of real-world banking networks. We thank the associate editor for pointing this out.

⁷Mean field game structures and their computational advantages are briefly introduced and described in Appendix B.

⁸This is reasonable as long as the proportion of active players multiplied by M remains large. We restrict our analysis to finite horizons for which this is the case.

⁹We use N here to talk about a general game and $N \rightarrow \infty$ so as not to confuse with M which stays finite in our mining game.

structure can admit simpler solutions, along with analysis of the convergence for $N \rightarrow \infty$, was pioneered in Huang, Malhamé, Caines, et al. (2006); Lasry and Lions (2007). Whereas a finite N -player model requires solving N coupled equations, the mean field games system consists of only two equations. The N -dimensional system of nonlinear partial differential-difference equations with N state variables in the first case is numerically intractable for $N \geq 3$ players, while, as we demonstrate in this paper, the two continuum equations (in time and one state variable) are quite tractable for explicit and numerical resolution. The intuition is as follows.

In a game of many, but finitely many, players, anyone considering the average of the others will observe a quantity very close to the true mean, by the law of large numbers. As the number of players increases to infinity, the observed average converges to the true mean. Hence, in each player’s optimization problem, instead of solving for all possible combinations of randomness influencing others, they need only account for the true mean where individual fluctuations average out.

In a Markovian setting, the best response to the population average behavior can be written as a function of the state. As two different players with the same state respond identically to the same population average, we simultaneously solve for the response of all players. Knowing the best response, we enter this into the Fokker–Planck (Kolmogorov forward) equation, whose solution is the evolution of the density of the population over the state space. With the density and best response, the population average can be computed. If this coincides with the population average for which we found the best response, we have found a (mean field) equilibrium. Such a mean field games equilibrium typically constitutes an ϵ -Nash equilibrium to the finite player game (with ϵ converging to 0 as $N \rightarrow \infty$). See e.g. Huang, Caines, and Malhamé (2007); Nourian and Caines (2013).

To summarize: By considering a continuum of players, they do not need to consider random fluctuations affecting other individuals, which reduces the number of dynamic programming equations to one. Instead, another equation accounts for the population dynamics and behavior. These equations are coupled by a fixed point equation, and at the fixed point no individual has anything to gain from deviating—an equilibrium.

We next describe the general structure of each individual miner’s problem, followed by its mathematical formulation. With the structure of each miner’s action, we then characterize their interaction and the equilibrium condition.

2.3 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 representative miner chooses a hash rate $\alpha_t \geq 0$, incurring a linear cost $c\alpha_t$ per unit of time, where $c > 0$, and $t \in [t_0, T]$. This cost is interpreted as the cost of electricity, and is thus proportional to their hash rates¹⁰, and, for simplicity, each miner incurs the same marginal cost c .¹¹

¹⁰Garratt and van Oordt (2020) find that fixed costs associated with acquisition and setup of mining equipment improves the resilience of the system. We leave this interesting extension of the model for future research.

¹¹We consider some models with cost heterogeneity in Section 4 and Appendix E.

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.¹² 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 $\mathcal{K} > 0$, where \mathcal{K}^{-1} is approximately 10 minutes. 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 $N = (N_t)_{t \geq t_0}$ with jump intensity $\lambda = (\lambda_t)_{t \geq t_0} > 0$.¹³ Let $M + 1$ be the total number of miners and $\bar{\alpha}_t \geq 0$ denote the mean hash rate across all miners. Here, our model for the reward intensity at time t as a function of an individual's hash rate α_t and the mean hash rate is

$$\lambda_t := \begin{cases} \mathcal{K} \frac{\alpha_t}{(\alpha_t + M\bar{\alpha}_t)}, & \alpha_t > 0, \\ 0 & \alpha_t = 0, \end{cases} \quad (2.2)$$

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

Each miner is modeled as having negligible impact on the population's mean production. However, the factor M in front of the mean hash rate in (2.2) 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 2.6.

The miners have initial wealth x , distributed at time t_0 according to an initial density function m_0 . Then, an individual miner's wealth process $X = (X_t)_{t \geq t_0}$ follows

$$dX_t = -c\alpha_t dt + r dN_t, \quad (2.3)$$

where r is the value of the mining reward. Successfully mining a block, i.e., appending the next block to the ledger, grants the miner a reward as compensation. This reward has two parts. The so-called block reward is set by the system as a fixed number per block. In addition to the block reward, the miner also receives any transaction fees from transactions included in the appended block. The value of the total reward is interpreted as the product of the cryptoasset price and its quantity.¹⁴

Since our focus is on the strategic decision of miners and the centralization in the competition, we treat the reward as a constant.¹⁵ Nevertheless, the number of bitcoins as a reward

¹²In reality, in the case of Bitcoin, this number is adjusted every 2016 blocks (about every two weeks). We make the simplifying assumption that this happens continuously in our model. Similarly, we shall assume that also the miners' hash rates may change continuously.

¹³Formally, $P[N_{t+\Delta t} - N_t = 1] = \lambda_t \Delta t + o(\Delta t)$ and $P[N_{t+\Delta t} - N_t \geq 2] = o(\Delta t)$, see Brémaud (1981).

¹⁴The miners are also rewarded the transaction fees in successfully mined blocks. We use the term reward and mining reward to refer to the total amount received, i.e., the block reward plus the transaction fees.

¹⁵Although in reality the mining reward is not constant due to regular decreases of the block reward, fluctuations in transaction fees, and fluctuations in the value of the cryptocurrency relative to the unit

is set to decrease geometrically with 50% reduction every 4 years approximately, and the current block reward is 6.25 bitcoins plus transaction fees of a few percent of the block reward. Although not presented, we have numerically considered this decreasing reward (and, for the sake of completeness, increasing) without seeing a qualitative change in the behavior.

2.4 The miners' optimization problem

With the dynamics of Section 2.3, we are ready to formulate the miners' risk-reward problems. The setup of the continuum model is to consider the optimization problem of an individual miner, in response to any given action of the rest. Let $\alpha = (\alpha_t)_{t \geq t_0}$ be 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. 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 .¹⁶ Let U denote a strictly increasing and concave utility function.¹⁷ The miner's value function is

$$v(t_0, x; \bar{\alpha}) = \sup_{\alpha \geq 0} \mathbb{E}[U(X_T) \mid X_{t_0} = x]. \quad (2.4)$$

We give a basic property of the value function.

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

The proof is given in Appendix A.1. The result is intuitive, because more wealth gives more flexibility to miners to choose their hash rates, and it will be useful below.

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

$$\partial_t v + \sup_{\alpha \geq 0} \left(-c\alpha \partial_x v + \frac{\mathcal{K}\alpha}{(\alpha + M\bar{\alpha}_t)} \Delta v \right) = 0, \quad (2.5)$$

with terminal condition $v(T, x) = U(x)$, and where

$$\Delta v = v(t, x + r; \bar{\alpha}) - v(t, x; \bar{\alpha}).$$

denominating costs, we believe our base model is a tractable starting point from which a lot of features of the cryptocurrency market can be seen. As a first step to seeing the effect of stochastic reward, in Appendix C, we introduce a model with stochastic reward and demonstrate that, in the setting considered there, miners react only to the current level of the price and do not factor in reward volatility. Moreover, we have run the numerical experiments also with deterministically halving rewards, and qualitative results such as the cutoff point in Proposition 3.5 are still observable.

¹⁶Many economic models choose an infinite horizon utility of consumption criterion to remove time as a dimension in the problem. That saving does not occur with mean field games, but the state dimension has already been decreased from $M + 1$ to 2: one plus the time dimension. In addition, one could argue that cryptocurrency miners are not operating on an infinite horizon.

¹⁷The miners are endowed with the same utility function and costs in our analysis, although we do consider some cases of different risk aversions and costs in Section 4. Miners with heterogeneous risk aversion parameters could be considered similarly to the heterogeneous costs in Appendix E, at some added complexity.

We assume that v is a classical solution of (2.5).¹⁸ By Lemma 2.1, $\Delta v > 0$ and $\partial_x v > 0$, and so the optimal hash rate is given by

$$\alpha^*(t, x; \bar{\alpha}) = \begin{cases} -M\bar{\alpha}_t + \sqrt{\frac{\mathcal{K}M\bar{\alpha}_t\Delta v(t, x; \bar{\alpha})}{c\partial_x v(t, x; \bar{\alpha})}}, & \text{if } \bar{\alpha}_t < \frac{\mathcal{K}\Delta v(t, x; \bar{\alpha})}{Mc\partial_x v(t, x; \bar{\alpha})}, \\ 0 & \text{otherwise.} \end{cases} \quad (2.6)$$

We observe in this expression that an individual miner at least temporarily stops hashing at time t if the aggregate competition $M\bar{\alpha}$ is above the threshold $\frac{\mathcal{K}}{c} \frac{\Delta v(t, x; \bar{\alpha})}{\partial_x v(t, x; \bar{\alpha})}$, but may resume as the game evolves. In this sense, there is endogenous entry and exit.

A key question here is whether at a given time t there exists a cutoff point $x_{\text{off}}(t)$ that separates the population into those who are hashing and those who are not, at that time. We observe that such a threshold $x_{\text{off}}(t)$ is determined by when $\frac{\Delta v(t, x; \bar{\alpha})}{\partial_x v(t, x; \bar{\alpha})}$ is small enough. We will see that with CARA utility, the answer is no, whereas with CRRA utility the answer is yes, as shown in Proposition 3.5.

Finally, we simplify the HJB equation (2.5) by plugging in (2.6) to give:

$$\begin{cases} \partial_t v + \left(\sqrt{Mc\bar{\alpha}_t\partial_x v} - \sqrt{\mathcal{K}\Delta v} \right)^2 = 0, & \text{if } \bar{\alpha}_t < \frac{\mathcal{K}\Delta v}{Mc\partial_x v}, \\ \partial_t v = 0, & \text{otherwise.} \end{cases} \quad (2.7)$$

2.5 Equilibrium characterization

We now look for a Markovian equilibrium of the mining game. It is useful to think of the continuum of miners as being labeled 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

$$\bar{\alpha}_t^* = \int_{\mathbb{R}} \alpha^*(t, x; \bar{\alpha}^*) m(t, x; \bar{\alpha}^*) dx, \quad \forall t \in [t_0, T]. \quad (2.8)$$

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 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.¹⁹ 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

$$\int m_0(x) dx = 1.$$

¹⁸In the case of exponential utility, this is indeed verified in Section 3.1, so this is a reasonable assumption.

¹⁹This holds for any cost c , as for any small $\varepsilon > 0$, a miner can hash at rate $\alpha = \varepsilon$ and, on average, receive $r\mathcal{K} - c\varepsilon$ as net rewards. As ε is arbitrarily small, this is clearly positive with arbitrarily small risk.

Substituting (2.6) into the equilibrium condition (2.8) gives

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

where

$$E_t = \{x : \alpha^*(t, x) > 0\} \quad (2.9)$$

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,

$$\bar{\alpha}_t^* = \frac{M}{(1 + M\eta(t))^2} \left(\int_{E_t} \sqrt{\frac{\mathcal{K}\Delta v(t, x)}{c\partial_x v(t, x)}} m(t, x) dx \right)^2, \quad (2.10)$$

while the Fokker-Planck equation for m is given by

$$\partial_t m - \partial_x (c\alpha^*(t, x)m) - \mathcal{K} \left(\frac{\alpha^*(t, x-r)}{\alpha^*(t, x-r) + M\bar{\alpha}_t^*} m(t, x-r) - \frac{\alpha^*(t, x)}{\alpha^*(t, x) + M\bar{\alpha}_t^*} m(t, x) \right) = 0, \quad (2.11)$$

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

Proving existence and uniqueness of an equilibrium in the general case is beyond the scope of this paper. However, in Propositions 3.1 and 4.1, we find the unique equilibrium (in a given class) explicitly. In the next section we introduce a numerical method whose results in Sections 3.2 and 4.3 give strong indication of existence of a unique equilibrium in those cases. Moreover, in the case where we have an explicit solution, the numerical solution is observed to reproduce it (Section 3.1.3).

2.6 Numerical method

Solving the equations (2.7) for v and (2.11) for m , *individually*, are classical PDE problems. The difficulty here lies in finding a solution that satisfies the equilibrium condition (2.8). The equilibrium can be represented as a fixed point equation involving the solutions v and m , formally written as

$$\bar{\alpha} \mapsto \Psi(\bar{\alpha}), \quad \text{where} \quad \Psi(\bar{\alpha}) = (\text{integrate } \alpha^*, m) \circ (\text{solve for } m) \circ (\text{solve for } v, \alpha^*)(\bar{\alpha}),$$

and where \circ as usual denotes composition. Our method for finding an equilibrium is an iteration of the fixed point mapping Ψ .

Each component of Ψ can be readily computed with standard finite difference methods. Although a direct fixed point iteration for Ψ often does work, in this case, the large factor M causes some overshoot when iterating. To resolve this, we use an (under-) relaxation of the fixed point iteration. Such relaxations are well known to help convergence and do not alter the fixed point (equilibrium) itself. H. Li, Fan, and Ying (2021) analyze convergence of this type of relaxed iterations in the particular case of mean field game problems. We refer to them for complete technical details on this type of scheme and the mathematical benefits of relaxation.

What follows is a more detailed account of the iteration, the PDE solutions, and the relaxation technique.

1. Initialize with a mean hash rate $t \mapsto \bar{\alpha}_t$, for instance as constant.
2. Solve for the value function v and the hash rate α^* :

At time T , the value function is v known, so (2.6) yields $\alpha^*(T, x; \bar{\alpha})$. This value is used as an approximation of $\alpha^*(T - dt, x; \bar{\alpha})$, which allows solving for v at $T - dt$, using the HJB equation:

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

The Δv term is calculated explicitly using $v(T, x + r; \bar{\alpha}) - v(T, x; \bar{\alpha})$, while the other part is discretized by an implicit finite difference scheme. This separation is made to balance stability with computational efficiency. 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 solving (2.11) and (3.6) for $m(t, x)$. In doing so, the following terms are discretized by an implicit finite difference scheme

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

while m in

$$-\frac{\mathcal{K}\alpha^*(t, x - r)}{\alpha^*(t, x - r) + M\bar{\alpha}_t}m(t, x - r)$$

is evaluated in the previous time step, again separated to balance stability and efficiency.

4. Due to the large factor M scaling up the errors away from the equilibrium, so a direct iteration will tend to overshoot. We therefore use an under-relaxation technique to reduce this effect. For some parameter $w \in [0, 1)$ and for each time t , the relaxed iteration is given by

$$\bar{\alpha}_t^{\text{new}} = w\bar{\alpha}_t + (1 - w) \int_{\mathbb{R}} \alpha^*(t, x; \bar{\alpha})m(t, x; \bar{\alpha})dx = w\bar{\alpha}_t + (1 - w)\Psi(\bar{\alpha})_t. \quad (2.12)$$

This type of (under-) relaxation technique is common in iteration schemes and stabilizes the iteration at the expense of slower iteration, cf. successive over-relaxation, which instead speeds up stable iterations with $w \geq 1$. The choice of w has no impact on the equilibrium fixed point.

5. Finally, repeat from the first step with $\bar{\alpha} = \bar{\alpha}^{\text{new}}$, unless the difference is sufficiently small to terminate.²⁰

²⁰As usual with fixed point iterations, convergence requires the initial guess to lie in the basin of attraction of the fixed point mapping. Outside this set, the iterations soon result in $\bar{\alpha} = 0$, which can be discarded as a non-equilibrium, as mentioned in Section 2.5. The choice $w = 1 - \frac{1}{M}$ makes the set large enough to find after a few guesses, and this is the value used for all figures. We emphasize that in all our experiments the limiting equilibrium has always been unique, regardless of initial guess, for any particular set of parameters.

We note that H. Li et al. (2021) study relaxation of the problem by relaxing the updating of both v and m . Adding relaxation to the updating of v and m (steps 2 and 3) could have further benefits in improving convergence properties, but because the scheme here successfully solved all problems we considered, such further relaxation was not explored.

3 Explicit (CARA utility) and Numerical (CRRA utility) Solutions & Analysis

When the miners' utility functions are of CARA-type, we give explicit formulas for the equilibrium Section (3.1). These provide evidence that the model is well-posed, but, as usual, wealth effects are missing with this choice of utility. We demonstrate in Section 3.1.3 that the numerical method reproduces the explicit solution. We also address the problem with miners having CRRA utilities numerically, in which case wealth effects become apparent.

3.1 Exponential utility (CARA)

With exponential utility,

$$U(x) = -\frac{1}{\gamma}e^{-\gamma x}, \quad \gamma > 0, \quad x \in \mathbb{R}, \quad (3.1)$$

where γ is the risk aversion parameter,²¹ a miner's current wealth X_t can take any value in \mathbb{R} . It is convenient to define²²

$$w(r; \gamma) = \frac{1 - e^{-\gamma r}}{\gamma}. \quad (3.2)$$

3.1.1 Explicit solution

The following proposition gives formulas for the explicit solution of the mining game.

Proposition 3.1. *With exponential utility, (3.1), in the equilibrium, all miners are always active, with constant hash rate*

$$\alpha^*(t, x) \equiv \bar{\alpha}_t^* \equiv \frac{\mathcal{K}M}{c(1+M)^2}w(r; \gamma), \quad (3.3)$$

and their individual reward rate is

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

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

$$v(t, x) = U(x)e^{-\gamma w(r; \gamma) \frac{\mathcal{K}(T-t)}{(1+M)^2}}. \quad (3.4)$$

²¹With a slight abuse of notation, we use $\gamma = 0$ to label the risk-neutral case of $U(x) = x$. The results of Section 3.1 hold for risk-neutral miners with formal replacement of γ by 0.

²²In the risk-neutral case, $w(r; 0) = r$.

Remark 3.2. Due to the wealth independence of optimal hash rates, and thus the wealth distribution, this is the same solution as a model without the mean field game approximation. Indeed, with $M + 1$ players all using the same (wealth-independent) strategy α^* , the total hash rate in the mean field games model $\alpha^* + M\bar{\alpha}^* = (1 + M)\alpha^*$ is equal to total hash rate in the finite player model: $\sum_{i=0}^M \alpha^* = (1 + M)\alpha^*$. This exact correspondence is specific to this particular setup.

Remark 3.3. By the above solution, we observe that the total mining $\sum_{i=1}^M \alpha^*(t, x)$ is bounded by, and for large M approximately equal to $\frac{\mathcal{K}}{c}w(r; \gamma)$. This is in turn bounded by $\mathcal{K}/c\gamma$ for CARA utility, and, as a consequence, the total mining is thus also bounded by $\mathcal{K}/c\gamma$, regardless of r .

Proof. We guess the form $v(t, x) = U(x)h(t)$ and then the HJB equation (2.7) becomes

$$\begin{cases} \partial_t h - \gamma \left(\sqrt{Mc\bar{\alpha}_t^*} - \sqrt{\mathcal{K}w(r; \gamma)} \right)^2 h = 0, & \text{if } \bar{\alpha}_t^* < \frac{\mathcal{K}}{Mc}w(r; \gamma), \\ \partial_t h = 0, & \text{otherwise,} \end{cases}$$

with terminal condition $h(T) = 1$. It is an equation in t only, which validates our ansatz. Moreover, this ansatz implies that $\Delta v / \partial_x v$ does not depend on x . Hence, from (2.10),

$$\bar{\alpha}_t^* = \frac{\eta^2 M}{(1 + \eta M)^2} \frac{\mathcal{K} \Delta v}{c \partial_x v} \leq \frac{M}{(1 + M)^2} \frac{\mathcal{K} \Delta v}{c \partial_x v} < \frac{\mathcal{K} \Delta v}{Mc \partial_x v},$$

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

The equilibrium mean hash rate is obtained from plugging in the ansatz into (2.10). Thereafter, (2.6) yields (3.3). The value function then satisfies

$$\partial_t h - \frac{\gamma \mathcal{K}}{(1 + M)^2} w(r; \gamma) h = 0,$$

with terminal condition $h(T) = 1$. Thus we have (3.4). \square

3.1.2 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

$$X_{t_0+t} = X_{t_0} - c\alpha^*t + r(N_{t_0+t}^* - N_{t_0}^*) = X_{t_0} - \frac{\mathcal{K}M}{(1 + M)^2} w(r; \gamma)t + r(N_{t_0+t}^* - N_{t_0}^*),$$

where N^* has the constant jump rate $\lambda_t = \frac{\mathcal{K}}{(1+M)}$ and thus is a Poisson process. Then we can compute the expectation and variance of X_{t_0+t} :

$$\mathbb{E}[X_{t_0+t} - X_{t_0}] = \left(r - \frac{Mw(r; \gamma)}{1 + M} \right) \frac{\mathcal{K}t}{(1 + M)}, \quad \text{Var}(X_{t_0+t} - X_{t_0}) = \frac{\mathcal{K}r^2t}{(1 + M)}.$$

If γr is small, the expected wealth change is approximately $\frac{\mathcal{K}rt}{(1+M)^2}$, which decays with competition as M^{-2} . However, the variance is only discounted by a linear factor $1 + M$, and in particular the standard deviation decays as $1/\sqrt{M}$. Hence, as more miners join the game, the expected gain shrinks very fast, but the potential risk decreases slowly: a more crowded competitive mining environment is less attractive in terms of the risk-reward tradeoff than a sparser one.

Iteration	$\frac{1}{T} \sum_{t=0}^T \bar{\alpha}_t - \bar{\alpha}_t^{\text{new}} $	$\frac{1}{T} \sum_{t=0}^T (\Psi(\bar{\alpha})_t - \bar{\alpha}_t^{\text{explicit}}) / \bar{\alpha}_t^{\text{explicit}}$
1	266244	69.673
2	139090	36.398
3	71891	18.813
4	36485	9.5473
5	18382	4.8095
6	9228	2.4139
7	4624	1.2089
8	2314	0.60440
9	1157	0.30166
10	578.6	0.15020
11	289.28	0.07446
12	144.60	0.03659
13	72.281	0.01766
14	36.130	0.00820
15	18.060	0.00347
16	9.0285	0.00111
17	4.5135	0.000068
18	2.2565	0.00065
19	1.1282	0.00095
20	0.5641	0.00110
30	0.00055	0.00124
40	5.5184e-07	0.00124
50	4.2964e-09	0.00124

Table 1: *Example of observed convergence of the numerical method to the explicit solution for CARA utility. Following the notation in Section 2.6, $\bar{\alpha}$ denotes the output of the previous iteration, $\Psi(\bar{\alpha})$ denotes the fixed point mapping output of the current iteration, $\bar{\alpha}^{\text{new}}$ is the relaxed output (2.12), and $\bar{\alpha}^{\text{explicit}}$ denotes the explicit solution from Proposition 3.1. The error remaining in the last column is due to the discretization error of the PDEs and decreases with the grid size.*

3.1.3 Numerical convergence

We use the explicit solution (3.3) to test the numerical scheme from Section 2.6. Table 1 lists errors for one run. The errors are observed to converge toward zero, and as the iteration closes in on the correct solution, the convergence is of order 1 and rate 0.5. Similar numbers are obtained for other starting points and parameter values, although the initial iterations can be slower than rate 0.5 far away from the solution, due to the strong under-relaxation. The relative error of approximately 0.1% in the last column is due to discretization and decreases with the grid size.

This observed convergence indicates that the numerical method is sound, and is therefore valuable indication of the correctness in subsequent sections with CRRA utility, where no explicit solution exists.

3.2 Power Utility (CRRA)

In this section, we consider the mining problem where miners have CRRA utility functions U , defined on $\mathbb{R}_{\geq 0}$, namely

$$U(x) = \frac{1}{1-\gamma} x^{1-\gamma} \quad \text{for } \gamma \in (0, 1), \quad x > 0, \quad (3.5)$$

and admissible strategies α are such that the wealth process X remains positive. The HJB equation (2.7) holds on $x > 0$ with the boundary condition $V(t, 0) = U(0) = 0$. We naturally assume that the initial density m_0 has strictly positive support, so all the miners start with positive wealth. This means that the problem is fully characterized on $\mathbb{R}_{\geq 0}$.

In contrast to the case of exponential utility in Section 3.1, the value function cannot be found explicitly, even with power utility, so we must solve (2.7) numerically. The mean hash rate is still given by (2.10), while the Fokker–Planck equation has two parts to account for the fact that miners' wealths stay strictly positive. The density m solves (2.11) (at least in a weak sense) for $x > r$. On the other hand, if $0 < x < r$, there is no density at $x - r$ jumping to x , because there are no miners with negative wealth. So, for $0 < x < r$, the density m solves

$$\partial_t m - \partial_x (c\alpha^* m) + \frac{\mathcal{K}\alpha^*}{(\alpha^* + M\bar{\alpha}^*)} m = 0, \quad (3.6)$$

where the optimal individual hash rate is given by (2.6). The initial condition on all $x > 0$ is $m(t_0, x) = m_0(x)$.

Remark 3.4. *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,*

$$\begin{aligned} \partial_t \int_{\mathbb{R}_{\geq 0}} m(t, x) dx &= \int_{\mathbb{R}_{\geq 0}} \partial_t m(t, x) dx, \\ &= c\alpha^*(t, \cdot) m(t, \cdot)|_0^r - \int_0^{+\infty} \frac{\mathcal{K}\alpha^*}{(\alpha^* + M\bar{\alpha}_t^*)} m dx + c\alpha^*(t, \cdot) m(t, \cdot)|_r^{+\infty} \\ &\quad + \int_r^{+\infty} \frac{\mathcal{K}\alpha^*(t, x-r)}{(\alpha^*(t, x-r) + M\bar{\alpha}_t^*)} m(t, x-r) dx, \\ &= c\alpha^*(t, \cdot) m(t, \cdot)|_0^{+\infty} = 0, \end{aligned}$$

where the second integral term cancels with the first by making the change of variable $x' = x - r$. The last equation holds because no one can obtain infinite wealth in finite time, and the condition $\alpha^*(t, 0^+) = 0$ holds.

3.2.1 Concentration of wealth and mining effort

In this section, we numerically solve for the equilibrium with power utility functions (3.5). This structure leads to strategic decisions of the miners that are very different from those found with exponential utility, as those with larger wealth tend to hash more. This also puts competitive pressure on lower wealth miners, and as a consequence, the wealthier miners receive a disproportionately large share of profits, a form of *preferential attachment*.

Preferential attachment effects appears 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 empirical study points out that it also appears in the Bitcoin network (Kondor et al., 2014): “we find that the wealth of already rich nodes increases faster than the wealth of nodes with low balance.” This empirical observation is consistent with our model’s numerical results.

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, gradually forming a significant spike. 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 1: *The distribution of miners’ wealth at different times. Parameters: $\mathcal{K}^{-1} = 0.007$, $r = 3$, $c = 2 \times 10^{-5}$, $T = 90$, $\gamma = 0.8$, $M = 1000$. The initial distribution m_0 is normal with mean 90 and standard deviation 5. Note that as time passes, there is mass concentrating on the left hand side of the peak at $x = x_{\text{off}} \approx 77$, cf. Proposition 3.5. In light of Proposition 3.5, this is expected. Below the peak, it is optimal for miners to not participate, whereas above they do. Hence, miners to the right of x_{off} move to the left, whereas those at or below are stationary, causing concentration of the density around x_{off} . We emphasize that this does not cause numerical problems, because derivatives are only evaluated on one side of this point.*

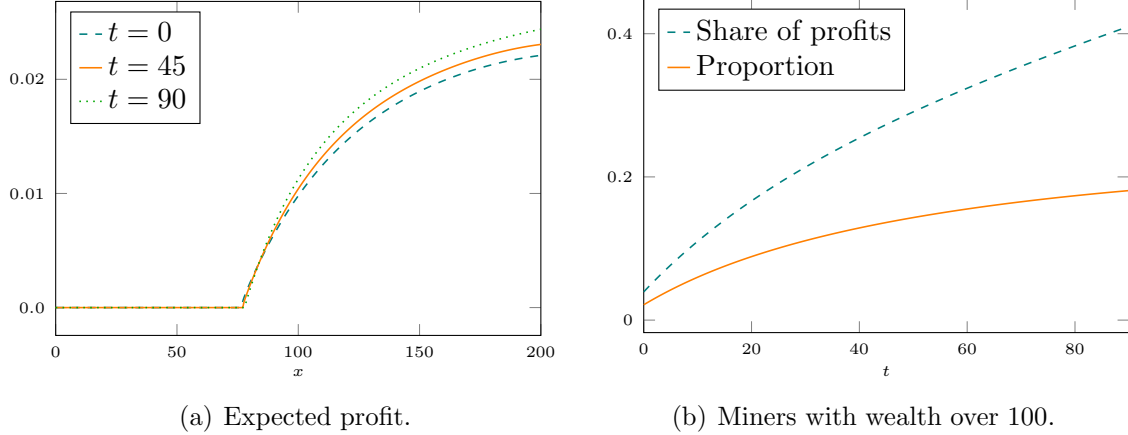


Figure 2: The left plot shows the expected instantaneous profit miners at different wealth levels and times. The right one gives the proportion of miners with wealth over 100 and their share of the total instantaneous profits. Parameters are the same as Figure 1.

From (2.3), the expected profit $\mathbb{E}[\Delta X_t]$ for an individual miner over a short time Δt is given, approximately by $\mathbb{E}[\Delta X_t] \approx -c\alpha_t^* \Delta t + r\lambda_t \Delta t$, where λ_t is given by (2.2). Figure 2(a) shows the instantaneous profit rate, namely

$$-c\alpha^*(t, x) + r\mathcal{K} \frac{\alpha^*(t, x)}{\alpha^*(t, x) + M\bar{\alpha}_t^*}, \quad (3.7)$$


as a function of miners' wealths at various times. We see that the more wealth a miner has, the higher is their expected profit rate. Miners with lower wealth hash at lower rates or even zero rate below x_{off} introduced in Proposition 3.5 below. This pattern holds at all times. In addition, it is interesting to see that the miners are blockaded (0 hash rate) around and below wealth level 77 (see Figure 2(a)). Their risk aversion prevents miners with small wealth from participating in the mining game. This is explained by the following lemma, whose proof is provided in Appendix A.

Proposition 3.5. *For any time t and equilibrium hash rate $\bar{\alpha}_t^* > 0$, there exists a wealth level $x_{\text{off}}(t) > 0$ such that $\alpha^*(t, x) = 0$ for $x \leq x_{\text{off}}(t)$, i.e., the optimal mining rate is zero at wealth levels below x_{off} .*


To better understand the preferential attachment effect, we compute the proportion of miners whose wealth is over 100 and 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 reward r and competition parameter M affect this phenomenon. To avoid repetition, we show only the plots at $t = 30$, but a similar pattern is present also at other times.


A larger reward r 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




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



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



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



(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 r takes multiple values.*

higher. At the same time, the density for $x \geq 110$ is clearly higher for the larger reward in Figures 3 (a)(b)(c). Figure 3(d) shows the expected instantaneous profits (3.7), which leads to the same conclusion.

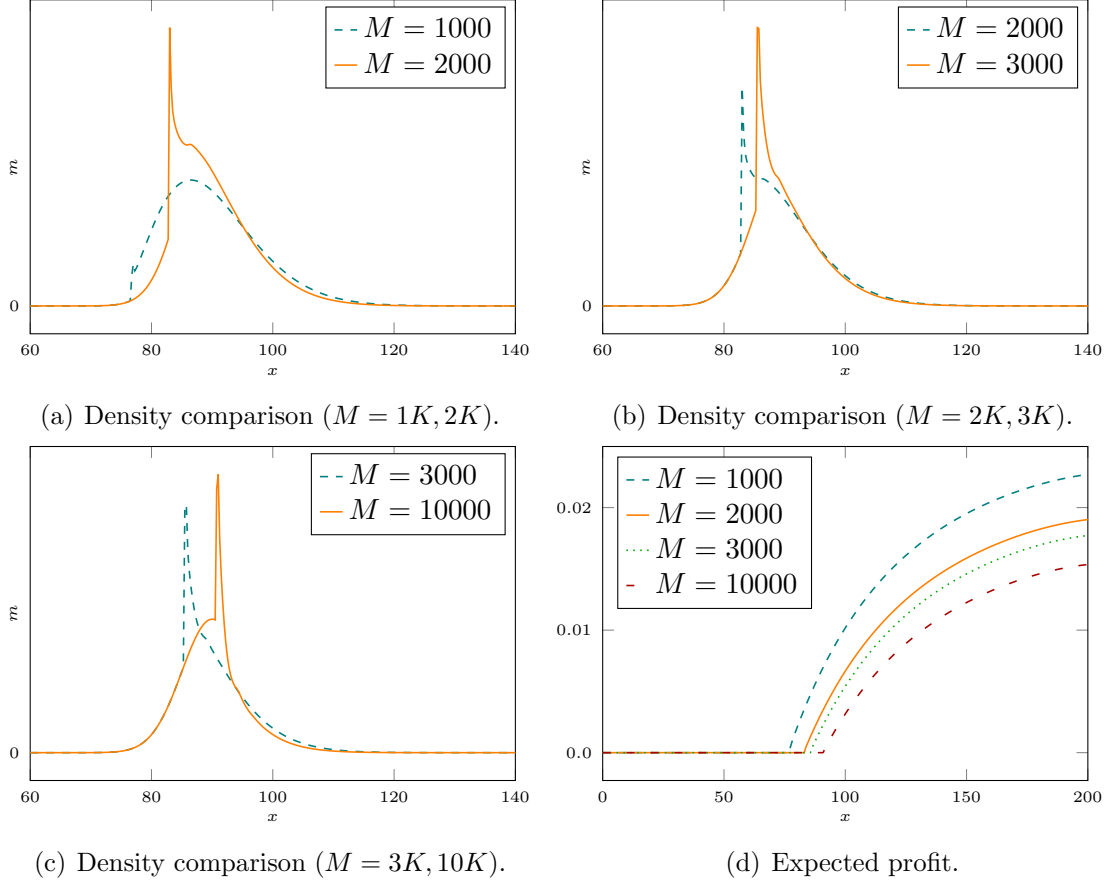


Figure 4: *The competition effects at $t = 30$. The first three plots show the distribution of miners' wealth. The last one shows the expected instantaneous profit for miners at different wealth levels, for four different competition levels. Parameters are the same as Figure 1 except that M takes multiple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

The competition parameter M reduces the preferential attachment effect. This is shown in Figures 4 (a)(b)(c), as the density for $x \geq 100$ is lower for larger M . Additionally, the expected profit rate (3.7) decrease with respect to M , as is shown in Figure 4(d). Meanwhile, as the competition becomes fierce and the participation threshold for miners to be active becomes larger. When $M = 1000$, miners need wealth around 75 to have incentive to hash, but this increases to about 90 for $M = 10000$. Hence, the competition makes the mining game less lucrative, and makes it harder for miners to stay active, which reduces the preferential attachment.

4 Competition with cost advantages

In this section, we consider a model in which one miner has 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 advantaged miner. It utilizes cheaper electricity in China, like the hydropower stations in Sichuan during the rainy season, and also Hydro Quebec in Canada, which offers some of the lowest electricity rates in North America.²³ 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.

Because the advantaged miner is distinct from the other miners, the general structure of Section 2 is adapted in Section 4.1. As in the earlier model, we are able to obtain the equilibrium explicitly when the smaller miners have exponential utility (Section 4.2), and we can demonstrate wealth effects numerically when they have power utility (Section 4.3).




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>.*

4.1 Problem formulation

We consider a cost-advantaged miner competing with $M + 1$ individual miners who are approximated, as in Section 2.3, by a continuum of miners.²⁴ This miner chooses a hash rate β_t , and incurs cost $c_1\beta_t = k_c c\beta_t$ per hash, where $0 < k_c \leq 1$ is the relative cost efficiency. Given the (continuum) mean hash rate $\bar{\alpha}_t$ of the individual miners, the counting process N_t^1

²³<https://www.reuters.com/article/us-canada-bitcoin-china/chinese-bitcoin-miners-eye-sites-in-energy->

²⁴This 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.

with intensity

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

denotes the number of rewards received by the cost-advantaged miner.

We think of the advantaged miner as a profit-maximizing firm, or, more generally, having preferences described by a utility function $U_1(x)$ coupled with the appropriate set of admissible hash rate strategies. The objective for the cost-advantaged miner is

$$\sup_{\beta_t \geq 0} \mathbb{E} \left[U_1 \left(\int_0^T -c_1 \beta_t dt + r dN_t^1 \right) \right]. \quad (4.1)$$

As in Section 2.3, the model for individual miners remains the same except that the intensity for N_t in (2.3) becomes

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

given the cost-advantaged miner's hash rate β_t . Here the denominator consists of both the cost-advantaged miner's hash rate and the total of individual miners.

In this generality, we would need an additional state variable for the wealth of the cost-advantaged minor, as well as their value function, increasing the dimension and complexity of the PDE system. However, in the case where the players have exponential utility (or are risk-neutral), the equilibrium hash rates are wealth-independent, and we can solve explicitly (Section 4.2). We can further handle numerically the case of the individual miners having power utility when the advantaged miner is profit-maximizing, i.e. risk-neutral, (Section 4.3), in which case we do not need an additional state variable. In either of these cases then, any Markov control can be identified by a function $\beta(t; \bar{\alpha})$, i.e., the control only depends on time.

4.1.1 HJB equation

The value function defined as in (2.4) now depends on both $\bar{\alpha}$ and β , i.e., $v(t_0, x; \bar{\alpha}, \beta)$. For a fixed choice of $\bar{\alpha} > 0$ and $\beta \geq 0$, the HJB can be written as

$$\partial_t v + \sup_{\alpha \geq 0} \left(-c\alpha \partial_x v + \frac{\mathcal{K}\alpha}{(\alpha + M\bar{\alpha}_t + \beta_t)} \Delta v \right) = 0,$$

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

$$\alpha^*(t, x; \bar{\alpha}, \beta) = \begin{cases} -(M\bar{\alpha}_t + \beta_t) + \sqrt{\frac{(M\bar{\alpha}_t + \beta_t)\mathcal{K}\Delta v}{c\partial_x v}}, & \text{if } M\bar{\alpha}_t + \beta_t < \frac{\mathcal{K}\Delta v}{c\partial_x v}, \\ 0 & \text{otherwise,} \end{cases} \quad (4.2)$$

and the HJB equation is

$$\begin{cases} \partial_t v + \left(\sqrt{c(M\bar{\alpha}_t + \beta_t)\partial_x v} - \sqrt{\mathcal{K}\Delta v} \right)^2 = 0, & \text{if } M\bar{\alpha}_t + \beta_t < \frac{\mathcal{K}\Delta v}{c\partial_x v}, \\ \partial_t v = 0 & \text{otherwise.} \end{cases} \quad (4.3)$$

4.1.2 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 a cost-advantaged 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 $\beta_t^* = \beta^*(t; \bar{\alpha}^*)$, from (4.1). Henceforth, let $\bar{\alpha}^*$ and β^* denote the equilibrium mean hash rate and equilibrium hash rate for the cost-advantaged 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 2.5, it is meaningful to consider $\bar{\alpha}_t^* > 0$ for all t . Thus in the equilibrium, we have coupled equations $\beta_t^* = \beta^*(t; \bar{\alpha}^*)$ and

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

by integrating (4.2) in x over the set (2.9). The Fokker-Planck equation is given by

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

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

4.2 Exponential and risk-neutral utility miners: explicit solution

Define

$$w(r; \gamma) = \begin{cases} \frac{(1 - e^{-\gamma r})}{\gamma} & \gamma > 0, \\ r & \gamma = 0. \end{cases}$$

When the higher cost miners have exponential utility ($\gamma > 0$) or are risk-neutral ($\gamma = 0$), and the cost-advantaged miner has exponential utility, or is risk-neutral ($U_1(x) = x$), we have the following result.

Proposition 4.1. *Suppose the individual miners have exponential utility $U(x) = -\frac{1}{\gamma}e^{-\gamma x}$ with $\gamma \geq 0$ and the cost-advantaged miner either i) has exponential utility with risk-aversion coefficient $\gamma_1 > 0$; or ii) is risk-neutral ($\gamma_1 = 0$). Suppose the relative cost efficiency satisfies*

$$k_c < \frac{w(r; \gamma_1)}{w(r; \gamma)} \frac{M+1}{M}, \quad (4.4)$$

and let

$$\kappa = \frac{\mathcal{K}w(r; \gamma)}{c}, \quad \kappa_1 = \frac{(M+1)\mathcal{K}w(r; \gamma_1)}{c_1}.$$

Then, in equilibrium, all miners are active with

$$\alpha^*(t, x) \equiv \bar{\alpha}_t^* \equiv \frac{\kappa^2 \kappa_1}{(\kappa + \kappa_1)^2} > 0, \quad \beta_t^* \equiv \frac{\kappa \kappa_1 (\kappa_1 - M\kappa)}{(\kappa + \kappa_1)^2} > 0, \quad (4.5)$$

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

Proof. With the ansatz $v(t, x) = U(x)h(t)$, the HJB (4.3) reduces to solving

$$\begin{cases} \partial_t h - \gamma \left(\sqrt{c(M\bar{\alpha}_t^* + \beta_t^*)} - \sqrt{\mathcal{K}w(r; \gamma)} \right)^2 h = 0, & \text{if } M\bar{\alpha}_t^* + \beta_t^* < \frac{\mathcal{K}}{c}w(r; \gamma), \\ \partial_t h = 0 & \text{otherwise,} \end{cases}$$

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 (4.7), and, using the ansatz in (4.2), we find

$$\alpha^*(t, x; \bar{\alpha}^*, \beta^*) = -(M\bar{\alpha}_t^* + \beta_t^*) + \sqrt{\frac{\mathcal{K}}{c}w(r; \gamma)(M\bar{\alpha}_t^* + \beta_t^*)}.$$

Since α^* does not depend on the wealth x , all individual miners are active. Therefore, we have $\bar{\alpha}_t^* = \alpha_t^*$. This, together with (4.7), yields (4.5). It is direct that α^* is positive, and β^* is positive if and only if (4.4) holds. Thus we have found the equilibrium in which everyone is active. \square

Cost advantage and its effect on mining power concentration

Proposition 4.1 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 (4.5) is increasing in κ_1 and hence decreasing in c_1 . Similarly, the hash rate α^* in (4.5) of the individual miners increases with respect to c_1 . Thus, a smaller c_1 —a larger cost advantage—makes the advantaged miner more dominant. As a consequence, individual miners with higher cost have to decrease their hash rates to regulate their risk exposure, as the advantaged miner gets a larger share of the mining rewards.

To quantify this, we write $\rho = k_c \frac{w(r; \gamma)}{w(r; \gamma_1)}$. Then $\kappa_1 = (M + 1)\kappa/\rho$. The probability for the cost-advantaged miner to get the reward is

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

for large M , whereas the remaining miners have a collective success probability ρ and individual probability $\rho/(M + 1)$.

If $\gamma \approx \gamma_1$, this simplifies further because $\rho \approx k_c$, i.e., the advantaged miner's share of the market is approximately their advantage $1 - k_c$. If the cost-advantaged miner is 10% efficient ($k_c = 0.9$), then $\rho \approx 0.9$, which gives a probability around 10% for the cost-advantaged miner to get the reward.

This analysis shows that the advantaged miner receives an disproportionately large share of block rewards. But because this miner pays a lower cost per hash, the profit increases even further. Both of these effects remain in the power utility setting of the next section and are explored in more detail in Figure 7 and the accompanying text.

4.3 Power utility

We now consider wealth effects by endowing the individual miners with power utility preferences (3.5). As discussed in the final paragraph of Section 4.1, we assume the cost-advantaged miner is profit-maximizing, which keeps the problem tractable without the need




Figure 6: *The distribution of miners' wealth at $t = 0, 30, 45$, with $k_c = 0.8$ for the distributions with an advantaged miner. The parameters are the same as in Figure 1.*

to introduce an additional state variable. Therefore (4.1) becomes

$$\sup_{\beta_t \geq 0} \mathbb{E} \left[\int_0^T -c_1 \beta_t dt + r dN_t^1 \right]. \quad (4.6)$$

Given $\bar{\alpha} > 0$, the maximizer in (4.6) satisfies the first-order condition


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

which yields the best response

$$\beta^*(t; \bar{\alpha}) = \begin{cases} -(M+1)\bar{\alpha}_t + \sqrt{\frac{r\mathcal{K}(M+1)\bar{\alpha}_t}{c_1}}, & \text{if } \bar{\alpha}_t < \frac{r\mathcal{K}}{c_1(M+1)}, \\ 0, & \text{otherwise.} \end{cases} \quad (4.7)$$

Like in Section 3.2, we solve the model numerically. The numerical method correspondingly follows the procedure in Section 2.6 with appropriate changes adding β .

Figure 6 shows the wealth distribution of individual miners when $k_c = 0.8$ for $t = 0, 30, 45$. Two time points from Figure 1 are also plotted, which show two subtle differences. Recall



(a) The advantaged miner's share of total re- (b) The advantaged miner's share of total profits.
wards.

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 in Figure 1.*

from Section 3.2.1 and Figure 1 that the left peak of each curve is the accumulation of miners at the point where mining is no longer optimal. Most clearly visible is the right-shift of this peak, in the presence of an advantaged miner. This means that the minimum wealth necessary for non-zero mining to be incentivized is greater with an advantaged miner, thus increasing centralization. The other effect is the slightly slower dispersion of the distribution, which is explained by the lower mining rate across the board as a result of hesitancy due to the added competition from the advantaged miner. Nevertheless, the type of evolution in Figure 1 is observable also in Figure 6.

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. Despite moving to power utility mining, the relationship is almost identical to that found analytically with exponential utility in Section 4.2: The hash rate share taken by the advantaged miner is approximately $1 - k_c$. We see that as the cost efficiency moves from $k_c = 1.0$ to 0.65, the advantaged miner's hash rate increases from around 1% to 35%. This is a strong effect, and hence the cost advantages could be one explanatory factor for the concentration of mining power. A similar idea also appears in (Arnosti & 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. Alsabab and Capponi (2020) 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 higher 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 small difference to the risk aversion, as these numbers are for $\gamma = 0.8$ and $\gamma_1 = 0$.

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.

5 Conclusion and further research

This paper develops stochastic models for the mining of cryptocurrencies that implement a proof-of-work protocol. As miners compete through their computation efforts—their hash rates—a rational utility maximization framework is analyzed as a novel mean field game of intensity control. Remarkably, the equilibrium is found explicitly for certain utility functions—one of the very few explicit solutions of mean field games outside the linear quadratic framework. In other cases, the equilibrium can be found numerically and efficiently by methods described in this paper.

Our main finding is that, as a result of initial wealth heterogeneity among the miners, more rewards tend to be collected by those who have more wealth, while miners with lesser wealth may be disincentivized from participating in mining activity at all. This leads to a feedback effect of greater wealth and more concentrated mining. Concentration goes against the very fundamental principle of mining being decentralized, which is a requisite for the security of the cryptocurrency.

Incorporating a player with even a slight cost-advantage into our game shows that they become relatively dominant. This is consistent with how these professional miners with more advanced equipment or access to cheaper electricity have come to account for a significant share of mining power in recent years.

There are plenty of directions for future research that can build on the approach to cryptocurrency mining presented in this paper. We present some preliminary findings on three of these directions in the appendices, as we now describe.

At the end of 2010, the first mining pool was announced. 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. Miners can join a pool, which collects their computational power to do the mining. Once the pool gets a reward, it is shared among the miners. Meanwhile, these companies also contribute a significant proportion of hash rates to their own pools. Although pools are often public, meaning miners may participate and receive a share of pool revenue, the pool operators control what block data (the ‘work’) is processed by its participants. This means that a few entities account for selecting the block data (transactions) processed by a large share of hash rates in the world. An initial discussion of how our model can be adapted to incorporate mining pools is presented in Appendix D.

A notorious feature of cryptocurrencies, in particular bitcoin, is wild changes in their prices, which can run to speculative highs followed by sudden crashes. In Appendix C, we explore the effect of stochastic rewards, for instance due to price fluctuations, in an extension of our model. In the case of exponential utility miners, we find miners react to the level but not the volatility of prices.

We have assumed homogeneous costs and preferences in our model, which may be relaxed as we demonstrate (for costs) in Appendix E. In general, costs would depend on miners’ geographic location, access to silicon and hardware, etc. Bill Tai of Hut 8 Mining Corp. has

stated that big miners (like Bitfury) can buy at discount thanks to being able to “buy silicon in large quantities and commit to the electricity grid in chunk sizes.”²⁵ This suggests that per unit costs should be smaller for large miners. Such concave costs would amplify the preferential attachment effects we find already in the linear costs model. Another avenue for further research is to include fixed costs of entering the mining game, as motivated in Garratt and van Oordt (2020).

Appendix A Proofs

A.1 Proof of Lemma 2.1

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

$$\mathbb{E}[U(X_T)|X_t = x] \leq U(\mathbb{E}[X_T|X_t = x]) \leq U(x + r\mathbb{E}[N_T - N_t|X_t = x]) \leq U(x + r\mathcal{K}T) < \infty.$$

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 $x_2 - x_1$. 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 Proposition 3.5

Since α can be chosen identically zero on $[t, T]$ and the optimal strategy does at least as well as doing nothing, for any $\varepsilon > 0$, we have

$$\frac{v(t, \varepsilon) - v(t, 0)}{\varepsilon} \geq \frac{U(\varepsilon) - U(0)}{\varepsilon} \xrightarrow{\varepsilon \rightarrow 0} \infty,$$

from the derivative of the power utility function (3.5). Hence, $\partial_x v(t, x) \rightarrow \infty$ as $x \rightarrow 0$ and consequently, $\Delta v / \partial_x v$ is arbitrarily small in some neighborhood of 0, because $\Delta v(t, 0) \leq v(t, r) - v(t, 0) = v(t, r) < \infty$. Thus, by (2.6), there exists a $x_{\text{off}}(t)$ such that zero rate mining is optimal for $x \leq x_{\text{off}}$.

Appendix B Mean field games (of controls)

We here describe the ideas behind mean field games and the arguments leading to the continuum limit. Most mean field games models exhibit interaction between the players through their state. In contrast, the reward rate to each player in the proof-of-work mining game depends only on the hash rate of the population. Consequently, the interaction between players is through their *actions* as opposed to their state. This is commonly referred to as

²⁵<https://www.bloomberg.com/news/articles/2018-04-18/bitcoin-miners-facing-a-shakeout-as-profitability>

extended mean field games or *mean field games of controls*. Conceptually, the state and action interaction types of mean field games are very similar, but the structure of the resulting equations are very different.

Consider a game of N players, each ($i \in 1, \dots, N$) with control α^i and state $X^i = X^{i, \alpha^i}$ described by the evolution

$$dX_t^i = f\left(\alpha_t^i, X_t^i, \sum_{j=1}^N \alpha_t^j / N\right) dt + dW_t^i, \quad X_0^i = x^i, \quad (\text{B.1})$$

where W^i is a Brownian motion representing player i 's idiosyncratic noise (W^i and W^j are independent for $i \neq j$). Given a strategy profile $\alpha^{-i} = (\alpha^1, \dots, \alpha^{i-1}, \alpha^{i+1}, \dots, \alpha^N)$ of the other players, player i is optimizing some quantity $\mathbb{E}[U(X_T^{i, \alpha^i})]$, and we assume that the optimizer $\alpha^{i, *}$ can be found in the class of Markovian controls, i.e., that it is a function of player i 's state (for fixed α^{-i}). In other words, we may write $\alpha_t^{i, *} = \alpha^*(t, X_t^i, \alpha_t^{-i})$.

An equilibrium $t \mapsto \alpha_t = (\alpha^*(t, X_t^1, \alpha_t^{-1}), \dots, \alpha^*(t, X_t^N, \alpha_t^{-N}))$ of this form is characterized by the property that for all i , $t \mapsto \alpha^*(t, X_t^i, \alpha_t^{-i})$ is an optimizer to player i 's optimization problem, i.e., no player has anything to gain from deviating. We will not delve into detail about why these games become very difficult to solve, but the gist of it is that each player must solve an HJB equation that depends on every other player's solution, thus creating a system of N coupled equations. The idea of mean field games enters to circumvent this dimensionality issue.

The first step is to consider a sequence of games with increasingly many players. Each additional player must be assigned an initial state (see state x^i for player i in (B.1)). This is done by sampling the initial states from a distribution (independent from W^i) with density m_0 . We now consider what happens to the problem in the limit $N \rightarrow \infty$. Because the probability of two players starting at the same state is zero, we may index each player by their state x^i instead of i . In the limiting problem we simply write x and remember that the collection of initial states x is distributed according to m_0 . As time passes, the players' states evolve. We denote by $m(t, \cdot)$ the density of players at time t .²⁶

For any strategy profile $\alpha = (\alpha^{x^1}, \dots, \alpha^{x^N})$, we consider the problem of player x^i . The first observation is that if player x^i deviates from α^{x^i} , the impact on the other players is of order $1/N$, due to the averaging effect in (B.1). In particular, as $N \rightarrow \infty$, we see that the effect of any one player on the others is vanishing, as

$$f\left(\alpha_t^{x^i}, X_t^{x^i}, \sum_{j=1}^N \alpha_t^{x^j} / N\right) \longrightarrow f\left(\alpha_t^{x^i}, X_t^{x^i}, \int_{\mathbb{R}} \alpha_t^{x^j} m(t, x^j) dx^j\right).$$

Hence, in the limit, player x^i 's action and state alone do not impact the population average. Moreover, the dependence of player x^i on the population is only through a statistical property: the mean. Dropping the superscript, we may thus consider the dynamics of player x given the mean control $\bar{\alpha}$ of the population:

$$dX^{x, \alpha^x} = dX_t^x = f(\alpha_t^x, X_t^x, \bar{\alpha}_t) dt + dW_t^x, \quad X_0^x = x.$$

Player x seeks to optimize $\mathbb{E}[U(X_T^x)]$, given the (mean) strategy profile $\bar{\alpha}$. For any fixed $\bar{\alpha}$, this is a standard control problem, and we again write the optimizer as a function of the current state and the actions of the other players: $\alpha^*(t, X_t^x, \bar{\alpha}_t)$.

²⁶This density does depend on the actions, but for now we omit this in the notation.

If every player chooses the action $\alpha_t^* = \alpha^*(t, X_t^x, \bar{\alpha}_t)$, we denote by $m(t, \cdot; \alpha^*) = m(t, \cdot; \alpha^*, m_0)$ the resulting density at time t . The process $\bar{\alpha}$ is the equilibrium mean control if

$$\bar{\alpha}_t = \int_{\mathbb{R}} \alpha^*(t, x, \bar{\alpha}_t) m(t, x; \alpha^*) dx. \quad (\text{B.2})$$

In other words, if all other players are using the strategy α^* , then no players can improve their situation by deviating from this strategy.

Finally, given any process $\bar{\alpha}$ and a strategy $\alpha(t, x)$ shared by all players, the density $m(t, \cdot; \alpha)$ of the population can be shown to satisfy the Fokker–Planck equation

$$\partial_t m(t, x; \alpha) - \partial_x (f(t, \alpha(t, x), x, \bar{\alpha}) m(t, x; \alpha)) - \frac{1}{2} \partial_{xx} m(t, x; \alpha) = 0, \quad m(0, \cdot, \alpha) = m_0. \quad (\text{B.3})$$

Furthermore, the value function v of each player's optimization problem is characterized by the HJB equation

$$\partial_t v + \sup_{\alpha} f(t, \alpha, x, \bar{\alpha}) \partial_x v + \frac{1}{2} \partial_{xx} v = 0, \quad v(T, \cdot) = U. \quad (\text{B.4})$$

To summarize: by considering the limit as $N \rightarrow \infty$, the system of N coupled HJB equations reduces to the two coupled equations (B.3) and (B.4) along with the equilibrium fixed point condition (B.2). Although the structure of this smaller system is still mathematically very complex, it is often amenable to numerical computations, in contrast to the computational intractability of the N -player system.

Appendix C Stochastic reward & exponential utility

We consider the base model of Section 2, but when the reward r is stochastically varying over time, driven, for instance, by the price of bitcoin. Similar to Section 3.1, we can solve analytically for the equilibrium when the miners have exponential utility. A miner's wealth process $X = (X_t)_{t \geq t_0}$ follows: $dX_t = -c\alpha_t dt + r_t dN_t$, where now $r > 0$ is modeled as a Markov process with infinitesimal generator \mathcal{L}_r , whose increments are independent of the contemporaneous increments dN . For instance, if r is fit to a geometric Brownian motion model: $dr_t = \mu r_t dt + \sigma r_t dW_t$, where W is an independent Brownian motion and σ is the reward volatility, then $\mathcal{L}_r = \frac{1}{2} c_1^2 r^2 \partial_{rr} + c_0 r \partial_r$.

Consider exponential utility miners, i.e., $U(x) := -\frac{1}{\gamma} e^{-\gamma x}$. Then the value function of an individual miner, which is a function of t , x , and r , is defined by

$$v(t, x, r; \bar{\alpha}) = \sup_{\alpha \geq 0} \mathbb{E}[U(X_T) | X_t = x, r_t = r].$$

The associated HJB equation is

$$\partial_t v + \mathcal{L}_r v + \sup_{\alpha \geq 0} \left(-c\alpha \partial_x v + \frac{\mathcal{K}\alpha}{(\alpha + M\bar{\alpha}_t(r))} \Delta v \right) = 0,$$

with terminal condition $v(T, x, r) = U(x)$, and where $\Delta v = v(t, x + r, r; \bar{\alpha}) - v(t, x, r; \bar{\alpha})$.

Adapting the analysis of Section 2.4, in equilibrium the optimal hash rate α^* and $\bar{\alpha}^*$ are related by

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

and the HJB equation can be written as:

$$\partial_t v + \mathcal{L}_r v + \left(\sqrt{Mc\bar{\alpha}_t^*(r)}\partial_x v - \sqrt{\mathcal{K}\Delta v} \right)^2 = 0.$$

With exponential utility miners, we have $v(T, x, r) = U(x) := -\frac{1}{\gamma}e^{-\gamma x}$. With the ansatz $v(t, x, r) = g(t, r)U(x)$, we find that g must solve the linear PDE problem

$$\partial_t g + \mathcal{L}_r g - \gamma \left(\sqrt{Mc\bar{\alpha}_t^*(r)} - \sqrt{\mathcal{K}w(r; \gamma)} \right)^2 g = 0, \quad g(T, r) = 1,$$

where w was defined in (3.2). The ansatz implies that $\Delta v / \partial_x v = w(r; \gamma)$ does not depend on x . Then, as in the proof of Proposition 3.1, all miners are active and so

$$\bar{\alpha}_t^*(r) = \frac{\mathcal{K}M}{c(1+M)^2}w(r; \gamma).$$

This yields the linear PDE problem for g :

$$\partial_t g + \mathcal{L}_r g - \gamma \frac{\mathcal{K}w(r; \gamma)}{(1+M)^2} g = 0,$$

with terminal condition $g(T, r) = 1$. The solution can be written as a Feynman–Kac expectation:

$$g(t, r) = \mathbb{E} \left[\exp \left(-\gamma \int_t^T \frac{\mathcal{K}w(r_s; \gamma)}{(1+M)^2} ds \right) \mid r_t = r \right],$$

but is not needed for the optimal mining rate, which is given by

$$\alpha^*(t, x, r; \bar{\alpha}^*) \equiv \bar{\alpha}_t^*(r) \equiv \frac{\mathcal{K}M}{c(1+M)^2}w(r; \gamma). \quad (\text{C.1})$$

The individual reward rate is $\lambda_t \equiv \frac{1}{D(1+M)}$, which does not depend on r . So, even as the price appreciates or falls, the hash rates go up and down respectively, following the formula (C.1), but the individual rewards rates remain the same. The reward affects the value function as higher reward gives higher reward, but reward volatility σ does not enter into the equilibrium hash rates (C.1), which are simply reactive to the present reward.

Appendix D Mining pools

An important property of mining pools is their use for risk sharing. Miners may join pools to reduce risk by sharing rewards among all pool contributors, in some proportion to

their contribution. Each miner thus receives smaller but more frequent rewards,²⁷ thereby reducing the risk associated with reward arrival. In fact, a miner may join multiple pools and contribute hashes to the different pools. This enables further risk reduction.

However, pool miners no longer have the power to select which transactions to process as part of their blocks.²⁸ Instead the block data—the ‘work’—is handed down to them by pool operators. So even if the risk sharing mechanism of mining pools improves the resilience of the blockchain to outside attacks, it does come at a cost to decentralization, and thus lesser resilience to malicious behavior among its participants.

We here consider a model with pools and miners who desire the risk sharing offered by pools. The decision making problem of pools becomes a problem very similar to that studied in the paper, suggesting that similar competition and centralization effects would appear on a pool level.

There exists a variety of reward distribution schemes, fee structures, and other features that differentiate pools. To not overcomplicate the model, we consider pools that distribute rewards to their contributors proportional to their hash rate contribution.²⁹ We also assume that all pools are public, i.e., anybody may join them, and that all miners who are not operating a pool distribute their contributions across pools to reduce risk. There are thus two groups of miners: pool operators and individual miners.

The assumptions that all pools are public and that everybody mines for pools allows individual miners perfect risk sharing, except for the inherent randomness of the arrival of the next block. This is accomplished by mining for each pool in proportion to the pool’s total contribution to the global hash rate. We denote by α^{ind} the aggregate hash rate of the individual miners, and by α^p the hash rate of the operator of pool $p = 1, \dots, P$. The hash rate contribution of individual miners to pool p is then

$$\alpha^{\text{ind}} \frac{\alpha^p}{\sum_{p'=1}^P \alpha^{p'}}.$$

Given the individual miners’ contributions to the pool, each pool’s hash rate is

$$\alpha^p + \alpha^{\text{ind}} \frac{\alpha^p}{\sum_{p'=1}^P \alpha^{p'}} = \alpha^p \left(1 + \alpha^{\text{ind}} \frac{1}{\sum_{p'=1}^P \alpha^{p'}} \right),$$

and the aggregate is

$$\sum_{p=1}^P \alpha^p \left(1 + \alpha^{\text{ind}} \frac{1}{\sum_{p'=1}^P \alpha^{p'}} \right).$$

A pool’s block arrival rate is (proportional to) the ratio of the two which, after simplification, is

$$\frac{\alpha^p \left(1 + \alpha^{\text{ind}} \frac{1}{\sum_{p'=1}^P \alpha^{p'}} \right)}{\sum_{p'=1}^P \alpha^{p'} \left(1 + \alpha^{\text{ind}} \frac{1}{\sum_{p'=1}^P \alpha^{p'}} \right)} = \frac{\alpha^p}{\sum_{p'=1}^P \alpha^{p'}} = \frac{\alpha^p}{\alpha^p + \sum_{\substack{p'=1 \\ p' \neq p}}^P \alpha^{p'}}.$$

²⁷Some pools even employ a system in which members are paid for hash contributions regardless of reward arrival, thus moving all risk to the pool. Such pools charge higher fees to offset this risk. As mentioned towards the end of this section, such a structure fits well into the analysis in the main text.

²⁸The decentralized mining pool P2Pool is an example of an exception to this.

²⁹More commonly, pools use a ‘Pay Per Last N Shares’ (PPLNS) system that discourages contributing to multiple pools. Such a system would see more centralization than the model set up here.

Hence, when individual miners strive for risk sharing, the reward arrival for pools has the same structure as without pools.

Suppose a pool charges a fee ψr when a block is mined. The pool's reward is then

$$\begin{aligned} r(\text{share in pool}) + r\psi(\text{share of participants}) &= r \frac{\alpha^p}{\alpha^p + \alpha^{\text{ind}} \frac{\alpha^p}{\sum_{p'=1}^P \alpha^{p'}}} + r\psi \frac{\alpha^{\text{ind}} \frac{\alpha^p}{\sum_{p'=1}^P \alpha^{p'}}}{\alpha^p + \alpha^{\text{ind}} \frac{\alpha^p}{\sum_{p'=1}^P \alpha^{p'}}} \\ &= r \left(1 + (\psi - 1) \frac{\alpha^{\text{ind}}}{\sum_{p'=1}^P \alpha^{p'} + \alpha^{\text{ind}}} \right). \end{aligned}$$

This reflects that the pool operator must share the reward with participants in proportion to the relative hash contributions.

We thus have a model for the pool manager's wealth as

$$dX_t^p = -c\alpha_t^p dt + \tilde{r}_t^p dN_t^p,$$

where N_t^p has intensity

$$\lambda^p = \frac{\alpha_t^p}{\alpha_t^p + \sum_{\substack{p'=1 \\ p' \neq p}}^P \alpha_t^{p'}} \approx \frac{\alpha_t^p}{\alpha_t^p + (P-1)\bar{\alpha}_t}$$

and

$$\tilde{r}_t^p = r \left(1 + (\psi - 1) \frac{\alpha_t^{\text{ind}}}{\sum_{p'=1}^P \alpha_t^{p'} + \alpha_t^{\text{ind}}} \right) \approx r \left(1 + (\psi - 1) \frac{\alpha_t^{\text{ind}}}{\alpha_t^p + (P-1)\bar{\alpha}_t + \alpha_t^{\text{ind}}} \right),$$

where $\bar{\alpha}$ is the continuum approximation of the average pool manager hash rate.

The continuum mean field game version is to consider a representative pool manager with initial wealth $X_0 = x$ and hash rate α_t , and wealth evolving according to

$$dX_t = -c\alpha_t dt + \tilde{r}_t dN_t, \quad \lambda_t = \frac{\alpha_t}{\alpha_t + (P-1)\bar{\alpha}_t}, \quad \tilde{r}_t = r \left(1 + (\psi - 1) \frac{\alpha_t^{\text{ind}}}{\alpha_t + (P-1)\bar{\alpha}_t + \alpha_t^{\text{ind}}} \right). \quad (\text{D.1})$$

The pool game (D.1) differs slightly from our model in Section 2.3, where r is constant. However, because in the expression for \tilde{r} , α only appears in the denominator, along with both α^{ind} and all other pool operators' hash rates, the representative operator has very little control over the reward size. Thus, the choice of α primarily impacts the reward arrival rate λ , where it appears also in the numerator. In other words, this model is very similar to that studied in this paper and can be treated in the same way.

The similarity of this model to that without pools suggests the competition and centralization would be similar among pool operators after taking pools into account. In particular, the previous results suggest the emergence of dominant pools, which we also observe in practice. This section assumes a proportional distribution of rewards to participants. If the pools instead employ a pay per share (PPS) system, which means miners receive a guaranteed payment per hash, pools essentially buy hashes at a cost equal to that payment. This cost can be interpreted as c in our model, which leads to a game even more similar to that in the main text.

Of course, a more complete model of pools might include competition in setting the fee ψ . For further discussion and analysis of pooled mining in a static setting, we refer to Cong et al. (2021).

Appendix E Cost heterogeneity

Recall that in the model of Section 3.1, when the individual miners have CARA utility $\Delta v / \partial_x v$ is constant. Consequently, α^* in (2.6) becomes

$$\alpha^* = \begin{cases} -M\bar{\alpha} + \sqrt{\frac{w(r;\gamma)\mathcal{K}M\bar{\alpha}}{c}}, & \text{if } \bar{\alpha} < \frac{w(r;\gamma)\mathcal{K}}{cM}, \\ 0, & \text{otherwise,} \end{cases} \quad (\text{E.1})$$

with w defined in (3.2), which is independent of wealth.

Suppose now that the continuum of miners have heterogeneous costs. In Section 2.4 miners were indexed by their wealth x , whereas here the miners are characterized by their wealth x and the parameter c indexing their cost. Because of the wealth-independence of the optimal hash rates described above, we may consider only c when indexing the miners. For simplicity, consider costs c in some bounded interval $(c_0, c_1) \subset \mathbb{R}_+$. Denote by $f_c(c)$ the associated density of miners. Then, the individual miners hash at rates (E.1). We use the notation $\alpha^*(c)$ to stress the dependence on c .

We see from (E.1) that a miner is active if its cost

$$c < \frac{w(r;\gamma)\mathcal{K}}{\bar{\alpha}M}.$$

Define $c_{\max} = \min\{\frac{w(r;\gamma)\mathcal{K}}{\bar{\alpha}M}, c_1\}$. By this definition and (E.1), miners with costs $c < c_{\max}$ are active, while the rest do not mine. Then, averaging (E.1) over active miners, i.e., over $c < c_{\max}$ with respect to the density $f_c(c)$, leads to the following equation for $\bar{\alpha}$:

$$\bar{\alpha} = \int_{c < c_{\max}} \left(-M\bar{\alpha} + \sqrt{\frac{w(r;\gamma)\mathcal{K}M\bar{\alpha}}{c}} \right) f_c(c) dc. \quad (\text{E.2})$$

In a case where all miners are active, i.e., $c_{\max} = c_1$, we can solve (E.2) to find

$$\bar{\alpha} = \frac{w(r;\gamma)\mathcal{K}M}{\bar{c}(M+1)^2}.$$

where \bar{c} is the inverse square root averaged cost

$$\frac{1}{\sqrt{\bar{c}}} = \int_{c_0}^{c_1} \frac{1}{\sqrt{c}} f_c(c) dc.$$

Again using (E.1), the condition for a miner to be active is therefore

$$c < \left(\frac{M+1}{M} \right)^2 \bar{c},$$

and so costs have to be increasingly homogeneous for larger M for everyone to be active. The cost distribution may have a small tail to the left, but on the right miners would not participate. This suggests that homogeneous costs, possibly with few cost-advantaged miners, is the right place to start, as we did in this paper.

References

- Abadi, J., & Brunnermeier, M. (2018). *Blockchain economics* (Tech. Rep.). National Bureau of Economic Research.
- Alsabah, H., & Capponi, A. (2020). *Pitfalls of bitcoin’s proof-of-work: R&D arms race and mining centralization* (Tech. Rep.).
- Arnosti, N., & Weinberg, S. M. (2018). Bitcoin: A natural oligopoly. In *Itcs*.
- Arrow, K. J., & 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., & Vicsek, T. (2002, August). Evolution of the social network of scientific collaborations. *Physica A: Statistical Mechanics and its Applications*, 311(3-4), 590–614.
- Bertucci, C., Bertucci, L., Lasry, J.-M., & Lions, P.-L. (2020). Mean field game approach to bitcoin mining. *arXiv preprint arXiv:2004.08167*.
- Biais, B., Bisière, C., Bouvard, M., & Casamatta, C. (2019, 04). The blockchain folk theorem. *The Review of Financial Studies*, 32(5), 1662–1715.
- Brémaud, P. (1981). *Point processes and queues: martingale dynamics* (Vol. 50). Springer.
- Brown-Cohen, J., Narayanan, A., Psomas, A., & Weinberg, S. M. (2019). Formal barriers to longest-chain proof-of-stake protocols. In *Proceedings of the 2019 acm conference on economics and computation* (pp. 459–473).
- Chan, P., & Sircar, R. (2017). Fracking, renewables, and mean field games. *SIAM Review*, 59(3), 588–615.
- Cong, L. W., & He, Z. (2019, 04). Blockchain disruption and smart contracts. *The Review of Financial Studies*, 32(5), 1754–1797.
- Cong, L. W., He, Z., & Li, J. (2021). Decentralized mining in centralized pools. *The Review of Financial Studies*, 34(3), 1191–1235.
- Dai, M., Jiang, W., Kou, S., & Qin, C. (2019). From Hotelling to Nakamoto: The economic meaning of Bitcoin mining.
(In preparation)
- Deshmukh, S. D., & Pliska, S. R. (1980). Optimal consumption and exploration of nonrenewable resources under uncertainty. *Econometrica*, 48(1), 177–200.
- Easley, D., O’Hara, M., & Basu, S. (2019). From mining to markets: The evolution of Bitcoin transaction fees. *Journal of Financial Economics*, 134(1), 91–109.
- Fanti, G., Kogan, L., Oh, S., Ruan, K., Viswanath, P., & Wang, G. (2019). Compounding of wealth in proof-of-stake cryptocurrencies. In *International conference on financial cryptography and data security* (pp. 42–61).
- Gallego, G., & Hu, M. (2014). Dynamic pricing of perishable assets under competition. *Management Science*, 60(5), 1241–1259.
- Gallego, G., & van Ryzin, G. (1994, August). Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science*, 40(8), 999–1020.
- Gallego, G., & van Ryzin, G. (1997). A multiproduct dynamic pricing problem and its applications to network yield management. *Operations Research*, 45(1), 24–41.
- Garratt, R., & van Oordt, M. R. (2020). Why fixed costs matter for proof-of-work based cryptocurrencies. *Available at SSRN 3572400*.
- Guéant, O., Lasry, J.-M., & Lions, P.-L. (2011). Mean field games and applications. In

- Paris-princeton lectures on mathematical finance 2010* (pp. 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.
- Huang, M., Caines, P. E., & Malhamé, R. P. (2007). Large-population cost-coupled lqg problems with nonuniform agents: individual-mass behavior and decentralized ε -nash equilibria. *IEEE transactions on automatic control*, 52(9), 1560–1571.
- Huang, M., Malhamé, R. P., Caines, P. E., et al. (2006). Large population stochastic dynamic games: closed-loop mckean-vlasov systems and the nash certainty equivalence principle. *Communications in Information & Systems*, 6(3), 221–252.
- Kondor, D., Pósfai, M., Csabai, I., & Vattay, G. (2014). Do the rich get richer? An empirical analysis of the Bitcoin transaction network. *PLoS ONE*, 9(2), e86197.
- Lasry, J.-M., & Lions, P.-L. (2007). Mean field games. *Japanese Journal of Mathematics*, 2(1), 229–260.
- Li, H., Fan, Y., & Ying, L. (2021). A simple multiscale method for mean field games. *Journal of Computational Physics*, 439, 110385.
- Li, L. (1988). A stochastic theory of the firm. *Mathematics of Operations Research*, 13(3), 447–466.
- Ludkovski, M., & Sircar, R. (2012). Exploration and exhaustibility in dynamic Cournot games. *European Journal of Applied Mathematics*, 23(3), 343–372.
- Ludkovski, M., & Sircar, R. (2015). Game theoretic models for energy production. In R. S. R. Aid M. Ludkovski (Ed.), *Commodities, energy and environmental finance*. Springer.
- Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system. (<https://bitcoin.org/bitcoin.pdf>)
- Narayanan, A., Bonneau, J., Felten, E., Miller, A., & Goldfeder, S. (2016). *Bitcoin and cryptocurrency technologies: A comprehensive introduction*. Princeton, NJ, USA: Princeton University Press.
- Nourian, M., & Caines, P. E. (2013). ϵ -Nash mean field game theory for nonlinear stochastic dynamical systems with major and minor agents. *SIAM Journal on Control and Optimization*, 51(4), 3302–3331.
- Perc, M. (2012, July). Evolution of the most common English words and phrases over the centuries. *Journal of The Royal Society Interface*, 9(77), 3323–3328.
- Roşu, I., & Saleh, F. (2021). Evolution of shares in a proof-of-stake cryptocurrency. *Management Science*, 67(2), 661–672.
- Simon, H. A. (1955). On a class of skew distribution functions. *Biometrika*, 42(3-4), 425–440.
- Sockin, M., & 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.