

An Equilibrium Model of the Market for Bitcoin Mining Julien Prat, Benjamin Walter

▶ To cite this version:

Julien Prat, Benjamin Walter. An Equilibrium Model of the Market for Bitcoin Mining. Journal of Political Economy, 2021, 129 (8), pp.2415-2452. 10.1086/714445 . hal-03506522

HAL Id: hal-03506522 https://hal.science/hal-03506522

Submitted on 2 Jan 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

An Equilibrium Model of the Market for Bitcoin Mining

Julien PRAT^{*} Benjamin WALTER[†]

March 8, 2021

Abstract

We propose a model which uses the exchange rate of Bitcoin against the US dollar to predict the computing power of Bitcoin's network. We show that free entry places an upper-bound on mining revenues and explain how it can be identified. Calibrating the model's parameters allows us to accurately forecast the evolution of the network computing power over time. We find that a significant share of mining rewards was invested in mining equipment and that the seigniorage income of miners was limited.

*CNRS, CREST, Ecole Polytechnique. Contact details: CREST, 5 avenue Henry Le Chatelier, 91120 Palaiseau, France, julien.prat@ensae.fr.

[†]CREST, Université Paris-Saclay. Contact details: CREST, 5 avenue Henry Le Chatelier, 91120 Palaiseau, France, benjaminjwalter@gmail.com. We are grateful to Daniel Augot, Giuseppe Bertola, Bruno Biais, Christophe Bisière, Vincent Danos, Hanna Halaburda, Joshua Miller, Harald Uhlig as well as three anonymous referees for their insightful suggestions and feedback. We also thank seminar participants at the University of Paris Dauphine, University of Nottingham, BlockSem seminar, University of Basel, University of Zurich, Cerg-talks, University of Chicago, New-York University, European University Institute and CREST for their comments. We acknowledge the support of the Investissements d'Avenir grant (ANR-11-IDEX- 0003/Labex Ecodec/ANR-11-LABX-0047) and of the Academic Chair Blockchain and B2B Platform.

1 Introduction

Bitcoin is the first payment system that operates without a central authority. Its protocol replaces trusted third parties with a network of computers, commonly referred to as "miners", that guarantees the immutability of past transactions. Miners compete for the right to add blocks of new transactions to the public ledger. Winners are rewarded with freshly minted bitcoins. Hence, as the value of Bitcoin skyrocketed, so did the resources devoted to mining. What started as a hobby for a few miners using their personal computers, eventually blossomed into an industry which consumes around 0.5% of the world's electricity through its network of mining farms, each one of them operating thousands of machines specially designed for mining.¹

The runaway growth of Bitcoin's carbon footprint is widely perceived as one of the most convincing argument against its long-run sustainability. Since the amount of electricity allocated to mining is increasing in Bitcoin price, it seems that the currency cannot appreciate much further without causing intolerable environmental damage.² Assessing the robustness of this prediction requires a proper understanding of the factors that shape the relationship between Bitcoin price and miners' investment in computing power. For instance, if price stability discourages miners entry, Bitcoin could very well stabilize at a higher price than today without triggering an increase in energy consumption. To rule out this possibility, a dynamic model of the mining industry is needed.

We provide such a framework by devising a dynamic model that accurately captures the evolution of miners' computing power. Our model builds on the following two key observations. First, investment in mining hardware cannot easily be reversed: machines have little to no resale value because they become obsolete very quickly, and, from 2014 onward, have no use outside of the market for Bitcoin mining since they have been optimized for this task only. Second, miners face a lot of uncertainty about future revenues due to the tremendous volatility of Bitcoin price. This combination generates a range of inaction where expected revenues are too low to justify entry but sufficient to prevent incumbents from exiting the market.

The main challenge is that we cannot consider the problem of each miner in isolation or treat revenues as exogenous. Instead, we have to take into account how returns are endogenously determined by the number of active miners. Our framework takes the demand for bitcoins as given. Combining

 $^{^1\}mathrm{See},$ among other sources, digiconomist.net/bitcoin-energy-consumption .

 $^{^{2}}$ In one of the most pessimistic forecasts, Mora et al. (2018) argue that, if Bitcoin adoption continues unabated, it could push global warming above 2 Celsius degrees within the next three decades.

the exchange rate of Bitcoin against the US dollar with the total computing power of Bitcoin network, we construct a new measure for miners' payoffs. We show that miners enter the market only when this measure reaches a reflecting barrier. Payoffs never exceed this threshold because new entries push payoffs down by triggering additional increases in mining costs. The characterization of the equilibrium is complicated by the fact that mining hardware benefit from a high rate of embodied technological progress. We show how one can adapt the canonical model of Caballero and Pyndick (1996) to account for this trend, and prove that the entry threshold decays at the rate of technological progress.

We calibrate the model and find that it forecasts remarkably well how miners respond to changes in the price of Bitcoin. The accuracy of our simulations is a testament to the fact that miners operate in an environment which verifies many properties that are often assumed but rarely verified in practice.³ First, free entry holds because mining is an unregulated activity with a streamlined set of tasks. Anyone can buy the appropriate hardware online and join the mining race. Second, all miners, independently of their geographical location, face the same challenge and earn the same amount of bitcoins for solving it. Third, the mining technology exhibits returns to scale that are constant by nature because Bitcoin protocol ensures that the odds of finding new coins remains proportional to the size of one's computing power. Fourth, the elasticity of revenues with respect to the network computing power is commonly known since it is encoded in Bitcoin protocol, and is therefore observable by all parties. The conjunction of all these features is extremely rare, if not unique. It makes the market for Bitcoin mining a perfect laboratory for models of industry dynamics, especially since all transactions are public, giving anyone access to perfectly clean and exhaustive data.

Although our baseline model is fairly accurate in the medium to long run, it sometimes temporarily deviates from the data. To identify the origins of these deviations, we devise and calibrate a series of extensions. First, we allow for discontinuities in miners' rewards that take into account the reductions in the rate of Bitcoin creation which are triggered by the protocol every four years. Second, we introduce a time-to-build and show that it explains the sluggish response of miners to sudden surges

in Bitcoin price.

³The market for Bitcoin mining is not affected by the kind of regulations, local oligopolies and search frictions that impede trade in most industries. See, for instance, Collard-Wexler (2012) for an analysis of local oligopolies in the ready-mix concrete industry and how they respond to demand fluctuations; or Buchholz (2017) for a characterization of the dynamic spatial equilibrium of taxicabs in New-York city.

Finally, we introduce congestion effects and non-linearities in the adjustment cost function at the industry level. This extension allow us to match the data during Bitcoin's 2017 bubble. Comparing our calibrated costs to available information about the price of mining hardware, we find that they are consistent. Our model indicates that the bubble triggered an increase in the demand for mining hardware which was so massive that it stretched the manufacturing capacity of hardware producers, resulting in a spectacular jump in the price of hardware that prevented entrants from flooding the market.

Having established the accuracy and robustness of our framework, we use it to identify which actors, if any, have been able to extract significant rents. In contrast to platforms that charge monopoly fees for their intermediation, Bitcoin ensures that all rewards are passed on to miners.⁴ We use our estimates to infer the profits, or seigniorage revenues, of miners as a function of the number of mining firms. Our results support the notion that Bitcoin harvested market forces and price signals to create a quasi-competitive environment. This implies in turn that miners channeled most of their rewards towards the producers of their input factors, namely hardware manufacturers and electricity suppliers. The distribution of rents therefore followed a similar pattern as that from the Californian gold rush during which, according to Clay and Jones (2008), most of the profits were reaped by individuals who pursued other occupations than mining.

The rents of hardware manufacturers have been eroded by the entry of new competitors in 2018. We use our model to quantify the long-run impact of this trend, showing that it increases the electricity consumption of the network as more agents find it profitable to enter the mining race. Our model also indicates that, among the other factors that could worsen the carbon footprint of Bitcoin, a slowdown in the rate of technological progress of the mining technology would be particularly detrimental.

Related literature. — Our paper uses insights from the literature on irreversible investment to contribute to the nascent field of *cryptoeconomics*. Bitcoin was created a decade ago when Nakamoto's paper ((Nakamoto, 2008)) was made public on October 31st 2008. It did not immediately attract much attention and it took a few years for Bitcoin to become the focus of academic research. Early works analyzed the reliability of Bitcoin network (Karame et al., 2012; Decker and Wattenhoffer, 2013). Reid and Harrigan (2012) examined the anonymity of users, which enabled Athey et al. (2017) to quantify the different ways bitcoins are used and Foley et al. (2018) to precisely identify

⁴See Huberman et al. (2017) for a description of Bitcoin protocol as a "monopoly without a monopolist".

illegal Bitcoin users. Grunspan and Pérez-Marco (2017) and Bowden et al. (2018) both corrected mathematical approximations made by Nakamoto in his seminal paper.

It is only recently that papers studying the economic implications of cryptocurrencies have started to emerge. Most articles focus on the monetary implications of Bitcoin. Schilling and Uhlig (2018), Biais et al. (2018) and Hong et al. (2017) study the interactions between fiat money and Bitcoin, providing formulas for the fundamental value of Bitcoin and testing their implications. Observing the plethora of existing cryptocurrencies, Fernández-Villaverde and Sanches (2016) characterize the conditions under which currency competition is economically viable and efficient. Gandal et al. (2017) analyze exchange rate manipulations, while Chiu and Koeppl (2017) assess the calibration of the parameters that underlie Bitcoin's design. Cong and He (2018) question the disclosure of information which results from the use of public blockchains.

A series of recent papers is more closely related to our research since they study the market for mining. Rosenfeld (2011), Houy (2016) and Biais et al. (2017) investigate miners' incentives to behave cooperatively, as expected in Bitcoin protocol, or to play "selfish". Ma et al. (2018) model the market for mining as a game between miners. Cong et al. (2018) study the rise of mining pools which allow miners to share their computing power in exchange of a fair allocation of the mining rewards. Although Cong et al. (2018) find that mining pools do not necessarily undermine the decentralization of Bitcoin's network, they stress that risk sharing significantly escalates the arms race between miners. We bypass risk diversification by considering that miners are risk neutral, a specification which can be rationalized as describing the current state of the industry where pooled mining has become the norm. Alsabah and Capponi (2019) identify another force that intensifies the arms race: When R&D is endogenous, higher investments in research translate into a more aggressive mining game. Even though we take the evolution of the mining technology as given, we characterize in Section 5.1 the impact that market concentration has on the computing power deployed by miners. Arnosti and Weinberg (2018) show that small cost asymmetries among miners can result in highly concentrated ownership of mining equipment, whereas Bertucci et al. (2020) use mean field game theory to characterize the strategic interactions between miners. Finally, Huberman et al. (2017) and Easley et al. (2019) analyze how users solve a queuing problem to select their optimal fees. They raise concerns about the sustainability of Bitcoin in the long-run, when miners will be rewarded in transaction fees only.

Our paper also models the market for mining but, unlike the aforementioned articles, we focus on miners' entry decisions. We show that their behavior can be captured using real options theory. Our findings should therefore be of interest to the broad community of researchers working on industry dynamics. Since it would be impossible to cover all the major contributions to this field, we refer to Dixit and Pyndick (1994) for a comprehensive overview, as well as to Thomas (2002), Caplin and Leahy (2010) and Bachmann et al. (2013) for more recent surveys of the related literature. Our model being devised in an equilibrium setting, it builds on the work of Bertola and Caballero (1994) and Caballero and Pyndick (1996). We find that, despite its apparent novelty, the market for Bitcoin mining behaves very much like the canonical industries described in these two seminal papers. In particular, we show how one can observe the entry threshold using industry-level data only. As far as we are aware, no other industry has yet been used to construct such a direct measure and to document the accuracy of the entry rule with real-time data.

Structure of the paper. — The article is organized as follows. Section 2 lays out the baseline model along with two extensions. Section 3 presents the data and explains how we calibrate the models described in the previous section. The 2017 bubble is analyzed in Section 4. The implications of our findings are discussed in Section 5, while Section 6 concludes. The proof of the main Proposition is relegated to the Appendix.

2 Equilibrium

We propose a framework that takes the demand for bitcoins as given. We use the trajectory of Bitcoin exchange rate against the US dollar, i.e. the dollar nominal price of Bitcoin, to predict the computing power of the network. Explaining how Bitcoin achieves decentralization is beyond the scope of this paper. Hence we only cover the elements that are required for the understanding of our model, namely the tasks accomplished by miners and the rewards they get in return. We refer readers interested in a more comprehensive treatment to Nakamoto (2008) and Narayanan et al. (2016).

2.1 Baseline Model

The mining technology.— The main challenge for a decentralized currency is to maintain consensus among all participants in order to prevent double spendings. To avoid such conflicts, transactions are bundled together into blocks which are incrementally appended to the public ledger. The resulting data structure is called a Blockchain because blocks of transactions are chained according to their dates of creation. Bitcoin miners continuously compete for the right to add the next block. To win the competition, miners have to stamp their block with a "proof-of-work". Generating a valid proof-of-work boils down to generating a block whose header is such that $S(header) \leq \underline{s}$, where \underline{s} is an arbitrary threshold and S is a hash function.⁵ Hash functions are such that the only way to find a valid header is to randomly hash guesses until the above condition is met.⁶

The value of the threshold \underline{s} determines the difficulty of finding a valid block, which we denote by D. The target \underline{s} can be adjusted so that every computed hash will lead to a valid block with probability 1/D. The probability that a hash yields a valid block is, for all practical purposes, independent of the number of trials already done. This memory-less property implies that the event of mining a block is captured by a Poisson process.⁷ The Poisson arrival rate $\lambda(h, D) = h/D$ is linear in the number, h, of computed hashes per period or hashrate, a requirement known as fairness in the computer science literature.⁸

We normalize the length of a period to 10 minutes and set the hashrate of each miner equal to one hash per period. Let Q_t denote the *total hashrate* of the network, i.e. the overall number of mining units currently competing for the next block. Since fairness implies that the mining technology exhibits constant returns to scale, the average waiting time between blocks is equal to $\lambda(Q_t, D_t)^{-1} = D_t/Q_t$. To ensure that block generation proceeds at a steady pace, the difficulty of the hashing problem is adjusted by the protocol until new blocks are created on average every ten minutes. Given our normalization, this objective is attained when $\lambda(Q_t, D_t) = 1$, so that $D_t = Q_t$.

In practice, the hashrate of the network Q_t is not directly observable. Bitcoin circumvents this problem by relying on an adaptive expectation algorithm. Every two weeks on average, Bitcoin protocol uses the block generation rate over the last 2016 blocks to infer the average value of Qduring the previous period. Then the difficulty parameter is adjusted until it matches the estimated

⁵More precisely, each block possesses a header that includes the hash of the previous block, the root hash of the merkle tree of all transactions in the block, the current time, the target \underline{s} and a nonce. The nonce is a number that can be arbitrarily chosen by miners so as to meet the target. For further details on the mining process, see Narayanan et al. (2016).

⁶This property is often referred to as puzzle-friendliness in the cryptographic literature.

⁷See the online Appendix A and Rosenfeld (2011) for a derivation.

⁸Our model builds on the premise that Bitcoin works as intended by its protocol. Whether this is actually true is the subject of active research among cryptographers. For instance, it has been shown by Eyal and Sirer (2014) that selfish mining allows miners to mine more than their "fair share" of blocks.

value of Q. This updating procedure guarantees that, if the network hashrate does not deviate too much from its estimated average, the block generation rate will remain close to its target of 10 minutes. Since we devise our model in continuous time, assuming that difficulty is adjusted periodically would greatly complicate the analysis, making it impossible to derive tractable results. This is why we slightly idealize the actual protocol by assuming that D_t is continuously adjusted.⁹ We check in the online Appendix E that the number of blocks mined every day mostly remained within the confidence interval centered on the protocol's target. In other words, the block generation rate was not significantly different from one block every 10 minutes, and so did not deviate much from the idealized state that would prevail under Assumption 1.

Assumption 1. The difficulty parameter D_t is continuously updated and set equal to the current network hashrate Q_t , so that $D_t = Q_t$ for all t.

Miners' revenues. — Building a valid block is costly in terms of hardware and electricity. Miners are compensated when they win the competition: The first miner who finds a valid block earns a predetermined amount of new coins (12.5 bitcoins at the time of writing), and the sum of the fees granted by the transactions included in the block. Whereas the amount of new bitcoins is fixed by the protocol, fees are freely chosen by users. So far transaction fees have accounted for only 2.1% of average block rewards. We use R_t to denote the block rewards in dollars, i.e. the $\beta/\$$ exchange rate multiplied by the sum of new coins and fees. Then, as shown in the online Appendix A, the flow payoff P_t of a miner is equal to the block rewards, R_t , times the Poisson arrival rate, $1/D_t$, of a valid block¹⁰

$$P_t = R_t / D_t. \tag{1}$$

Under Assumption 1, the flow payoff is given by an isoelastic function of the network hashrate

$$P_t = R_t / Q_t. \tag{2}$$

⁹From a formal standpoint, this hypothesis is equivalent to assuming that Q_t is observable and that D is set equal to Q with a delay ε , so that $D_{t+\varepsilon} = Q_t$. Our model arises in the limit as ε converges to zero.

¹⁰We implicitly assume that miners value the block rewards at the current market price of Bitcoin. This premise is consistent with Schilling and Uhlig (2018) since they show that agents should be indifferent between holding bitcoins or fiat money. Actually, miners are more likely than other agents to prefer fiat money because they have already tied up the value of their investment to that of Bitcoin. Hence, converting their Bitcoin rewards into fiat money allows them to hedge part of their investment risk. In practice, miners have to wait on average 16 hours and 40 minutes, i.e. 100 blocks, before being able to transfer their newly earned coins.

The microfoundation of (2) is rather unique since the decreasing relationship between revenues and industry capacity does not stem from the satiation of consumers' demand, but is instead generated by the increase in mining costs encoded in Bitcoin protocol.

We do not attempt to endogenize the demand for bitcoins, and thus take its exchange rate against the dollar as given. Following much of the literature on irreversible investment, we assume that revenues follow a Geometric Brownian Motion (GBM hereafter).

Assumption 2. The block rewards $(R_t)_{t\geq 0}$ follow a Geometric Brownian Motion, so there is an $\alpha \in \mathbb{R}$, and a $\sigma \in \mathbb{R}_+$, such that

$$dR_t = R_t \left(\alpha dt + \sigma dZ_t\right),\tag{3}$$

where $(Z_t)_{t\geq 0}$ is a standard Brownian motion.

Given that newly minted coins account for the bulk of block rewards, changes in R are almost fully proportional to changes in Bitcoin price. The GBM specification is consistent with the equilibrium pricing formula for bitcoins of Schilling and Uhlig (2018). The stochastic term captures the martingale component arising from the "exchange rate indeterminacy result" of Kareken and Wallace (1981); while the deterministic trend is proportional to the correlation between the pricing kernel and the price of Bitcoin. As we will see below, the estimated α are always positive, which indicates that the pricing kernel and Bitcoin price were negatively correlated. Additional factors were probably at work during our period of study because Schilling and Uhlig (2018) focus on the equilibrium situation where Bitcoin is used as a medium of exchange. Whether or not this is true today remains open to debate, but most people would agree that Bitcoin was not widely used as a medium of exchange during its adoption phase. Then, as shown by Biais et al. (2018), the rate of return on Bitcoin for our geometric specification as transactional benefits are proportional to the price of Bitcoin, a property that distinguishes cryptocurrencies from other assets and which yields a pricing equation with a multiplicative structure.

The GBM specification narrows down the class of equilibrium price processes by requiring that they exhibit independently and normally distributed returns with constant variance. Assessing these requirements, we do find that returns are not linearly autocorrelated, and that their distribution is well approximated by a normal distribution (see online Appendix F). However, we also find that tail events are too common, and that the volatility of returns varies over time. These shortcomings of the GBM model are not specific to Bitcoin but common to most financial assets. Yet, GBM processes are still widely used to price assets because they provide a reasonable first-order approximation, while being much more convenient to handle than jump-diffusion processes. We adopt this pragmatic approach and leave the study of more complex price processes to further research.

Knowing the law-of-motion followed by the reward process is not sufficient to compute the expected payoffs because they also depend on the hashrate of the network Q, whose level is endogenously determined. To solve for the equilibrium, one has to simultaneously derive the process followed by Qand the entry policy of miners.

Market entry. — Mining is a costly activity. To operate a unit of hashpower bought at time τ ,¹¹ miners incur the flow operating cost C_{τ} . The electricity consumption of mining hardware accounts for a significant share of the operating costs.¹² The costs vary with the vintage of the hardware because they benefit from embodied technological progress, as newer machines are able to perform more hashes with the same amount of energy. We assume that investment in mining hardware is irreversible, and explain why this is a reasonable premise in Section 3.1.

Assumption 3. Market entry is irreversible. Once installed, mining hardware is never switched off.

Assumption 3 implies that the value of a unit of hashpower of vintage τ reads

$$V(P_t,\tau) = \mathbb{E}_t \left[\int_t^\infty e^{-r(s-t)} P_s ds \right] - \frac{C_\tau}{r},\tag{4}$$

where r is the yearly discount rate and $t \ge \tau$.¹³ We consider that all the miners of a given vintage face the same problem. In practice, operating costs may differ across locations but only those miners that have access to the cheapest sources of electricity will find it profitable to enter the market.

Entrants have to buy a unit of hardware whose price we denote by I_t . Both investment and operating costs decrease over time because hardware becomes more efficient. Let A_t measure technological

¹¹The hashpower measures the number of hashes that can be performed per period.

¹²We implicitly assume that the price of electricity remains constant. It is easy to relax this restriction by letting C also depend on the current date t. However, changes in electricity costs can be ignored in the empirical analysis because they are dwarfed by variations in Bitcoin price.

 $^{^{13}}$ See the online Appendix A for a derivation of (4). It is straightforward to generalize (4) to include an exogenous rate at which hardware breaks down. We do not take it into account because its calibration returns non significant values. Intuitively, failures seem to occur at a much slower rate than technological obsolescence since we do not observe that the network hashrate decreases in the absence of market entry.

efficiency, so that buying A_t units of hashpower at date t costs the same amount than buying one unit of hashpower at date 0. For the reasons explained below, we assume that technological improvements accrue at a constant pace, i.e. $A_t = \exp(at)$ with a > 0.

Assumption 4. Machines become more efficient at the constant rate a > 0. Hence the investment and the operating costs satisfy $I_t = I_0/A_t = \exp(-at)I_0$ and $C_t = C_0/A_t = \exp(-at)C_0$.

Anyone can enter the mining race: All that is required is to buy the mining hardware and to connect it to a steady supply of electricity. Hence free entry is likely to prevail, ensuring that

$$I_t \ge V(P_t, t) = \mathbb{E}_t \left[\int_t^\infty e^{-r(s-t)} P_s ds \right] - \frac{C_t}{r}, \text{ for all t.}$$
(5)

When new miners enter the market, (5) holds with equality. Since the exchange rate follows a Markov process, it is natural to conjecture that miners' decisions will only depend on the current realization of P: Whenever payoffs reach some endogenously determined threshold \overline{P}_t , a wave of market entries ensures that the free entry condition (5) is satisfied.

To see why such a mechanism defines a competitive equilibrium, it is helpful to decompose the law of motion of P. Reinserting (3) into (2) and using Ito's lemma, we find that

$$d\log(P_t) = \left(\alpha - \frac{\sigma^2}{2}\right)dt + \sigma dZ_t - d\log(Q_t).$$
(6)

Payoffs are decreasing in Q because the response of the protocol to an increase in total hashrate is to raise the difficulty parameter, thus making it less likely for each miner to earn a reward. This is why free entry places an upper bound on payoffs. Their value can never exceed the threshold \overline{P}_t as more miners would find it profitable to enter the market, which would push payoffs further down.

Industry equilibrium. — So far, the main takeaway from our analysis is that the market for mining can be described as a standard industry because Bitcoin protocol generates a cost function that is increasing in aggregate capacity. We define a competitive equilibrium as a symmetric Nash equilibrium in entry strategies. If all other miners follow a policy of entry at \overline{P}_t , no individual miner can find it optimal to follow any other policy.

Definition 1 (Industry equilibrium).

An industry equilibrium is a payoff process P_t and an upper threshold \overline{P}_t such that:





(i) $P_t \in [0, \overline{P}_t].$

(ii) The network hashrate Q_t increases only when $P_t = \overline{P}_t$.

(iii) The free entry condition (5) is satisfied at all points in time, and it holds with equality whenever $P_t = \overline{P}_t$.

Conditions (i) and (ii) ensure that entry keeps P_t below the entry threshold \overline{P}_t ; while condition (iii) ensures that no individual miner will find it profitable to deviate from the entry policy. Conjecturing the properties of the equilibrium greatly simplifies the analysis since we only have to verify that they are indeed satisfied by the entry strategies. From a formal standpoint, the fundamental difference between our equilibrium and the one studied by Caballero and Pyndick (1996) is that, due to embodied technological progress, the investment and the operating costs decrease over time. Hence the entry threshold \overline{P}_t cannot remain constant. However, if we impose Assumption 4, so that mining efficiency improves at a constant rate, we can solve for the equilibrium in the space of detrended payoffs and recover a flat threshold.

Proposition 1. Assume that assumptions 1, 2, 3 and 4 hold. Then there exists an industry equilibrium (P_t, \overline{P}_t) such that P_t is a GBM reflected at $\overline{P}_t = \overline{P}_0/A_t$ where¹⁴

$$\overline{P}_{0} = \frac{(r-\alpha)\beta}{\beta-1} \left[I_{0} + \frac{C_{0}}{r} \right], \text{ and } \beta = \frac{\frac{\sigma^{2}}{2} - \alpha - a + \sqrt{\left(\alpha + a - \frac{\sigma^{2}}{2}\right)^{2} + 2\sigma^{2}\left(a + r\right)}}{\sigma^{2}} > 0.$$

$$(7)$$
¹⁴Note that, when $\alpha = r, \overline{P}_{0} = \left(I_{0} + \frac{C_{0}}{r} \right) \left(\alpha + a + \frac{\sigma^{2}}{2} \right).$

A typical equilibrium is illustrated in Figure 1. The upper-panel reports an arbitrary sample path for the payoff process $(P_t)_{t\geq 0}$. Payoffs follow the changes in block rewards and thus behave as a GBM until they hit the reflecting threshold \overline{P}_t . Such events trigger market entry, as shown in the lower-panel. The resulting increase in hashrate raises the difficulty of the mining problem and lowers payoffs until market entry is not anymore profitable. The entry threshold decreases at the rate of technological progress because it corresponds to the pace at which both investment and operating costs fall over time.

Comparative statics. — To get some intuition about the impact that each parameter has on the entry threshold, it is useful to consider the hypothetical situation where further entries are precluded. Then the marginal miner is also the last one to ever enter the market. Provided that $r > \alpha$, the value of the last entrant is positive whenever $P_0 > \overline{P}_0^{last} \equiv (r - \alpha) [I_0 + C_0/r]$.¹⁵ Comparing the thresholds with and without entry, we see that $\overline{P}_0 = [\beta/(\beta - 1)]\overline{P}_0^{last} > \overline{P}_0^{last}$. The break-even payoff is higher under free entry because the arrival of new miners ensures that future payoffs are reflected downwards when they reach the entry threshold \overline{P}_0 . The term $\beta/(\beta-1)$ measures the negative impact that entries have on the value of the marginal incumbent.

Differentiating the expression of \overline{P}_0 in (7), we find that $\partial \overline{P}_0/\partial a > 0$ and $\partial \overline{P}_0/\partial r > 0$. If technological progress accelerates, miners' revenues shrink more rapidly because there will be even more entries in the future. Hence entrants have to earn more early on, which raises the entry threshold. A similar mechanism explains the impact of r since future revenues are discounted at a higher rate when r goes up. Not surprisingly, an increase in the average growth rate α of the block rewards incentivizes entry as $\partial \overline{P}_0/\partial \alpha < 0$. Finally, the volatility of payoffs σ discourages entry since $\partial \overline{P}_0/\partial \sigma > 0$. Note that this is not due to an increase in the value of waiting because the perfectly competitive structure of the industry rules out such an option: Competitors would preempt any procrastination beyond the zero expected profit threshold. Instead, the negative impact of σ on entry is mechanical. Given that payoffs are truncated from above by the reflecting barrier, an increase in their spread automatically lowers their expected value. Quantitatively, the rate of technological progress a has, by far, the largest effect on \overline{P}_0 .

$$V_0^{last} = \int_0^\infty e^{-rt} \mathbb{E}_0[P_t] dt - \left[I_0 + \frac{C_0}{r}\right] = \frac{P_0}{r - \alpha} - \left[I_0 + \frac{C_0}{r}\right].$$

 $^{^{15}\}mathrm{If}$ market entry is forbidden, P_t obeys the same law of motion as R_t so that

2.2 Extensions

We now generalize our model so as to take into account the delivery lags for mining hardware and the halving of block rewards every four years.

Time-to-build. — We have assumed that miners can enter the market immediately. In practice, however, new hardware have to be delivered and installed. Each step increases the lapse of time separating entry from actual production. When it requires δ years to effectively become operational, prospective entrants at date t have to forecast their revenues starting from $t + \delta$. Hence they have to take into account the price fluctuations that will occur during the delivery period, as well as the amount and arrival times of hardware in the delivery pipeline.

To reduce the dimensionality of the state space, we follow the approach proposed by Grenadier (2000). Let H_t denote the amount of "committed hashrate", that is all the mining units which are either already operational or on their way to being delivered. Given that all orders will be installed δ years from now, the hashrate of the network when today's orders become operational will be equal to the current amount of committed hashrate, i.e. $Q_{t+\delta} = H_t$. Hence the relevant state variable from the standpoint of entrants is not anymore $P_t = R_t/Q_t$, but instead $P_t^{\delta} \equiv R_t/H_t$. We show in the online Appendix L that equilibrium strategies are functions of P_t^{δ} only, and that the mining market is in equilibrium when P_t^{δ} is a reflected GBM.

Proposition 2. Assume that assumptions 1, 2, 3 and 4 hold. Furthermore, assume that market entry is delayed by the time-to-build δ . Then there exists an industry equilibrium $\left(P_t^{\delta}, \overline{P}_t^{\delta}\right)$ such that $P_t^{\delta} = R_t/H_t$ is a GBM reflected at $\overline{P}_t^{\delta} = \overline{P}_0^{\delta}/A_t$. The entry threshold is related to that of the model without time-to-build by the following equation

$$\overline{P}_{t}^{\delta} = e^{(r-\alpha)\delta}\overline{P}_{t}\left[K_{t}/K_{t}^{\delta}\right],\tag{8}$$

where $K_t \equiv I_t + C_t/r$ denotes the overall costs of entry in the model without delay, and $K_t^{\delta} \equiv K_t - (1 - e^{-r\delta})C_t/r$.

Proof. See the online Appendix L.

The expression of the entry threshold with time-to-build differs from that of the baseline model in two respects. First, the overall costs of entry K_t^{δ} are slightly lower because they are evaluated at the time of the entry decision. Since entrants have to wait δ years to start mining, their operating costs C_t/r are multiplied by the discount factor $e^{-r\delta}$. Second, the threshold without delay is rescaled by $e^{(r-\alpha)\delta}$ because it is optimal to enter when the expected value of payoffs in δ years is equal to the discounted threshold $e^{r\delta}\overline{P}_t$. Since $\mathbb{E}_t[P_{t+\delta}] = e^{\alpha\delta}P_t^{\delta}$, setting $\mathbb{E}_t[P_{t+\delta}] = e^{r\delta}\overline{P}_t$ indeed implies that $P_t^{\delta} = e^{(r-\alpha)\delta}\overline{P}_t$.¹⁶

The models with and without time-to-build are not as similar as their descriptions might suggest. The solution of the baseline model is Markovian since knowing the current hashrate and Bitcoin price is enough to forecast the evolution of the network hashrate. By contrast, the solution of the model with time-to-build is path dependent since forecasts over the next δ years are conditional on all the purchase orders that were placed over the previous δ years.

Halvings.— Another limitation of our baseline model is that it ignores the inclusion in Bitcoin protocol of a feature which divides by two the number of coins issued per block. These so-called halvings are triggered every 210,000 blocks to ensure that the supply of bitcoins converges to a finite limit, namely 21 millions. Halvings generate discontinuities in the paths of R_t that are inconsistent with the GBM specification. To take them into account, one has to replace Assumption 2 with Assumption 5 according to which block rewards are halved every four years.

Assumption 5. The block rewards are given by $R_t = h_t \tilde{R}_t$. \tilde{R}_t follows a GBM while $h_t = \left(\frac{1}{2}\right)^{\lfloor t/4 \rfloor}$, where t measures the number of years elapsed since the inception of Bitcoin, and $\lfloor x \rfloor = \max_{n \in \mathbb{N}} \{n \leq x\}$.

Assumption 5 slightly simplifies the halving process. First, the reward a miner gets when she finds a block is not exactly divided by two after each halving because it includes transaction fees on top of new coins. But the discrepancy is not very important in practice, as transaction fees accounted for a residual share of block rewards in our samples. Second, halvings do not occur every four years, but instead every 210,000 blocks. Counting blocks is a way to approximate elapsed time because Bitcoin protocol adjusts the difficulty of the hashing problem every two weeks on average. It is shown in the online Appendix E that, as expected, the updating rule managed to keep the generation rate close to one block every 10 minutes.

$$\mathbb{E}_t\left[P_{t+\delta}\right] = \mathbb{E}_t\left[\frac{R_{t+\delta}}{Q_{t+\delta}}\right] = \frac{\mathbb{E}_t\left[R_{t+\delta}\right]}{H_t} = e^{\alpha\delta}P_t^{\delta}.$$

¹⁶The relationship between the expectation of $P_{t+\delta}$ and P_t^{δ} holds true because $Q_{t+\delta} = H_t$, so that





Note: R_t is computed using information collected on coindesk.com and btc.com. Q_t is measured in Terahash per second. Its value is inferred using the procedure described in Appendix G.

Halvings render the optimization problem of miners non-stationary: the closer they are to the halving date, the lower their expected payoffs. This implies that we have to rely on numerical methods because the entry threshold does not anymore admit a closed-form solution. Starting from the expression in Proposition 1, we proceed by backward induction and use a finite-difference procedure to approximate the entry rule. As the horizon increases, our algorithm quickly converges towards an entry threshold that is independent of the number of future halvings.¹⁷

3 Calibration

3.1 Data

We now show that feeding our model with exchange rate data allows one to accurately predict the evolution of the network hashrate. For this purpose, we need to infer the miners' payoffs $P_t = R_t/Q_t$. Remember that the numerator, R_t , is equal to the value of new coins plus the transaction fees. The number of created coins per block is specified by the protocol while Bitcoin exchange rate against the dollar is directly available from coindesk.com.¹⁸ The transaction fees are recorded in Bitcoin's

¹⁷More precisely, the entry threshold turns out to be stable after four iterations. We use finite-difference methods to approximate the Hamilton-Jacobi-Bellman equations satisfied by the value function of miners. We rely on the implicit Euler scheme in order to ensure that the approximation is stable. The system of linearized equations is solved using a generalization of the Gauss-Seidel iterative method known as the successive-over-relaxation method.

¹⁸There are many different exchanges and the exchange rate varies a bit across them. In the online Appendix D, we check the validity of Coindesk data by comparing them to a weighted average measure over 17 exchanges. Since the

blockchain and can easily be retrieved from btc.com. Thus all the components of R_t are readily available. By contrast, the network hashrate Q_t is not directly observable. It must be estimated using the theoretical probability of success and the number of blocks found each day. Given that we are not primarily interested in statistical inference, we relegate the description of our estimation procedure to the online Appendix G and save on notation by using Q_t to denote our estimate, although its time series only approximates the true hashrate. We show in the online Appendix G that the approximation is accurate. We update the value of Q_t on a daily basis and, since there are on average 144 blocks mined every day, the expected payoffs per day are given by $P_t = 144 \times R_t/Q_t$.

We report the series followed by R_t and Q_t in Figure 2. There is a clear correlation between the two variables. Our model suggests that their structural relation should become apparent if one takes the ratio of the two series and detrend it at the rate of technological progress a. Then the resulting series should behave as a reflected GBM. For many years, improvements in the semiconductor industry have followed Moore's law according to which processor speed doubles every two years. We expect improvements in the mining technology to outpace those in processing speed because miners came up with a series of innovations which allowed them to leverage their computing power. Thus, at this exploratory stage, we pick a rate of technological progress that is slightly faster than Moore's law (1.5 times faster). We will refine our guess later on by calibrating the value of a.

The detrended payoff series is reported in Figure 3. It exhibits two stationary regimes, with a break in the middle where payoffs decreased regularly until they reached a lower plateau. At first, this pattern does not seem to square with our model. But if we focus on the date at which the break initiates, we realize that it coincides with a fundamental change in the mining technology.

Early on, miners used to mine with their own computers. Around mid-2010, they realized that Graphical Processing Units (GPU) were much more efficient. One year later, miners started to use Field Programmable Gate Arrays (FGPA) and, since 2013, they mostly mine with Application Specific Integrated Circuits (ASIC). ASICs are also called mining rigs because their sole purpose is to solve Bitcoin's hash-puzzle.

The first ASIC was delivered to Mr. Jeff Garzik on January 30th 2013.¹⁹ Since this revolution in the mining technology boosted the rate of technological progress well above its long-run trend, its propagation among miners violates Assumption 4, and so, one should not expect the predictions of

two series are virtually identical, we select the one that is more easy to access.

¹⁹See https://bitcoin.stackexchange.com/questions/40944/when-did-the-asic-mining-era-begin





Note: P_t is equal to the daily network revenues $144 * R_t$ divided by Q_t . The vertical dashed lines indicate the period where miners transitioned from GPUs to ASICs.

our model to be verified during the transition phase. We therefore leave aside the lapse of time where miners switched from GPUs and FPGAs to ASICs, and focus instead on the subperiods where miners used the same technology. More precisely, during the first period, which ranges from 04/01/2011 to 01/31/2013,²⁰ miners mainly mined with GPUs; while they mostly relied on ASICs from 10/01/2014 onwards. We also exclude the winter of 2017 and the first semester of 2018 from our sample because they feature an episode of trading frenzy during which Bitcoin experienced a giant bubble followed by a sudden burst. We will analyze this event and its aftermath in Section 4.

Buying an ASIC is an irreversible decision because it can be used for cryptocurrency mining only. Hence, if the price of Bitcoin falls, ASICs cannot be resold for profit as all miners face the same returns. The irreversibility assumption is less obvious for GPUs. Yet, the calibrated values of *a* reported below in Table 1 show that GPUs were facing a very high rate of obsolescence. This suggests that irreversibility is also a sensible approximation for GPUs, as nobody would buy them second hand without a tremendous discount. The conjecture is confirmed by the analysis in the online Appendix J where we calibrate a model with reversible investment and find that it fails to match the data of the first subperiod.

 $^{^{20}{\}rm We}$ exclude the very early history of Bitcoin because it features an unstable block generation rate (see the online Appendix E.

3.2 Calibration strategy

We calibrate the parameters for each subperiod. The baseline models is parsimonious enough to rely on six parameters only: the deterministic trend α of rewards and their volatility σ^2 , the rate of technological progress a, the discount rate r, the price I_0 of one unit of hashpower bought at the beginning of the period, and the operating cost C_0 of that same unit. The first two parameters can be directly estimated using R_t only. Under Assumption 2, the log returns are independent and follow a normal distribution with mean $\mu \equiv \alpha - \sigma^2/2$, and variance σ^2 , which we estimate by maximum likelihood (see the online Appendix F).²¹

The rate of technological progress, a, and the reflecting barrier, \overline{P}_0 , are set to minimize a (pseudo)distance between the observed and the simulated paths of the hashrate. A direct consequence of our equilibrium definition is that $Q_t = \max\left(Q_{t-1}, \frac{R_t A_t}{P_0}\right)$ for all t. This condition provides us with a straightforward way to simulate the hashrate for any sample with T observations:

1. Set the initial value of the simulated hashrate Q_0^{sim} equal to its empirical counterpart, i.e. $Q_0^{sim} := Q_0.$

2. Update the simulated hashrate as follows $Q_t^{sim} := \max\left(Q_{t-1}^{sim}, \frac{R_t A_t}{P_0}\right)$, for $t = 1, \ldots, T$.

Since $(R_t)_{t\geq 0}$ and Q_0 are observed, the minimization procedure boils down to finding the value of aand \overline{P}_0 such that

$$(\hat{a}, \hat{\overline{P}}_0) \in \underset{(a, \overline{P}_0) \in \mathbb{R} \times \mathbb{R}_+}{\operatorname{argmin}} \sum_{t=1}^T \left(\frac{Q_t - Q_t^{sim}(a, \overline{P}_0)}{Q_t} \right)^2.$$
(9)

The other three parameters $\{r, I_0, C_0\}$ cannot be separately identified.²² Thus we fix r and recover the overall costs of entry at the beginning of each subperiod, $K_0 \equiv I_0 + C_0/r$, by equating the expression of \overline{P}_0 in (7) with the calibrated value $\hat{\overline{P}}_0$. Fortunately, the choice of discount rate turns out to be relatively neutral because the term $(r - \alpha)\beta/(\beta - 1)$ in (7), and thus the entry costs, are rather inelastic with respect to $r.^{23}$

²¹Estimating α and σ using adaptive instead of rational expectations does not significantly impact our estimates for the cost of market entry.

²²We estimate in the online Appendix M a model where investment is reversible and explain how it enables us to separate the investment cost, I, from the operating cost, C.

²³In the second period, setting r = 0.2 yields $K_0 =$ \$1639, while r = 0.05 yields $K_0 =$ \$1934.

3.3 Results

Calibrated parameters.— The parameters resulting from our calibration strategy are reported in Table 1, their values expressed as yearly rates whenever applicable.²⁴ The standard errors are obtained using block bootstrap, an estimation technique that is more suited to time series than standard bootstrap.²⁵

We first present the trend and volatility of the reward process. Both coefficients are independent of the modelling strategy since they are directly estimated by maximum likelihood on the rewards series R_t . The average growth rate of rewards, μ , decreased a lot between the two periods of study. As one would expect, early buyers of bitcoins earned higher returns. Information about their profits pushed the demand for bitcoins which raised the exchange rate even more. But these extremely high returns became harder to sustain as the market capitalization grew from a negligible amount to around \$ 200 billions by the end of our sample. In spite of this cooling process, investing in Bitcoin remained extremely profitable. These tremendous returns have led many observers to announce the imminent collapse of Bitcoin.²⁶ Whether or not such predictions will eventually be vindicated is beyond the scope of this paper, but our estimates for the volatility coefficient σ indicate that there was no obvious arbitrage opportunity; investors willing to bet on Bitcoin also had to bear a huge risk. Even though the volatility of rewards was divided by three in the second period, its value remained an order of magnitude higher than its counterpart for the S&P 500.²⁷

According to Moore's law, the price of one unit of hashpower should be divided by two every two years. Hence it implies that the rate of technological progress a should be close to $\log(2)/2 \approx 0.35$, a number well below the calibrated values of a reported in Table 1. The mining technology progressed at a faster pace than the one predicted by Moore's law because miners were able to implement innovations specific to the hash-puzzle on top of the raw increase in computing power. Our calibrated parameters also indicate that the rate of technological progress slowed down considerably in the second period, thus suggesting that improvements specific to the mining problem became harder to unearth as the technology matured.

Comparing the parameters across models, we see that introducing halvings lowers the overall costs of entry, K_0 , but raises the rate of technological progress, a. The decrease in K_0 is quite intuitive:

²⁴For example, the calibrated values of a means that the price of a new hardware has been on average divided by $\exp(a)$ every year during each subperiod.

²⁵The block bootstrap procedure is described in the online Appendix H.

²⁶According to the website bitcoinobituaries, by August 2019, 371 opinion pieces had predicted the death of Bitcoin. ²⁷We find that, for the S&P 500, $\sigma^2 = 0.053$ for the first period and $\sigma^2 = 0.027$ for the second period

Since halvings lower expected revenues, free entry holds when mining costs are smaller. The reason why a increases is more subtle. The adjustment corrects the misspecification of the baseline model that leads to an overestimation of the hashrate around the halving dates. This is why the minimization procedure, when applied to the baseline specification without halvings, generates a negative bias for a because it uses this parameter to reduce the discrepancies around the halving date.

The delivery lags of the model with time-to-build are relatively modest: 11.5 days during the first period, 46.5 and 43.5 days during the second and third periods. As expected, the lags are smaller during the first period since GPUs are more commonly available than ASICs. The total costs are lower with time-to-build than without, a finding that is in line with Proposition 2 and the impact of discounting on future profits. We also notice that the impact of delays on the calibrated rate of technological progress is negative in the first and third periods, but positive in the second period. Without further data, it is difficult to tell whether this ambiguity is structural or simply specific to our samples.

Finally, note that the parameters are identified with greater precision in the second and third periods. Three factors explain why the first period calibration is so fuzzy. First, Bitcoin price was extremely volatile. Second, Bitcoin experienced a long slump, a period known as the first crypto winter within Bitcoin community. This resulted in a nearly flat hashrate for most of the sample, as can be seen in Figure 4. From the standpoint of our calibration strategy, this means that there are relatively few data points where free entry binds, making it difficult to pinpoint the structural parameters. Finally, the technology was less homogenous during the first period. In particular, it witnessed the emergence of mining pools. Cong et al. (2018) explain why market entry was incentivized by this new opportunity to share risk, thus generating a positive bias in our calibration of the rate of progress. For all these reasons, we henceforth treat the second and third periods as our samples of reference.

Predicted vs actual hashrate. — The calibration procedure provides us with an estimate for the reflecting barrier, \overline{P}_0 , as well as for its trend, a. Using these two values, we can run the two-step algorithm described above to simulate the network hashrate Q^{sim} . We report the simulated series against their empirical counterparts in Figure 4. Since the model's fit in the second and third periods are quite comparable, we relegate the illustration of the third period to the online Appendix and focus on the first and second periods. In spite of its very parsimonious structure, the baseline model tracks the actual hashrate remarkably well.

Figure 4: Simulated vs Observed Hashrates First Period Second Period



We nonetheless notice some temporary discrepancies. The most striking is around the second halving date (07/09/2016). This should not be surprising since miners do not anticipate halvings in the baseline specification while they certainly do in reality. Figure 4 shows that this shortcoming is solved by the extended model with halvings. However, besides this specific period, the paths generated by the three models remain very close to each other. Due to the extreme volatility of the exchange rate, halvings affected miners' behavior only a couple of months ahead. It is actually more intriguing that such a disconnect between the simulation of the baseline model and the data is not apparent around the first halving date (11/28/2012). According to our model, miners had a very short investment horizon during the first period because the rate of technological progress was extremely high.

Another noticeable difference between the actual and simulated hashrates is that the former sometimes decreases, especially during the first period, while the latter never does. Our models cannot generate any decrease in hashrate because they are based on the premise that investment is totally irreversible.²⁸

These discrepancies do not invalidate our approach because its objective was to capture medium to long run adjustments in mining power, and it largely succeeds in that respect. Yet one may argue that such a conclusion is not warranted since our procedure would fit the data fairly well even if the model were misspecified because we minimize the distance between the simulation and the data. Remember, however, that the baseline model uses only two parameters to fit times series of 608 and

 $^{^{28}}$ We address this shortcoming in the online Appendix M by estimating an extension of our model where myopic miners are allowed to mothball and scrap their hardware.





913 data points. For each simulation, we start from the initial hashrate and let the model run until the end of the sample. Hence, any fundamental misspecification would generate a noticeable gap between the simulation and the data, at least over some time intervals. The fact that there is no deterioration of the models' accuracy is therefore a convincing argument in favor of their validity. We now provide additional evidence supporting this interpretation.

Our first validity check is to perform out-of-sample experiments. We assess the ability of the models to match out-of-sample data by dividing the second period into a fit period and a test period. We calibrate a and \overline{P}_0 on the fit period only and find that, even when the fit period is short, the calibrated values remain close to the ones based on the full sample. Hence, as shown in the online Appendix I, the predicted hashrate stays accurate several years after the end of the fit period.

Calibrated costs versus online prices. — The plausibility of our calibrated costs can be assessed by comparing them to online data on the selling price of mining rigs. To the best of our knowledge, there is no official source for hardware characteristics and availability dates. We therefore retrieved our data from different websites, selecting those which offered the most reliable information, namely Bitcoin wiki and the Bitcoin forum. We collected data on state-of-the-art mining hardware at the time it was put on the market.²⁹ We focus on the post-2014 period because there is too much uncertainty around the type of hardware that was used before the introduction of ASICs.

Figure 5 reports the electricity consumption and market price of hardware against the price series

 $^{^{29}\}mathrm{The}$ online data and their sources are reported in the online Appendix K.

consistent with our calibration. Remember that our model does not disentangle the flow costs from the fixed costs of entry.³⁰ Thus we equate the starting point of the simulated price series with its value in the data, so that the price of mining rigs accounts for 48% of the calibrated overall costs K = 3451\$. Extrapolating the model's prediction to cover all the sample where online data are available, we find that the calibrated rate of technological progress is almost identical to the one observed in the data. Figure 5 also validates Assumption 4 according to which the price of mining hardware and their energy consumption decrease at the same rate. There is, however, a specific time window where the assumption does not hold: Between August 2017 and February 2018, the price of mining rigs skyrocketed from around \$2,000 to \$5,200 before returning to its long-run trend. This temporary increase was triggered by the concurrent bubble in Bitcoin price. We explain in Section 4 how our model can be modified to capture this temporary deviation in the cost of market entry.

Inspecting the entry rule.— Besides assessing the accuracy of the simulated hashrate, we can also check whether the behavior of payoffs is consistent with the entry rule. Comparing the simulated with the observed payoffs series reported in Figure 6, we see that the model's accuracy is as convincing as for the hashrate.³¹ The upper-panel of Figure 6 focuses on the model with halvings whose threshold shifts down by 50% on the halving date. This drop is preceded by a period where the threshold slopes up because miners anticipate the fall in future revenues, and so, procrastinate further before entering the market. But the increase in the threshold becomes noticeable only a few months before the halving and is therefore not relevant for most of the preceding period. This might be surprising given that a division by two of revenues seems like a huge loss; yet one has to put it into perspective by comparing it to the very rapid obsolescence of hardware and to the extreme volatility of Bitcoin price. These two forces imply that a loss of 50% in rewards over a few months was not an implausible event.

As predicted by our model, payoffs remain below the threshold most of the time and tend to reflect downwards when they reach its vicinity. This is remarkable since \overline{P}_t was calibrated regardless of this requirement, fitting the hashrate only. Although the observed and simulated payoff series are often superimposed, they differ over some short time intervals. These discrepancies are usually triggered

³⁰We explain in the online Appendix M how one can use a model with reversible investment to separate the fixed costs from the flow costs. Its calibration enables us to pinpoint the starting value of the price series, yielding an estimate that is consistent with the data.

³¹For the sake of conciseness, we only report the payoff series in the second period. We show in the online Appendix B that the accuracy of the model is of comparable quality during the first and third periods.



Figure 6: Simulated vs. Observed Payoffs

by extreme increases in the exchange rate, as can be seen comparing the upper-panel of Figure 6 with the middle-panel that contains Bitcoin price series. Quite intuitively, when the exchange rate goes up by 10% or more in one day, miners cannot enter the market as quickly as the model predicts because they are facing, among many other frictions, delivery and manufacturing delays.

This conjecture is confirmed by the lower-panel of Figure 6 which reports the payoff series with time-to-build. The simulation now almost perfectly tracks the data. In particular, sudden price increases do not anymore drive a wedge between the model and the data. Instead, they push both series above the entry threshold for a short amount of time. This is possible in the model with time-to-build because the entry threshold acts as a reflecting barrier for the committed hashrate, H_t , and not for the actual hashrate Q_t . A sudden increase in the price of Bitcoin triggers a jump in H_t , as new miners decide to enter the market, but the impact of their decision is delayed by the time-to-build. This explains why the payoff series tends to revert after a big price surge: On impact, it follows the price trajectory, and decreases only a few weeks later, once the new equipment has been installed.

To take stock, out-of-sample experiments, data on the rate of technological progress of mining rigs and inspection of the entry rule all support the plausibility of our model. Its accuracy temporarily deteriorates around halving dates and after big price surges, two shortcomings that can be addressed by the introduction of halvings and time-to-build. These adjustments do not strongly affect the calibrated values of the parameters which remain rather stable across the three specifications. Having tested the soundness of our approach, we now consider an extension which enables us to fit the 2017 bubble and its aftermath.

4 The 2017 Bubble and its Aftermath

We have excluded the winter of 2017 and the subsequent months because Bitcoin experienced a period of trending frenzy. From three thousand dollars in September 2017, Bitcoin exchange rate shot up to nearly twenty thousand in December, and then, dropped back to six thousand in February 2018. This bubbly episode raises a significant challenge because it led to a structural break in the relation between the exchange rate and the network hashrate. As shown in Figure 7, if the relation had remained stable, the hashrate should have been five times higher than its actual value at the peak of the bubble. The discrepancy between the observed hashrate and the one that would have resulted from our frictionless model is explained by three different factors.

First, market entry was constrained by the manufacturing capacity of ASICs producers. In May 2017, there were approximately 230,000 active mining rigs. Between May and December 2017, the $\ddot{B}/\$$ exchange rate was multiplied by 12. To keep up with this pace, approximately 2,700,000 new mining rigs would have had to be installed within eight months only. Such a dramatic increase was bound to stretch the capacity of Bitmain, the main manufacturer of ASICs for Bitcoin mining. Second, Bitmain exercised his monopoly power and decided not to flood the market with new hardware in order to raise its selling price. Indeed, the price of an Antminer S9 mining rig was multiplied by three between the beginning and the climax of the bubble, and then divided by around four during the following crash.³² Third, as the bubble collapsed within a few months only, prospective miners simply cancelled their orders or backtracked on their decision to enter the market.

We take these constraints into account by assuming that investment costs are not constant but increasing in aggregate investment. More precisely, let q_t denote the *flow of entrants* at date t, so that $Q_t = Q_0 + \int_0^t q_s ds$. The investment costs for the marginal entrant are now given by

$$I\left(q_t; Q_t, A_t\right) = \frac{I_0}{A_t} \left[1 + \left(\frac{q_t}{bQ_t}\right)^{\eta}\right], \text{ for } I_0, b \in \mathbb{R}_+, \text{ and } \eta > 1.$$

$$(10)$$

Congestion externalities are captured by the convex function on the right-hand side of (10): An increase in the flow of entrants q_t stretches manufacturing capacities, thus raising the cost of entering the market. The parameter η controls the convexity of the cost function.³³ As η increases, I converges towards an hyperbolic function with an asymptote at bQ_t . Hence, one can think of bQ_t as the production capacity of ASICs manufacturers which is assumed to be proportional to the number of operational units.

For brevity, the derivation of the optimal entry rule under (10) is relegated to the online Appendix N. We demonstrate that, as in the baseline model, entry is a function of P_t only and that there exists a threshold \overline{P}_0 such that $dQ_t = 0$ whenever $P_t < \overline{P}_t = e^{-at}\overline{P}_0$. However, \overline{P}_t is not anymore a reflecting barrier. Due to the convexity of the cost function, aggregate investment is now absolutely continuous with respect to time.³⁴ Whenever $P_t > \overline{P}_t$, an Hamilton-Jacobi-Bellman equation pins

³²See the online Appendix K.

³³In particular, note that we recover our baseline specification when $\eta = 0$.

³⁴By contrast, aggregate investment was a singular control process in the baseline model with dQ_t being (infinitesimally) positive only on a measure-zero set of time points.



Figure 7: Model with Convex Costs

PARAMETERS: r = .1, a = 0.65, $\eta = 14.1$, b = 3. Note: Investment costs are multiplied by the efficiency coefficient A_t .

down the positive relation between q_t and P_t .

We include the most recent observations and select, as before, the parameters that minimize the distance between the simulated hashrate series and its empirical counterpart. We also include online price data for mining hardware in our set of targeted moments. This enables us to identify the convexity parameter η since it controls the elasticity of the hardware price with respect to aggregate investment.³⁵

Figure 7 shows that the baseline model vastly overestimates entry during the bubbly episode, and then, due to the irreversibility of past investment, remains well above the actual hashrate for the rest of the sample. By contrast, calibrating the model with convex costs enables us to match the hashrate over the full sample. The lower-panel which reports the normalized price series, $\tilde{I}_t = A_t I_t$, indicates that the simulated investment costs are also in line with the data. As predicted by the baseline model, the investment costs decrease at the rate of technological progress, thus generating a flat profile when they are normalized. There is, however, a notable exception during the height of the 2017 bubble

 $^{^{35}\}mathrm{See}$ the online Appendix N for further details on the calibration procedure.

where the entry costs increased dramatically. This means that the marginal cost function (10) is essentially flat until it nears the threshold bQ_t and starts to increase exponentially.

The calibrated parameters are reported in the legend of Figure 7. The value b = 3 implies that the congestion externality becomes relevant solely at very high rates of investment amounting to a twentyfold annual increase in the network hashrate. The calibration $\eta = 14.1$ confirms that the cost function is indeed extremely steep in the vicinity of bQ_t . This explains why the bubble triggered a sudden jump in the cost of entry that prevented miners from flooding the market during the bubbly episode. Interestingly, our calibration demonstrates that, even in the face of an event as extreme as Bitcoin bubble, one does not to need to abandon the efficient market hypothesis by assuming that miners refrained from investing because they anticipated the incoming crash. Instead, we find that their behaviour is explained by large variations in the price of their main input factor.

5 Discussion

Having calibrated our model and documented its accuracy, we now use its insights to address the two main questions that motivate our analysis. First, to which extent is the mining industry competitive and who has been able to appropriate the seigniorage income? Second, what forecasts can we draw about the evolution of Bitcoin's carbon footprint?

5.1 **Revenues allocation**

Oligopolistic industry.— Although we cannot reject the premise that miners operate under perfect competition, our results do not prove that the premise is true either. One should be careful when interpreting our findings because, as first established by Grenadier (2002), our entry rule holds even when the industry is oligopolistic. More precisely, assume that, instead of being populated by a continuum of atomistic miners, the market is controlled by n symmetric firms. Then, provided that Assumptions 1, 2, 3 and 4 hold, there exists a symmetric Nash equilibrium where each firm increases its mining power when P_t reaches the entry threshold $\overline{P}_t^n = e^{-at}\overline{P}_0^n$. As in the baseline model, P_t is a GBM reflected at \overline{P}_t^n , where

$$\overline{P}_{0}^{n} = \frac{n}{n-1} \frac{\beta(r-\alpha)}{\beta-1} K_{0}^{n} \text{ and } K_{0}^{n} = I_{0} + \frac{C_{0}}{r}.$$
(11)

Given that our calibration strategy does not use the analytical expression of the entry threshold to identify its level, it returns the same threshold independently of whether the industry is competitive or not. The degree of competition matters at the second stage, when we infer the cost parameters which are consistent with the threshold. Setting (7) equal to (11), we find that the costs in the oligopolistic and competitive models are proportional as $K_t^n = (1 - 1/n) K_t$. Calibrated costs are lower when the industry is oligopolistic because firms use their market power to extract some rents.

We can easily construct an intuitive measure for the oligopolistic rents. First, note that the net present value of an additional unit of hashpower is equal to $W(P_t) - K_t^n$, where $W(P_t) \equiv \mathbb{E}_t \left[\int_t^{\infty} e^{-r(s-t)} P_s ds \right]$ denotes the expected value of discounted payoffs. The expectation operator for P does not depend on the degree of competition because competition does not affect the calibrated threshold \overline{P} . Hence, free entry is satisfied if and only if $W(\overline{P}_t) = K_t$. It follows that, if we evaluate the net present value of entrants and divide it by the overall costs of entry, we find that the option premium reads

Option Premium =
$$\frac{W(\overline{P}_t) - K_t^n}{K_t^n} = \frac{1}{n-1}.$$
 (12)

As expected, the option premium converges to zero and free entry holds when n goes to infinity. Combining equations (11) and (12) with our estimates allows us to infer the seigniorage income of entrants. Its value as a function of the number of competing firms is reported in Figure 8, along with the selling price of mining rigs. The operating costs account for the remaining share of the expected payoffs.

We can reject high degrees of market concentration by comparing the operating costs to the electricity expenses required to run the mining hardware. Considering that miners pay 3 cents per kilowatt hour (a generously low guess since it amounts to around half of the average market price),³⁶ we find that the discounted electricity costs of entrants amount to 428\$.³⁷ Hence we can rule out the premise that less than 4 firms control the mining market since it would imply that the operating costs are smaller than the electricity costs.³⁸ This is a lower bound on the number of mining firms because operating costs include sizeable maintenance and opportunity costs on top of the electricity

³⁶See the International Energy Agency commentary "Bitcoin energy use-mined the gap" and references therein.

 $^{^{37}}$ The most recent mining rigs at the beginning of the second period were Antminer S4 whose electricity consumption amounted to 700 watts per unit of hashpower (see the online Appendix K). The discounted costs are obtained multiplying their daily consumption of 16.8 kilowatts by the price of electricity, and discounting the resulting daily costs over the lifetime estimated in the online Appendix M.

 $^{^{38}}$ According to the estimates reported in Figure 8, operating costs with 3 and 4 mining firms amount to 322\$ and 450\$, respectively.





Note: Allocations of entrants' revenues predicted by the model with time-to-build at the starting date of the second period (10/01/2014). Hardware costs are inferred from online data, see the online Appendix K for data sources.

consumption. Our model therefore suggests that the seigniorage income of entrants amount to at most 21.5% of their overall revenues. At first sight, this conclusion may appear to be at odds with the evidence in Hileman and Rauchs (2017) and Gencer et al. (2018) that a majority of mining power was controlled by about eight to eleven firms. However, Figure 8 shows that the seigniorage income rapidly decreases in the number of competitors, indicating that, for the levels of concentration supported by the aforementioned evidence, the allocation of revenues is not far from that of a competitive industry.

Input producers.— Since miners were not able to extract large rents, they channeled most of their income towards the producers of their input factors, namely hardware manufacturers and electricity suppliers. As large mining farms are scattered around the globe (with major hubs in China, North-America, Northern and Eastern Europe), evaluating their impact requires a geographic analysis that would go well beyond the scope of this paper. Rauchs et al. (2013) and Hileman and Rauchs (2017) provide the most comprehensive survey of mining locations but, as far as we know, no study has yet used their data to assess the effect that mining has on the revenues of local electricity providers.

By contrast, the production of ASICs was, until recently, a very concentrated activity, with Bitmain claiming a market share of 74.5% of 2017 sales revenues. In their 2018 application proof for an IPO on the Hong Kong Stock Exchange,³⁹ Bitmain indicated that it had been able to generate \$952.6

³⁹Bitmain's IPO prospectus is available at http://templatelab.com/bitmain-ipo-prospectus/. Note that Bitmain was drawing part of its revenues from proprietary mining and mining services. Yet, hardware production accounted for the bulk of Bitmain's activity, namely 90% of its overall revenues.

million in profits out of \$2.5 billion in revenues, thus reporting an healthy profit margin of 37.8% in 2017. To put these numbers into perspective, the overall mining rewards generated by Bitcoin over the same year were equal to \$3.18 billions. Hence, as predicted by our model, ASICs manufacturers managed to capture a significant share of mining revenues. The quasi-monopolistic position enjoyed by Bitmain was finally contested in 2018 by the arrival of new competitors. In particular, Pangolin entered the market in July 2018, proposing a mining rig called Whatsminer M10 that was 30% more efficient than the previous state of the art. The entry of this new competitor triggered a dramatic drop in the price of Bitmain's product (see Figure 5 and the online Appendix K).

5.2 Electricity consumption of Bitcoin

The erosion of Bitmain's dominant position is lowering the cost of entering the mining market. At the same time, price data indicate that Bitcoin is providing lower returns and not exhibiting as much volatility as in the past. We also expect the rate of technological progress to slow down and converge, in the best case scenario, to the value predicted by Moore's law. What will be the impact of these ongoing changes on Bitcoin's electricity consumption? Having a model enables us to answer this question in a quantitative manner.

The entry threshold fully characterizes the industry dynamics for any price trajectory. Most of the time, however, payoffs will be below the threshold. Thus we need to evaluate the payoffs probability distribution in the no-entry region. Fortunately, the long-run distribution of reflected Brownian motions admits a closed-form solution. In order to apply it to our setting, we first notice that the detrended payoff process, $\tilde{P}_t \equiv P_t A_t$, is a GBM reflected at \overline{P}_0 . It is well known (see for instance Grenadier (2002)) that \tilde{P}_t has a long-run stationary distribution whenever $\alpha + a > \sigma^2/2$, a condition which is comfortably satisfied by our calibrated parameters. Using $f_{\tilde{P}}$ to denote the long-run density of \tilde{P} , we find that, for all $y \in (0, \overline{P}_0]$,

$$f_{\tilde{P}}(y) = \frac{\gamma}{y} \left(\frac{y}{\overline{P}_0}\right)^{\gamma}, \text{ where } \gamma \equiv \frac{2\left(\alpha + a - \sigma^2/2\right)}{\sigma^2}.$$

In contrast to \tilde{P} , the network hashrate Q follows a non-stationary process and thus fails to have a long-run distribution. But a simple change-of-variable allows us to compute the steady-state distribution of Q conditional on R and A, as

$$f_Q(Q; R, A) = f_{\tilde{P}}\left(\tilde{P}(Q, R, A)\right) \frac{\partial \tilde{P}(Q, R, A)}{\partial Q}$$
$$= -\gamma Q^{-(\gamma+1)} \left(\frac{RA}{\overline{P}_0}\right)^{\gamma}.$$

Integrating f_Q over the consistent values of Q finally yields its conditional mean

$$\mathbb{E}[Q;R,A] = \int_{\infty}^{\frac{RA}{\overline{P}_0}} Qf_Q(Q;R,A) dQ = \left(\frac{\gamma}{\gamma-1}\right) \frac{RA}{\overline{P}_0}.$$
(13)

The electricity consumption of the network is inversely proportional to the efficiency parameter A. Hence

$$\frac{\mathbb{E}[Q;R,A]}{A} = \left(\frac{\gamma}{\gamma-1}\right)\frac{R}{\overline{P}_0} \tag{14}$$

is the best guess we can make about the long-run energy consumption of Bitcoin. Since (14) is linearly increasing in block rewards, our model confirms the common belief that halvings will lower Bitcoin's electricity consumption as long as the decrease in the number of minted bitcoins is not compensated by an increase in Bitcoin price or in transaction fees. Our contribution consists in characterizing the slope of the relation between R and Q. Figure 9 reports the impact of the parameters on the conditional expectation of Q as well as on its value at the entry threshold. For readability, we set R and A equal to one. We also normalize to one the average value of Q generated by the calibrated parameters. Hence, all changes can be interpreted as percentage deviations from the calibrated model.

First we lower a from its calibrated value to the one consistent with Moore's law. The results reported in the north-west panel of Figure 9 show that a decrease in the rate of technological progress significantly raises the average hashrate. When the rate of obsolescence of hardware decreases, miners are able to devote a greater share of their income to operating costs. Not surprisingly, the growth rate of block rewards α has a positive impact on the level of investment as more miners find it attractive to enter the market. The impact of a decrease in the selling price of mining hardware I is reported in the south-east panel. We use Bitmain's 2017 profit margin of around 30% as an upper-bound on the price correction. When hardware become cheaper, miners enter the market in greater numbers and devote more of their resources to electricity consumption.



Figure 9: Impact of Parameters on the Network's Mining Power

Note: The values at the end of the x-axes correspond to that of the baseline model at the beginning of the second period (10/01/2014).

The impact of the volatility coefficient σ is more intriguing since it has opposite effects on the entry threshold and average hashrate. When the variance of Bitcoin goes up, miners procrastinate more before entering because good shocks are truncated by the entry threshold while nothing prevents payoffs from falling under bad shocks. By the same token, the larger the shocks, the more often payoffs are far below the entry threshold, leaving miners burdened with too much mining power. In other words, the long-run distribution $f_{\tilde{P}}$ becomes less skewed towards the entry threshold when σ goes up. Given that the flattening of the long-run distribution is partially compensated by the increase in the entry threshold, σ has a positive but relatively modest effect on the average hashrate.

Our model also predicts that energy requirements are increasing in the degree of competition among miners. Let $\mathbb{E}^n[Q; R, A]$ denote the conditional expectation of Q when the mining market is oligopolistic with n symmetric firms. Reinserting (11) into (13), we find that an increase in the number of competing firms n raises the network hashrate since $\mathbb{E}^n[Q; R, A] = (1 - 1/n) \mathbb{E}[Q; R, A]$. Given that normative studies conclude that Bitcoin hashrate is too high (see for instance Huberman et al. (2017)), encouraging concentration in the mining market is likely to increase welfare.

What predictions can we draw from these experiments regarding the future of Bitcoin's electricity consumption? Mostly pessimistic ones since increased competition between hardware producers and, above all, a slowdown in the rate of technological progress would worsen Bitcoin's carbon footprint. For these trends to be contained, Bitcoin price will have to increase at a slower rate than the one observed up to now.

6 Conclusion

One of the most enticing promise of Blockchains is their ability to support the maintenance of their infrastructure through a decentralized network. Decentralization has received a lot of attention, becoming a byword for Blockchains. Yet the extent to which Blockchains truly achieve decentralization remains a bone of contention. We contribute to this ongoing debate by analyzing the dynamics of Bitcoin's mining industry. To the best of our knowledge, our paper is the first to structurally estimate the entry decisions of miners, and it supports the premise that the mining market operates under conditions which are close to being competitive. This finding has positive implications for the security of Bitcoin. Given that miners are not able to capture large seigniorage revenues, most of the mining rewards are reinvested into actual hashpower, thus making it more costly to mount a double-spending attack. Moreover, the hashrate of the network is quite resilient to crashes in Bitcoin price because the irreversibility of past investments induces some downward rigidity. This is good news for the security of Bitcoin transactions but bad news for their carbon footprint. Especially since our model predicts that the energy efficiency of the network will deteriorate further if the rate of technological progress decelerates from the high pace it has experienced so far.

Our model will be useful to Bitcoin practitioners since it provides a forecasting tool for investors willing to enter the mining industry. It should also be of interest to researchers studying the optimal design of proof-of-work protocols. For instance, we find that a significant share of mining revenues is not dissipated in electricity consumption, as often argued, but instead spent on mining hardware. Garratt and van Oordt (2020) show that this reinforces the resilience of Bitcoin because miners are more reluctant to join double-spending attacks that depreciate the value of their hardware. Cryptographers have also recognized the importance of taking into account changes in the mining power of the network when analyzing the consistency and liveness of Bitcoin's protocol (see for instance Garay et al. (2017) and Chan et al. (2020)).

Although our model is fairly accurate over the medium run, it remains rather stylized. Further research should strive to relax its assumptions, starting with the premise that the environment is stationary. Since this restriction is hard to maintain over a long horizon, a promising direction would be to embed our framework into a non-stationary environment and allow agents to update their priors.

Our modeling strategy is likely to apply to other cryptocurrencies based on proof-of-work. Taking into account the ability of miners to concurrently mine multiple cryptocurrencies would refine our understanding of their economic incentives (see Aggarwal and Tan (2019)). Future research should also seek to improve our granular understanding of the mining industry by building on the growing amount of geographical data to investigate whether, as suggested by Arnosti and Weinberg (2018), the clustering of mining facilities is explained by cost asymmetries.

References

- Aggarwal, V. and Tan, Y. (2019). A Structural Analysis of Bitcoin Cash's Emergency Difficulty Adjustment Algorithm. mimeo Michael G. Foster School of Business, University of Washington.
- Alsabah, H. and Capponi, A. (2019). Pitfalls of Bitcoin's Proof-of-Work: R and D Arms Race and Mining Centralization. mimeo Columbia University.
- Arnosti, N. and Weinberg, M. (2018). Bitcoin: A Natural Oligopoly. In Blum, A., editor, 10th Innovations in Theoretical Computer Science Conference (ITCS 2019). Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany.
- Athey, S., Parashkevov, I., Sarukkai, V., and Xia, J. (2017). Bitcoin Pricing, Adoption, and Usage: Theory and Evidence. Stanford University Graduate School of Business Research Paper No. 16-42.
- Bachmann, R., Caballero, R., and Engel, E. (2013). Aggregate Implications of Lumpy Investment: New Evidence and a DSGE Model. American Economic Journal: Macroeconomics, 5:29–67.
- Bertola, G. and Caballero, R. (1994). Irreversibility and Aggregate Investment. The Review of Economic Studies, 61:223-246.
- Bertucci, C., Bertucci, L., Lasry, J.-M., and Lions, P.-L. (2020). Mean Field Game Approach to Bitcoin Mining. arXiv:1801.07447.
- Biais, B., Bisière, C., Bouvard, M., Casamatta, C., and Menkveld, A. (2018). Equilibrium Bitcoin Pricing. TSE Working Paper, 18-973.
- Biais, B., Bisière, C., Bouvard, M., and Casamatta, C. (2017). The Blockchain Folk Theorem. TSE Working Paper No. 17-817.
- Bowden, R., Keeler, H., Krezinski, A., and Taylor, P. (2018). Block Arrivals in the Bitcoin Blockchain. arXiv:1801.07447.
- Buchholz, N. (2017). Spatial Equilibrium, Search Frictions and Efficient Regulation in the Taxi Industry. Princeton University.
- Caballero, R. J. and Pyndick, R. S. (1996). Uncertainty, Investment, and Industry Evolution. International Economic Review, 37:641–662.

- Caplin, A. and Leahy, J. (2010). Economic Theory and the World of Practice: A Celebration of the (S,s) Model. *Journal of Economic Perspectives*, 24:183–202.
- Chan, H., Ephraim, N., Marcedone, A., Morgan, A., and Pass, Rafael ANS Shi, E. (2020). Blockchain with Varying Number of Players. Cryptology ePrint Archive, Report 2020/677.
- Chiu, J. and Koeppl, T. (2017). The Economics of Cryptocurrencies Bitcoin and Beyond. Victoria and Queen's University Working Paper.
- Clay, K. and Jones, R. (2008). Migrating to Riches? Evidence from the California Gold Rush. The Journal of Economic History, 68:997–1027.
- Collard-Wexler, A. (2012). Demand Fluctuations in the Ready-Mix Concrete Industry. *The Journal* of *Economic History*, 81:1003–1037.
- Cong, L. W. and He, Z. (2018). Blockchain Disruption and Smart Contracts. NBER Working Paper No. 24399.
- Cong, L. W., He, Z., and Li, J. (2018). Decentralized Mining in Centralized Pools. George Mason University School of Business Research Paper No. 18-9.
- Decker, C. and Wattenhoffer, R. (2013). Information Propagation in the Bitcoin Network. 13-th IEEE International conference on peer-to-peer computing.
- Dixit, A. K. and Pyndick, R. S. (1994). Investment under Uncertainty. Princeton University Press.
- Easley, D., O'Hara, M., and Basu, S. (2019). From Mining to Markets: The Evolution of Bitcoin Transaction Fees. Journal of Financial Economics, 134:91–109.
- Eyal, I. and Sirer, E. G. (2014). Majority is Not Enough: Bitcoin Mining is Vulnerable. In Financial Cryptography and Data Security, vol. 8437 of Lecture Notes in Computer Science. Springer.
- Fernández-Villaverde, J. and Sanches, D. (2016). Can Currency Competition Work? NBER Working Paper No. 22157.
- Foley, S., Karlsen, J. R., and Puntnis, T. J. (2018). Sex, Drugs, and Bitcoin: How Much Illegal Activity Is Financed through Cryptocurrencies? University of Sydney and University of Technology Working Paper.

- Gandal, N., Hamrick, J., Moore, T., and Oberman, T. (2017). Price Manipulation in the Bitcoin Ecosystem. CEPR Working Paper No. DP12061.
- Garay, J., Kiayias, A., and Leonardos, N. (2017). The Bitcoin Backbone Protocol with Chains of Variable Difficulty. CRYPTO, pages 291–323.
- Garratt, R. and van Oordt, M. (2020). Why Fixed Costs Matter for Proof-of-Work Based Cryptocurrencies. Bankc of Canada, Staff Working Paper 2020-27.
- Gencer, A. E., Basu, S., Eyal, I., van Renesse, R., and Sirer, E. G. (2018). Decentralization in Bitcoin and Ethereum Networks. *Financial Cryptography and Data Security (FC)*.
- Grenadier, S. (2000). Equilibrium with Time-to-Build: A Real Options Approach. In Brennan, M. and Trigeorgis, L., editors, Project Flexibility, Agency, and Competition. Oxford University Press, Oxford.
- Grenadier, S. (2002). Option Exercise Games: An Application to the Equilibrium Investment Strategies of Firms. *The Review of Financial Studies*, 15:691–721.
- Grunspan, C. and Pérez-Marco, R. (2017). Double Spend Races. arxiv:1702.02867.
- Hileman, G. and Rauchs, M. (2017). Global Cryptocurrency Benchmarking Study. Cambridge University, Cambridge Center for Alternative Finance, Cambridge, UK.
- Hong, K., Park, K., and Yu, J. (2017). Crowding out in a Dual Currency Regime? Digital versus Fiat Currency. Bank of Korea Working Paper.
- Houy, N. (2016). The Bitcoin Mining Game. LEDGER, 1:53–68.
- Huberman, G., Leshno, J. D., and Moallemi, C. (2017). Monopoly without a Monopolist: An Economic Analysis of the Bitcoin Payment System. Columbia Business School Research Paper No. 17-92.
- Karame, G. O., Androulaki, E., and Capkun, S. (2012). Two Bitcoins at the Price of One? Doublespending Attacks on Fast Payments in Bitcoin. Proceedings of the ACM Conference on Computer and Communications Security (CCS'12), Raleigh, NC, USA.

- Kareken, J. and Wallace, N. (1981). On the Indeterminacy of Equilibrium Exchange Rates. Quarterly Journal of Economics, 96.
- Ma, J., Gans, J. S., and Rabee, T. (2018). Market Structure in Bitcoin Mining. NBER Working Paper No. 24242.
- Mora, C., Rollins, R., Taladay, K., Kantar, M., Chock, M., and Shimada, M. (2018). Bitcoin Emissions alone Could Push Global Warming above 2°c. Nature Climate Change, 8:931–933.
- Nakamoto, S. (2008). Bitcoin, a Peer-to-Peer Electronic Cash System. https://bitcoin.org/ bitcoin.pdf.
- Narayanan, A., Bonneau, J., Felten, E. W., Miller, A., and Goldfeder, S. (2016). Bitcoin and Cryptocurrency Technologies: A Comprehensive Introduction. Princeton University Press.
- Rauchs, M., Blandin, A., Klein, K., Pieters, G., Recanatini, M., and Zhang, B. (2013). 2nd Global Cryptoasset Benchmarking Study. Cambridge University Press.
- Reid, F. and Harrigan, M. (2012). An Analysis of Anonimity in the Bitcoin System. arXiv:1107.4524v2.
- Rosenfeld, M. (2011). Analysis of Bitcoin Pooled Mining Reward Systems. arXiv preprint arXiv, 1112.
- Schilling, L. and Uhlig, H. (2018). Some Simple Bitcoin Economics. NBER Working Paper No. 24483.
- Thomas, J. (2002). Is Lumpy Investment Relevant for the Business Cycle? The Journal of Political Economy, 110:508–534.

7 Appendix

Proof of Proposition 1 Let $W(P_t, \overline{P}_t, A_t) \equiv V(P_t, t) + C_t/r$ denote the value of an entrant net of operating costs as a function of the payoff P_t , the entry threshold \overline{P}_t and the efficiency of the technology A_t . Assumption 4 requires that $dA_t = -aA_tdt$. Assumptions 1 and 2 imply that $dP_t = P_t(\alpha dt + \sigma dZ_t)$ whenever $P_t < \overline{P}_t$ because Q_t remains constant in that region of the payoff space. Finally, the law-of-motion of the entry threshold \overline{P}_t is endogenous, and it is precisely the aim of this proof to show that the market for mining satisfies the equilibrium requirements stated in Definition 1 when \overline{P}_t decreases at the rate of technological progress. Thus we conjecture that $\overline{P}_t = \overline{P}_0/A_t$, with \overline{P}_0 as in Proposition 1, and proceed to show that it is indeed optimal for entrants to wait until $P_t = \overline{P}_t$.

Having specified the law of motion of the three state variables allows us to use Ito's Lemma to derive the Hamilton-Jacobi-Bellman equation satisfied by the value function

$$rW\left(P_{t},\overline{P}_{t},A_{t}\right) = P_{t} + \alpha P_{t}W_{1}\left(P_{t},\overline{P}_{t},A_{t}\right) - a\overline{P}_{t}W_{2}\left(P_{t},\overline{P}_{t},A_{t}\right) + aA_{t}W_{3}\left(P_{t},\overline{P}_{t},A_{t}\right) \\ + \frac{\sigma^{2}}{2}P_{t}^{2}W_{11}\left(P_{t},\overline{P}_{t},A_{t}\right),$$

when $P_t < \overline{P}_t$. Assume that $\alpha \neq r$,⁴⁰ then the general solution of the Hamilton-Jacobi-Bellman equation reads

$$W\left(P_t, \overline{P}_t, A_t\right) = \frac{P_t}{r - \alpha} + \frac{D_1}{A_t} \left(\frac{P_t}{\overline{P}_t}\right)^{\beta_1} + \frac{D_2}{A_t} \left(\frac{P_t}{\overline{P}_t}\right)^{\beta_2},$$

where D_1 and D_2 are constants whose values will be chosen so as to match some boundary conditions, while β_1 and β_2 are the two roots of the following quadratic equation

$$\mathcal{Q}(\beta) \equiv \frac{\sigma^2}{2}\beta(\beta-1) + (\alpha+a)\beta - a - r = 0.$$

Since $\mathcal{Q}(0) = -a - r < 0$ and the coefficient associated to the second order term is strictly positive, we know that one root, β_1 for instance, is strictly positive while the other root, β_2 , is strictly negative. The function W has to satisfy the following three boundary conditions. First, since $P_t = 0$ is an absorbing state, we must have $W(0, \overline{P}_t, A_t) = 0$. This implies that $D_2 = 0$, as otherwise the value

⁴⁰As r tends to α , \overline{P}_0 converges to $\left(I_0 + \frac{C_0}{\alpha}\right)\left(\alpha + a + \sigma^2/2\right)$ and $W\left(P_t, \overline{P}_t, A_t\right)$ tends to $\frac{I_0 + \frac{C_0}{\alpha}}{A_t}\left(\frac{P_t}{\overline{P}_t}\right)\left[1 - \log\left(\frac{P_t}{\overline{P}_t}\right)\right]$.

function would diverge to either minus or plus infinity when P goes to zero. Second, the left continuity of the value function at the entry threshold \overline{P}_t implies that there can be no arbitrage opportunity solely if the value function is flat at the contact point. This requirement, known as the smoothpasting condition, is satisfied when $W_1(\overline{P}_t, \overline{P}_t, A_t) = 0$, i.e. when $D_1 = -\frac{\overline{P}_0}{\beta_1(r-\alpha)}$. Finally, the entry threshold is pinned down by the free entry condition $W(\overline{P}_t, \overline{P}_t, A_t) = I_t + C_t/r$, which implies that $\overline{P}_0 = (I_0 + C_0/r) \frac{(r-\alpha)\beta_1}{\beta_1-1}$.⁴¹ Thus we have found a solution which satisfies all the requirements laid-out in Definition 1 for the existence of a competitive equilibrium.

⁴¹Alternatively, we could have solved the planner's problem and used the "super contact" condition $W_{11}\left(\overline{P}_t, \overline{P}_t, A_t\right) =$

Parameter	Baseline Model	Model with Halvings	Model with Time-to-build
1st period: $04/01/2011$ to $01/31/2013$			
α	2.38	2.38	2.38
σ^2	1.95	1.95	1.95
a	1.18	1.29	1.10
	(0.50)	(0.43)	(0.42)
K_0	\$5.6 mn	\$5.3 mn	\$4.7 mn
	(\$ 16 mn)	(\$ 8.9 mn)	(\$ 17.8 mn)
δ			$11.5 \mathrm{~days}$
			(9.24 days)
2nd period: $10/01/2014$ to $03/31/2017$			
α	0.46	0.46	0.46
σ^2	0.54	0.54	0.54
a	0.76	0.85	0.90
	(0.12)	(0.13)	(0.13)
K_0	1,825	1,655	\$1,465
	(\$ 199)	$(\$ \ 83)$	$(\$ \ 232)$
δ			$46.5 \mathrm{~days}$
			(27.8 days)
3rd period: $08/01/2018$ to $09/19/2020$			
α	0.27	0.27	0.27
σ^2	0.80	0.80	0.80
a	0.76	0.95	0.65
	(0.34)	(0.31)	(0.32)
K_0	\$ 182	\$173	\$160
	(\$ 203)	(\$ 122)	(\$ 152)
δ			$43.5 \mathrm{~days}$
			(21.3 days)
	Interpretation	Estimation method	
α	Trend of block rewards	Maximum likelihood	
σ^2	Volatility of block rewards	Maximum likelihood	
a	Rate of technical progress	Calibration	
K_0	Total Costs	Calibration	
δ	Time-to-Build	Calibration	

Table 1: Calibrated Parameters

Note: Calibrations based on an annual discount rate r = 10%. K_0 is the calibrated total cost per Terahash-second at the first day of each subperiod. All parameters expressed as yearly rates except the time-to-build, δ , which is expressed in days. Standard errors from block bootstrap in parenthesis.