

Energy Equilibria in Proof-of-Work Mining

AMOS FIAT, Tel Aviv University

ANNA KARLIN, University of Washington at Seattle

ELIAS KOUTSOPIAS, University of Oxford

CHRISTOS PAPADIMITRIOU, Columbia University

The Bitcoin protocol induces miners, through monetary rewards, to expend energy in order to add blocks to the chain. We show that, when energy costs are substantial and taken into account, counterintuitive and unintended strategic behavior results: In a simple bounded-horizon setting with two identical miners there is a unique pure symmetric equilibrium in which both miners first “slow down” in order to decrease the crypto complexity and then take advantage of this decrease. If miners have different energy efficiencies and are restricted to choose the same hash rate for many epochs, there is a unique pure equilibrium in which miners either participate at low levels that depend in intricate ways on all the other miners’ efficiencies, or choose to abstain from mining if their efficiency is too low. In the general setting in which miners can adapt their hash rates over time, we show that, unless the number of miners is very small, the only possible pure equilibria are rather chaotic, with miners quitting and starting again periodically — or there is no pure equilibrium at all. We discuss the implications of these results for the stability of proof-of-work protocols.

ACM Reference Format:

Amos Fiat, Anna Karlin, Elias Koutsoupias, and Christos Papadimitriou. 2019. Energy Equilibria in Proof-of-Work Mining. In *ACM EC '19: ACM Conference on Economics and Computation (EC '19), June 24–28, 2019, Phoenix, AZ, USA*. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3328526.3329630>

1 INTRODUCTION

The protocol described in the 2008 paper “Bitcoin: A peer-to-peer electronic cash system” by Satoshi Nakamoto [9] (hence called “the Satoshi protocol”) is a singular achievement: a bold, novel system design that has spawned, without much debugging a decade later, a global distributed system with millions of users that is surprisingly robust.

The backbone of Bitcoin and of similar cryptocurrencies is a proof-of-work blockchain protocol which attempts to keep a consistent list of transactions in a peer-to-peer network. This list of transactions, the public ledger, is maintained by the users of the network, who constantly attempt to extend the blockchain, a public data structure consisting of a sequence of blocks of transactions. To add a block to the blockchain and claim some reward, a user has to provide proof-of-work which takes the form of an easily verifiable solution to a hard cryptographic puzzle. This process is called mining and we will use the term miner to refer to the users and nodes of the distributed network. For excellent introductions to Bitcoin and cryptocurrency technologies, see [10, 14].

There is only one known major fault of the Satoshi protocol: The original paper by Nakamoto claimed that as long as no miner has a majority of the mining power, no one would have reason to deviate from the protocol — that is to say, it was claimed that the Satoshi protocol is *incentive compatible*. Recently, Eyal and Sirer [7] disproved this claim by showing that there are circumstances

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC '19, June 24–28, 2019, Phoenix, AZ, USA

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6792-9/19/06...\$15.00

<https://doi.org/10.1145/3328526.3329630>

where it is advantageous for a miner (with less than half the total power) to deviate from the Satoshi protocol and delay the publication of a new block. There has been much subsequent work along these lines, e.g., [8, 11–13]. Significantly, it was recently pointed out that in the regime in which transaction fees, as opposed to block reward by new Bitcoins, will be the main monetary reward for mining — as it is expected to be eventually — the increased variance of rewards will incentivize even small miners to occasionally fork the chain [5].

Our main contribution is *a new genre of strategic deviations from the intended function of the Satoshi protocol, related to energy costs and crypto puzzle difficulty*. Critically, the deviations we consider *are not* dependent upon having a large fraction of the computing power, any constant fraction will do. An important part of the Satoshi protocol is the *difficulty adjustment process*, whereby the complexity (inverse probability of success for a single hash) of solving the crypto puzzle is recomputed every two weeks, or more precisely, every 2016 blocks (an *epoch*), with the goal that the expected length of time between successive blocks in the upcoming epoch is kept to ten minutes, under the assumption that the rate with which hashes are computed does not change between the previous epoch and the next¹. The idea of using difficulty adjustment appears in [4, 6] but in a different context of replacing one history with another.

We show that, once energy costs are taken into consideration, Bitcoin mining entails strategic behavior deviating from the implicit intension of the Satoshi protocol — namely, that each miner exerts the same amount of effort each epoch, presumably using all of the resources they have available. It is remarkable however that the strategic behavior we describe is not an actual deviation from the Satoshi protocol; in the original protocol there is no explicit suggestion that users should use their full hashpower every round. Our results are most relevant in a regime in which energy costs are a substantial part of mining revenue, which seems to be the case now: it is estimated (see for example [1, 2]) that this has been the case during the past year, with current estimates of energy costs hovering higher than 90% of revenue.

Strategic deviations and equilibria described in this paper hold for arbitrary proof-of-work cryptocurrencies that include a difficulty adjustment feature. It is unclear how a proof-of-work system could avoid having such a feature — but see discussion at the end of the paper. Bitcoin is hardly the only proof-of-work cryptocurrency. A [very very] long list of proof-of-work cryptocurrencies (and associated market capitalization) appears in [3].

We model mining as a game between miners in which the utility for each miner is the miner's revenue minus the miner's energy costs **per unit time**; we assume that each miner i has a specific energy cost per hash, denoted α_i . We analyze several such games:

- In a warm-up toy game (Section 2) in which two miners with identical energy costs α compete by adjusting their hash rate for two epochs, there is a unique symmetric equilibrium in which both miners hold back their hashing effort in order to “game down” the difficulty adjustment process, and benefit from this decreased difficulty in the last epoch.
- In Section 3 we turn to the *fixed effort* game, in which n miners with hashing costs $\alpha_1 \leq \alpha_2 \leq \dots \leq \alpha_n$ must each choose an effort level which will be kept constant for a long run of epochs; we completely characterize the unique pure equilibrium of this game. At equilibrium, the first k^* miners will participate with some positive effort, and the rest will abstain. The number k^* of participating miners, and the effort level at which they will participate, depends on all α_i 's in a rather intricate way: the $k + 1$ st miner will participate if and only if its α_{k+1} is no larger than an “enhanced average” of the first k . At least two miners will participate — this

¹Because of this provision, puzzle difficulty and energy consumption has been increasing rapidly over the past decade (despite a recent drop), an issue that has justifiably attracted much attention: Bitcoin mining now costs the world a good fraction of one percent of total energy consumption, bringing it on par with Portugal.

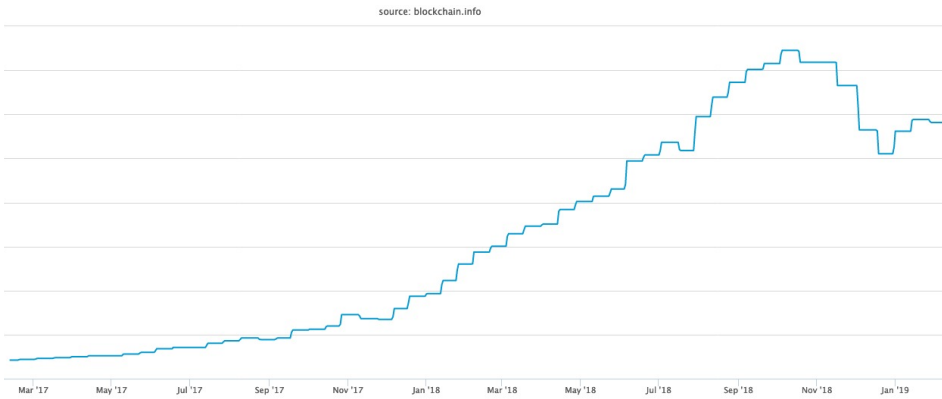


Fig. 1. The chart of the Bitcoin mining difficulty for the last two years

is intuitive, since a single very efficient miner wants to lower her effort, and so the second most efficient miner will “squeeze in”. Of course, if there are only two miners participating then the more powerful miner can also rewrite history. If there are more miners with equal or similar costs as the second one, they will also participate. At the other extreme, n miners with equal α_i 's will each get at equilibrium a revenue of $\frac{1}{n^2}$ per block, in contrast to the social optimum of $\frac{1}{n}$ per miner, were they able to cooperate through a contract.

- But of course the fixed effort assumption is not realistic. In Section 4, we consider the case where the n players are free to adapt their hashing rates from one epoch to the next. For this setting, we show that with sufficiently many miners, either there are no pure equilibria at all, or they must be somewhat bizarre: miners will vary their efforts from epoch to epoch. The number of miners needed for this result to kick in is parameterized by the maximum fraction of the total hash power available to a single miner. It seems that this captures a host of plausible scenarios regarding the number of agents and the fraction of the total effort that can be exerted by a single agent.
- Since equilibria are unlikely, we next look at the nature of the best response by one agent, who is free to adapt her hash rate arbitrarily, to fixed hash rates by other agents. In Section 5 we present closed form equations for the best response, whose accuracy depends on a mathematical conjecture that we articulate, derived after much experimentation and analysis through Mathematica. The results suggest that the effort level of the agents may vary dramatically and chaotically from epoch to epoch.

We conclude with a discussion and open problems.

2 MODEL

Definitions:

- (1) An epoch is a period during which 2016 blocks are generated. Epoch $t \geq 1$, starts immediately after the $(t - 1) \cdot 2016^{th}$ block is generated, and ends when the next 2016 blocks are generated. Let ℓ_t denote the length of epoch t in “time units”. One time unit is 10 minutes of elapsed time.
- (2) Let m_i , $1 \leq i \leq n$, be the maximum number of hashes the i^{th} miner can perform per time unit. We mainly use this to assign some meaning to the “full throttle” idea; in almost all of the games studied below we shall assume that the maximum number of hashes is not binding,

however we will generally not consider the possibility of a miner having a large percentage of the network hash power, and certainly not more than 50%.

- (3) Let $0 \leq h_i^t \leq m_i$, $1 \leq i \leq n$, $t \geq 1$, be the number of hashes per time unit actually performed by miner i during the t^{th} epoch. If $h_i^t = m_i$ we say that miner i is going *full throttle* during epoch t , and we say that miner i is *holding back* otherwise.
- (4) Let $H^t = \sum_i h_i^t$, the total number of hashes performed per time unit during epoch t .
- (5) The hardness of the puzzle in epoch t is expressed in terms of the probability, p_t that a single hash will solve the puzzle. This is recomputed when epoch $t - 1$ ends, so that the expected length of epoch t is one time unit, under the assumption that $H^t = H^{t-1}$. Now, $1/p_t$ is the expected number of hashes until the crypto puzzle is solved, and H_{t-1} is the total number of hashes in epoch $t - 1$. Thus, p_t is computed as follows:

$$\frac{1}{p_t} = H^{t-1}.$$

The *utility* for player i over the first T epochs is the profit over T epochs (i.e., the number of Bitcoins minus energy cost in Bitcoins) divided by the length of time for these T epochs.

DEFINITION 1 (ENERGY COST α_i). We denote by α_i the cost (in Bitcoin) of the electricity required for one hash using the technology available to miner i .

To derive a formula for utility, we observe that:

- (1) The length of epoch t is H^{t-1}/H^t (in time units).
- (2) Thus, the energy cost for miner i in epoch t is

$$\alpha_i \cdot \frac{H^{t-1}}{H^t} \cdot h_i^t.$$

Using these, we now define the utility of each agent.

DEFINITION 2 (UTILITY). The utility per time unit for agent i , averaged over the first T epochs, is given by

$$U_i = \frac{\sum_{t=1}^T \frac{h_i^t}{H^t} (1 - \alpha_i \cdot H^{t-1})}{\sum_{t=1}^T \frac{H^{t-1}}{H^t}}. \tag{1}$$

Note: In the definitions of utility and of puzzle difficulty we omit the factor 2016 reflecting the number of blocks in one epoch; this is inconsequential, akin to a change in currency from 1 Bitcoin to 2016 Bitcoins.

Example: A finite horizon game

We illustrate these definitions through a simple example, which also showcases the kind of phenomenon that we study. Suppose that two identical players with hash costs α and maximum hash power $m_1 = m_2 = \frac{1}{2}$, are about to play for only two epochs, starting from some fixed puzzle complexity at the first epoch. Rational players will go “full throttle” during the last epoch; this is so because the utility is increasing in h_i^T (in Equation 1, the numerator is increasing and the denominator is decreasing in h_i^T).

There are two unknowns, the hash rates at the first epoch, call them h_1 and h_2 by dropping superscripts. The length of epoch 1 is $1/(h_1 + h_2)$, while the length of the epoch 2 is $h_1 + h_2$. For $\alpha > 1$, the total duration of epochs 1 and 2 is $h_1 + h_2 + 1/(h_1 + h_2) \geq 2$.

It follows from equation 1 that utility of agent $i \in \{1, 2\}$ is

$$U_i = \frac{\frac{h_i}{h_1+h_2} - \alpha \cdot h_i \cdot \frac{1}{h_1+h_2} + \frac{1}{2} - \alpha \frac{1}{2}(h_1 + h_2)}{\frac{1}{h_1+h_2} + (h_1 + h_2)}. \quad (2)$$

We observe that

- When $\alpha < 1/2$, the utility U_i is increasing in h_i . This means that the unique Nash equilibrium is full throttle for both players.
- When $\alpha \in [1/2, 3/2)$, to find the Nash equilibria we consider the points with $\partial U_i / \partial h_i = 0$ or the extreme values 0 and $1/2$ of h_i .

By setting both derivatives to 0,

$$\frac{\partial U_1}{\partial h_1} = 0, \quad \frac{\partial U_2}{\partial h_2} = 0,$$

we get that

$$h_1 = h_2 = \frac{-\alpha + \sqrt{3 - 2\alpha + \alpha^2}}{2}.$$

For $\alpha > 3/2$ the solution to the above equation gives an infeasible value of h_i (negative) which means that players use ever smaller hash levels, and no equilibria is possible.

For $\alpha \in [1/2, 3/2)$, the above solution gives the unique Nash equilibrium. What this means is that equilibria exists only if $\alpha \leq 3/2$ and we have that

$$h_i = \begin{cases} 1/2, & \alpha \leq 1/2 \\ \frac{-\alpha + \sqrt{3 - 2\alpha + \alpha^2}}{2}, & 1/2 \leq \alpha \leq 3/2 \end{cases} \quad U_i = \begin{cases} (1 - \alpha)/2, & \alpha \leq 1/2 \\ \frac{-\alpha + \sqrt{3 - 2\alpha + \alpha^2}}{4}, & 1/2 \leq \alpha \leq 3/2 \end{cases}$$

In conclusion, if the miners have very low costs ($\alpha \leq \frac{1}{2}$) they will mine at full speed. But for higher cost in the range $(\frac{1}{2}, \frac{3}{2}]$ both miners hold back and mine during the first epoch below their hash capacity in order to bring the puzzle difficulty down and exploit it in the last epoch.

3 THE CONSTANT EFFORT GAME

We first consider a setting where miners do not change the number of hashes between different epochs, *i.e.*, $h_i^t = h_i^{t'} = h_i$ for all t, t' – this assumption is revisited and removed in the next sections. The n miners differ in their efficiency levels with energy costs (per single hash) of $\alpha_1 < \alpha_2 < \dots < \alpha_n$. The lower the miner index the better the technology and the smaller the cost. Formally, this is a game where the strategic decision of a miner is the choice of a single hash level that they will use throughout. For now, we also assume that there is no upper bound on the maximal number of hashes that can be performed by a miner ($m_i = \infty$).

In this section we prove the following theorem:

THEOREM 1. *Given agents with hash costs $\alpha_1 \leq \alpha_2 \leq \dots \leq \alpha_n$, $n \geq 2$, there is a unique equilibrium of the following form:*

(1) *The participating agents are those with the lowest costs, $\alpha_1, \dots, \alpha_k^*$, for some $2 \leq k^* \leq n$ where*

$$k^* = \max \left\{ 2 \leq k \leq n \mid \alpha_k \leq \frac{\sum_{\ell=1}^k \alpha_\ell}{k-1} \right\}.$$

(2) The hash rates in the equilibrium for all $i \in \{1, \dots, k^*\}$ are given by

$$h_i^* = \frac{k^* - 1}{\sum_{j=1}^{k^*} \alpha_j} \left(1 - \frac{k^* - 1}{\sum_{j=1}^{k^*} \alpha_j} \cdot \alpha_i \right) \quad (3)$$

$$= \sqrt{H_{-i}/\alpha_i} - H_{-i}, \quad \text{where} \quad H_{-i} = \sum_{1 \leq \ell \leq k^*, \ell \neq i} h_\ell^*. \quad (4)$$

(3) The hash rates h_i^* for $k^* < i \leq n$ are all zero.

(4) The utility of agent $i \in \{1, \dots, k^*\}$, in equilibrium (in each epoch) is:

$$U_i = (1 - \alpha_i H)^2 \quad \text{where} \quad H = \sum_{\ell=1}^{k^*} h_\ell^*.$$

Remarks:

- (1) Note that part (1) of the theorem implies that in equilibrium there are always at least two participants if $n \geq 2$. This is intuitively obvious since a single miner, no matter how powerful, will always attempt to reduce her hash rate so that the puzzle difficulty will drop to 0. However, this will always induce another miner, no matter how inefficient, to join in.
- (2) Of course, two participants implies that one of the miners has 50% of the mining power, which allows her to rewrite history. More generally, under the assumption that agents use constant effort, Theorem 1 can be used to derive conditions under which the system falls apart. If the unique equilibria has a very powerful miner, the same miner can go off and rewrite history.
- (3) We derive the hash power in equilibrium, h_i^* , in two ways: Equation (3) describes the equilibrium hash level of an agent as a function of all energy costs, whereas (4) describes the equilibrium hash level of an agent as a function of hash power used by the others and her own energy level.

We proceed with the proof of the theorem.

LEMMA 1. *The best response for agent i given that all other agents have (jointly) hash level H_{-i} , is*

$$br_i(H_{-i}) = \max(0, \sqrt{H_{-i}/\alpha_i} - H_{-i}),$$

moreover, the best response is > 0 iff $1 - \alpha H_{-i} > 0$.

PROOF. Consider the utility of agent i given in Equation (1). As the denominator is constant, ($H^t = H^{t'}$ for all t, t') it follows the best response for agent i to the others making effort H_{-i} ,

$$br_i(H_{-i}) = \operatorname{argmax}_{h_i} h_i / (h_i + H_{-i}) - \alpha_i h_i \quad (5)$$

$$\begin{aligned} &= \operatorname{argmax}_{h_i} 1 - H_{-i} / (h_i + H_{-i}) - \alpha_i h_i \\ &= \operatorname{argmin}_{h_i} \alpha_i h_i + H_{-i} / (h_i + H_{-i}). \end{aligned} \quad (6)$$

Taking the derivative of Equation (6) with respect to h_i and setting this to zero we get that the best response hash rate for miner i has

$$br_i(H_{-i}) = \sqrt{H_{-i}/\alpha_i} - H_{-i}.$$

Now, we have $br(H_{-i}) > 0$ iff $H_{-i}/\alpha_i > (H_{-i})^2$, which is equivalent to $1 - \alpha H_{-i} > 0$. \square

LEMMA 2. *Let P be an arbitrary subset of $\{1, \dots, n\}$ such that every miner $i \in P$ has hash rate $h_i > 0$ and all miners $j \in \{1, \dots, n\} \setminus P$ have hash rate $h_j = 0$. Then, h_1, \dots, h_n is a Nash Equilibrium if and only if*

(1)

$$H = \sum_{i \in P} h_i = \frac{|P| - 1}{\sum_{j \in P} \alpha_j}, \quad (7)$$

 (2) For all $i \in P$ we have that

$$h_i = H(1 - H \cdot \alpha_i) = \frac{|P| - 1}{\sum_{j \in P} \alpha_j} \left(1 - \frac{|P| - 1}{\sum_{j \in P} \alpha_j} \cdot \alpha_i \right). \quad (8)$$

(3)

$$1 - \alpha_i H > 0 \quad \text{if and only if} \quad i \in P \quad (9)$$

PROOF. Suppose that $\{h_i\}$ is an equilibrium and let $H_{-i} = \sum_{j \neq i} h_j$. Then by Lemma 1, for $i \in P$, we have that

$$h_i = \sqrt{\frac{H_{-i}}{\alpha_i}} - H_{-i} \Leftrightarrow \sqrt{\frac{H_{-i}}{\alpha_i}} = H \Leftrightarrow H_{-i} = \alpha_i H^2 \quad \text{for all } i \in P. \quad (10)$$

Therefore, for $i \in P$

$$h_i = H - H_{-i} = H - \alpha_i H^2 = H(1 - \alpha_i H). \quad (11)$$

Summing over $i \in P$, we get

$$H = \sum_{i \in P} h_i = H \left(|P| - H \sum_{i \in P} \alpha_i \right).$$

Cancelling H on both sides yields (7) and substituting (7) back into (11) yields (8). Moreover, for each $i \in P$, $h_i > 0$ and h_i is a best response to $\{h_{-j}\}_{j \neq i}$ only if

$$1 - \alpha_i H = 1 - \frac{|P| - 1}{\sum_{j \in P} \alpha_j} \cdot \alpha_i > 0.$$

Finally, by Lemma 1, we have that $h_i = 0$ is a best response only if $1 - \alpha_i H_{-i} \leq 0$. But for $i \notin P$, $H_{-i} = H$, so $h_i = 0$ is a best response if $1 - \alpha_i H \leq 0$ completing the claim that (9) must hold.

Next we argue that if (7), (8) and (9) hold, then $\{h_i\}$ is a Nash equilibrium. From (7), (8) and (9), we have that all $h_i \in P$ are positive and the best response for all $i \notin P$ is $h_i = 0$. So we have just left to verify that for $i \in P$

$$h_i = \sqrt{H_{-i}/\alpha_i} - H_{-i}. \quad (12)$$

Subtracting (8) from (7), we get that for $i \in P$,

$$H_{-i} = \alpha_i H^2$$

which by (10) is equivalent to (12), which completes the proof that $\{h_i\}$ is a Nash equilibrium. \square

These two lemmas imply that setting hash values for agents in the set $P^* = \{1, \dots, k^*\}$ as given in Equation (3) and zero hash values for others is a Nash equilibrium. Next, we verify that this equilibrium is unique. Suppose that there is another equilibrium $\{h'_i\}$ satisfying (7), (8) and (9). Let $P' := \{i | h'_i > 0\}$, $H' := \sum_{j \in P'} h'_j$.

First suppose that $P' = \{1, \dots, k'\}$. If $k' < k^*$, then $1 - \alpha_{k^*} H' > 0$ contradicting (9). Similarly, if $k' > k^*$, then $1 - \alpha_{k'} H' \leq 0$, again contradicting (9). Finally if P' has some miner i with $h'_i = 0$ and another miner j with $h'_j > 0$, where $\alpha_i < \alpha_j$, then since $1 - \alpha_j H' > 0$, it also holds that $1 - \alpha_i H'_{-i} > 0$ contradicting (9).

Finally, we observe that substituting (7) and (8) into the utility expression 5 yields part 4 of Theorem 1 and completes its proof. \square

On price of anarchy and technological innovation.

Assuming constant effort per agent, Theorem 1 has some interesting consequences. When all miners have equal hash costs $\alpha_i = \alpha$ for all i , we find that the equilibrium hash rate is $h_i = (n - 1)/n^2\alpha$ and each agent's utility per epoch is $U_i = 1/n^2$. In contrast, if the miners could be bound by a contract to mine at an agreed-upon effort level, they could drive the puzzle difficulty arbitrarily low and have utility arbitrarily close to $1/n$: *The price of anarchy in this case is n .*

Suppose now that a new miner with significantly better technology, say $\epsilon \cdot \alpha$, enters the game. The remaining n miners still participate, but at a lower level $h = 1 - \frac{1}{1+\epsilon/n} = \epsilon/n + O((\epsilon/n)^2)$, with the new miner contributing much larger effort $h' = 1 - \epsilon + O((\epsilon/n)^2)$. The utility of the innovator will be $1 - \epsilon + O(\epsilon^2/n)$ with the original players receiving $\epsilon^2/n^2 + O(\epsilon^4/n^4)$, an ϵ^2 fraction of their past utility, while the price of anarchy is elevated near one.

Finally, if a second miner with significantly better technology than the original miners enters, all the original miners will drop out of the game.

4 CONSTANT EFFORT IS NOT AN EQUILIBRIUM WHEN AGENTS CAN VARY THEIR EFFORT

To contrast with results in the previous section, which only hold under the constant effort strategy space, we now assume that miners are free to change effort from one epoch to another up to some $\gamma + 1/n$ times the combined effort of the other miners, for some constant γ . Clearly, if any individual miner has too much of the total power then history can be rewritten, and even if not so high, attacks of the type introduced by Eyal and Sirer [7] can come into play. Think of γ as some constant less than, say, 0.3. Clearly, reasonable values of γ are small.

Theorem 2 shows that the space of possible pure equilibria is limited.

THEOREM 2. *For any γ , there is a value $n_\gamma = 4(1 + \gamma)^2/\gamma$ such that there is no Nash Equilibrium where*

- $n > n_\gamma$ miner participate in every epoch.
- There is some miner that uses at most a $1/n_\gamma$ fraction of the total effort in every epoch, but has the capacity to increase his hash power to $\gamma + 1/n$ times the total effort of the others.

Before giving the proof of this theorem, we note that, for every epoch, there is always at least one agent that uses no more than $1/n$ of the total effort. The theorem requires something stronger, that in the pure equilibrium, there is a miner that consistently uses no more than $1/n$ in every epoch. Of course if miners use constant effort across all epochs then there is always one agent that uses no more than $1/n$ of the total effort, in all epochs. Thus we derive the following Corollary to Theorem 2.

COROLLARY 1. *There exists no constant effort equilibrium in which $n \geq n_\gamma$ miners have spare power to increase their power to $\gamma + 1/n$ times the total effort of the others.*

Note that this Corollary contrasts with but does not contradict Theorem 1, because the strategy space considered for Theorem 1 restricts miners to choosing one constant effort level across all epochs whereas the agents considered in Theorem 2 can choose arbitrary effort levels that differ from epoch to epoch.

We now give the proof of Theorem 2:

PROOF. Suppose that in the hypothesized equilibrium, on round t the cumulative hash power used is H^t . Suppose also without loss of generality that T is even and

$$\sum_{t=1}^{T/2} H^{2t} \geq \sum_{t=1}^{T/2} H^{2t-1} \quad (13)$$

Further denote the fraction of hash power used by agent 1 in epoch t by f_1^t which we assume to be at most $1/n$ for all t . Then it follows from Equation 1 that agent 1's utility is

$$\begin{aligned} U_1 &= \frac{\sum_{t=1}^{T/2} [f_1^{2t}(1 - \alpha_1 H^{2t-1}) + f_1^{2t+1}(1 - \alpha_1 H^{2t})]}{\sum_{t=1}^{T/2} \left(\frac{H^{2t}}{H^{2t-1}} + \frac{H^{2t+1}}{H^{2t}} \right)} \\ &\leq \frac{\frac{1}{n} \sum_{t=1}^{T/2} (2 - \alpha_1 H^{2t-1} - \alpha_1 H^{2t})}{\sum_{t=1}^{T/2} \left(\frac{H^{2t}}{H^{2t-1}} + \frac{H^{2t+1}}{H^{2t}} \right)}. \end{aligned} \quad (14)$$

Now consider the following agent 1 deviation:

$$\tilde{h}^{2t} = (\gamma + f_1^{2t})H^{2t} \quad \text{and} \quad \tilde{h}^{2t+1} = 0, \quad \forall t.$$

Then the total hash power used in each round post deviation is

$$\tilde{H}^{2t} = (1 + \gamma) \cdot H^{2t} \quad \text{and} \quad \tilde{H}^{2t+1} = (1 - f_1^{2t+1})H^{2t+1}, \quad \forall t.$$

Therefore, agent 1's utility post deviation is:

$$\begin{aligned} \tilde{U}_1 &= \frac{\sum_{t=1}^T \frac{\tilde{h}^t}{\tilde{H}^t} (1 - \alpha_1 \cdot \tilde{H}^{t-1})}{\sum_{t=1}^T \frac{\tilde{H}^{t-1}}{\tilde{H}^t}} \\ &= \frac{\sum_{t=1}^{T/2} \left(\frac{\gamma + f_1^{2t}}{1 + \gamma} \right) (1 - \alpha_1 (1 - f_1^{2t-1})H^{2t-1})}{\sum_{t=1}^{T/2} \left((1 + \gamma) \frac{H^{2t}}{(1 - f_1^{2t-1})H^{2t-1}} + \frac{(1 - f_1^{2t+1})H^{2t+1}}{(1 + \gamma)H^{2t}} \right)} \\ &\geq \frac{\sum_{t=1}^{T/2} \left(\frac{\gamma}{1 + \gamma} \right) (1 - \alpha_1 H^{2t-1})}{\sum_{t=1}^{T/2} \left(\frac{(1 + \gamma)H^{2t}}{(1 - 1/n)H^{2t-1}} + \frac{H^{2t+1}}{(1 + \gamma)H^{2t}} \right)}, \end{aligned}$$

since $0 \leq f_1^t \leq 1/n$ for all t . Therefore, the utility post deviation

$$\tilde{U}_1 \geq \frac{\left(\frac{\gamma}{1 + \gamma} \right) \sum_{t=1}^{T/2} (1 - \alpha_1 H^{2t-1})}{\left(\frac{1 + \gamma}{1 - 1/n} + \frac{1}{1 + \gamma} \right) \sum_{t=1}^{T/2} \left(\frac{H^{2t}}{H^{2t-1}} + \frac{H^{2t+1}}{H^{2t}} \right)}$$

which by (13) is at least

$$\frac{\left(\frac{\gamma}{1 + \gamma} \right) \sum_{t=1}^{T/2} \frac{(2 - \alpha_1 H^{2t-1} - \alpha_1 H^{2t})}{2}}{\left(\frac{1 + \gamma}{1 - 1/n} + \frac{1}{1 + \gamma} \right) \sum_{t=1}^{T/2} \left(\frac{H^{2t}}{H^{2t-1}} + \frac{H^{2t+1}}{H^{2t}} \right)}.$$

Combining this with Equation (14), we get that

$$\tilde{U}_1 / U_1 \geq \frac{\frac{\gamma}{1 + \gamma}}{\frac{1 + \gamma}{1 - 1/n} + \frac{1}{1 + \gamma}} \cdot \frac{n}{2}.$$

If the last expression is greater than 1, the miner has an incentive to switch to the new strategy. This happens for all values of n exceeding a threshold (that depends on γ). It is straightforward to verify that for every

$$n \geq \frac{4(1 + \gamma)^2}{\gamma},$$

the left hand side of the above inequality is at least 1, which shows that the deviation is profitable. \square

5 FURTHER RESTRICTIONS ON THE EQUILIBRIA OF THE ADAPTIVE EFFORT GAME

We already know that there is no pure Nash equilibria when all agents have constant effort (Corollary 1). But, if everybody else plays at constant effort, what precise form will the deviation of the last player take?

The characterization of the best response described in this section crucially depends on Conjecture 1 below. Note that this is stated as a conjecture and not as a lemma as we were unable to prove it. However, we do have ample evidence derived from experiments and Mathematica that offers strong support for the conjecture.

So, suppose that the total hashing power H_{-i} of all miners except miner i is the same in all epochs. What is the best response of the single miner? Perhaps not so surprisingly at this point in the paper, the best response varies in general with time. To analyze the situation and keep the notation simple, let's define

$$x^t = \alpha_i H_{-i}^t \qquad \beta = \sqrt{\alpha_i H_{-i}}.$$

The utility of player i given in Equation (1), can be rewritten as

$$U_i = \frac{\sum_{t=1}^T \frac{(x^t - \beta^2)(1 - x^{t-1})}{x^t}}{\sum_{t=1}^T \frac{x^{t-1}}{x^t}}.$$

The best response for player i is to select $h_i^t \geq 0$, $t = 0, \dots, T$, or equivalently $x^t \geq \beta^2$, to maximize the above quantity. This maximization problem is affected by the boundary conditions, that is the values of x^0 and x^T , but their effect is limited and it almost disappears as T tends to infinity. We don't know the solution to this maximization problem, but we have strong evidence, including experimental results, that it satisfies the following conjecture.

CONJECTURE 1. *For every $\beta \in (0, 1)$ and for every even T , there exists x^0 and x^T , such that the optimal values x^t that maximize the quantity*

$$U^* = \sup_{x^t \geq \beta^2} \frac{\sum_{t=1}^T \frac{(x^t - \beta^2)(1 - x^{t-1})}{x^t}}{\sum_{t=1}^T \frac{x^{t-1}}{x^t}},$$

are 2-periodic, that is, $x^t = x^{t+2}$ for every $t \leq T - 2$.

Note that the class of 2-periodic solutions includes the constant effort solutions. The following lemma characterizes the unique 2-periodic solution of the above maximization problem. It asserts that for small β , the optimal solution is the same for all t , while for large β , it alternates between the minimum value and some other value.

LEMMA 3. Assume that Conjecture 1 holds, that is, assume that the solution to the above maximization problem is 2-periodic. Then the optimal solution is

$$x^t = \begin{cases} \beta & \beta \leq 4\sqrt{2} - 5 \\ \beta^2 & t \text{ is odd} \\ (1 + \sqrt{2})\beta^2 & t \text{ is even} \end{cases} \quad \text{otherwise,} \quad (15)$$

and the maximum value of the expression is

$$U^* = \begin{cases} (1 - \beta)^2 & \beta \leq 4\sqrt{2} - 5 \\ \frac{\sqrt{2}-1}{2}(1 - \beta^2) & \text{otherwise.} \end{cases} \quad (16)$$

PROOF. Since we consider only 2-periodic solutions we can express the problem as follows: find x^0, x^1 in $[\beta^2, \infty)$ to

$$\max \frac{\frac{(x^1 - \beta^2)(1 - x^0)}{x^1} + \frac{(x^0 - \beta^2)(1 - x^1)}{x^0}}{\frac{x^0}{x^1} + \frac{x^1}{x^0}},$$

which can be solved by standard methods. \square

We now translate the above lemma to obtain the best response against constant effort by the other miners. If the above conjecture holds, this is the unique best pure response.

LEMMA 4. If the total hashing power H_{-i} of all miners except miner i is the same for all epochs, the best 2-periodic response of player i , when it has unlimited hashing power, is given by

- if $\alpha_i H_{-i} \leq (4\sqrt{2} - 5)^2 \approx 0.43$, player i uses the following hashing power on every epoch: $h_i = \sqrt{H_{-i}/\alpha_i} - H_{-i}$,
- if $(4\sqrt{2} - 5)^2 < \alpha_i H_{-i} \leq 1$, player i mines only in every second epoch with the following hashing power: $h_i = \sqrt{2}H_{-i}$,
- otherwise, player i does not mine at all.

The player's utility is given by

$$U_i = \begin{cases} (1 - \sqrt{\alpha_i H_{-i}})^2 & \alpha_i H_{-i} \leq (4\sqrt{2} - 5)^2 \\ \frac{\sqrt{2}-1}{2}(1 - \alpha_i H_{-i}) & (4\sqrt{2} - 5)^2 < \alpha_i H_{-i} \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

From this general characterization of the best 2-periodic response against fixed power we can arrive at some useful conclusions. The first corollary below shows that for two players there is always a Nash equilibrium in which both players use the same power in every epoch. On the contrary, the next corollary shows that for $n \geq 3$, there is no symmetric Nash equilibrium.

COROLLARY 2. For $n = 2$ miners and assuming that Conjecture 1 holds, there exists a Nash equilibrium in which player $i = 1, 2$ uses hashing power $h_i = \alpha_{3-i}/(\alpha_1 + \alpha_2)^2$ in every epoch.

PROOF. Suppose that player $3 - i$ plays the strategy of the corollary, $h_{3-i} = \alpha_i/(\alpha_1 + \alpha_2)^2$, in every epoch. Then by Lemma 4, it suffices to show that $\alpha_i H_{-i} \leq (4\sqrt{2} - 5)^2$, in which case the best response of player i is to use the same power $\sqrt{H_{-i}/\alpha_i} - H_{-i} = \sqrt{h_{3-i}/\alpha_i} - h_{3-i} = \alpha_{3-i}/(\alpha_1 + \alpha_2)^2$ in every epoch. Indeed, we have that $\alpha_i H_{-i} = \alpha_1 \alpha_2 / (\alpha_1 + \alpha_2)^2 \leq 1/4 \leq (4\sqrt{2} - 5)^2$, for every positive numbers α_1, α_2 . \square

COROLLARY 3. Consider the case of symmetric miners, all with efficiency parameter α . For $n \geq 3$, there is no pure symmetric Nash equilibrium in which the miners use the same power in each epoch.




Fig. 2. The utility of the best periodic response with periods 1 (parabola) and 2 (straight line), when the other players use the same power H_{-i} in every epoch (Lemma 4).

PROOF. Suppose that such an equilibrium exists. Observe first that there is no symmetric equilibrium in which all players use 0 power in some epoch. By Lemma 4, we must have $\alpha H_{-i} \leq (4\sqrt{2}-5)^2$, otherwise a varying 2-periodic better response exists.

For this range, the best response of miner i is $h_i = \sqrt{H_{-i}/\alpha} - H_{-i}$. By symmetry, $H_{-i} = (n-1)h_i$, which gives $\alpha H_{-i} = (n-1)^2/n^2$. Therefore, a symmetric Nash equilibrium exists only if $(n-1)^2/n^2 \leq (4\sqrt{2}-5)^2$. This does not hold for $n \geq 3$.

For the case of $n = 2$, $(n-1)^2/n^2 \leq (4\sqrt{2}-5)^2$ and, assuming that Conjecture 1 holds, there is a symmetric Nash equilibrium with $h_i = 1/(4\alpha)$; this also follows from Corollary 2. \square

6 DISCUSSION

Proof-of-effort mining with difficulty adjustment, proposed in Satoshi Nakamoto's paper a decade ago, is currently the engine of many blockchain systems. There is the tacit assumption that rational agents are incentivized by this protocol to outfit themselves with computational resources and mine at full speed as long as this is profitable. Here we point out that this is not so: rather sophisticated strategic considerations render the situation far more complicated, unstable, and hard to predict. Even though there are rather nice constant-effort pure Nash equilibria (still quite complex and surprising in their detail), they collapse when miners strategize from epoch to epoch. In most equilibria – or deviations from such – considered here, effort by each miner is held back to a strategic relatively low level, and it often oscillates from one epoch to the next.

We believe that these are issues that must be attended to, and explored further:

- Adjusting the adjustment: would a modified difficulty adjustment process – for example, a smoothed version based on a weighted sum of the effort levels of the few past epochs – create more manageable pure equilibria? We are pursuing this direction.
- Our results (both equilibria and deviations) assume that the agents possess full information about the game – and this is of course unrealistic. However, in most of our results the decision of agent i relies mainly on H_{-i} , the total effort exerted by others, and it is reasonable to

assume that this is observable. Still, it would be of interest to study incomplete information versions of the mining game.

- Similarly, in some parts of our work we assumed that users can vary their power from epoch to epoch substantially. It would be interesting to see how maximum energy levels can affect these results. We predict that maximum energy levels may cause even more chaotic participation in some cases.
- Do the strategic realities of mining pointed out in this paper suggest a better outlook for the energy footprint of blockchains? To answer, more research is required.
- Our results suggest that the rational behavior in many mining situations may be a mixed strategy by at least some of the miners. What would such a strategy be like? How would we know whether this is happening by looking at mining data? Note that, even in this case, a miner would be able to observe some aggregate of the other miners' current random choices.
- Incidentally, *is there always a mixed equilibrium in mining games?* Since the strategy space is continuous, Nash's theorem does not pertain, and more sophisticated analytic techniques (such as tracing the limits of the sequence of the equilibria of discretized strategy spaces) are required.
- Finally, we intend to pursue the Conjecture.

7 ACKNOWLEDGMENTS

The 1st author was supported in part by the Israel Science Foundation grant number 1841/14, the 2nd author was supported in part by NSF grant CCF-1813135, the 3rd author was supported in part by ERC Advanced Grant 321171 (ALGAME), and the 4th author was supported in part by NSF grant CCF-1763970.

REFERENCES

- [1] [n. d.]. Bitcoin Difficulty historical chart. <https://bitinfocharts.com/comparison/bitcoin-difficulty.html>. Accessed: 2019-02-13.
- [2] [n. d.]. Bitcoin Energy Consumption Index. <https://digiconomist.net/bitcoin-energy-consumption>. Accessed: 2019-02-13.
- [3] [n. d.]. List and Market Cap of Proof of Work Crypto Currencies. <https://cryptoslate.com/cryptos/proof-of-work/>. Accessed: 2019-02-13.
- [4] Lear Bahack. 2013. Theoretical Bitcoin Attacks with less than Half of the Computational Power (draft). *IACR Cryptology ePrint Archive 2013* (2013), 868.
- [5] Miles Carlsten, Harry Kalodner, S Matthew Weinberg, and Arvind Narayanan. 2016. On the instability of bitcoin without the block reward. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 154–167.
- [6] Nicolas T. Courtois and Lear Bahack. 2014. On Subversive Miner Strategies and Block Withholding Attack in Bitcoin Digital Currency. *CoRR* abs/1402.1718 (2014). arXiv:1402.1718 <http://arxiv.org/abs/1402.1718>
- [7] Ittay Eyal and Emin Gün Sirer. 2018. Majority is Not Enough: Bitcoin Mining is Vulnerable. *Commun. ACM* 61, 7 (June 2018), 95–102. <https://doi.org/10.1145/3212998>
- [8] Aggelos Kiayias, Elias Koutsoupias, Maria Kyropoulou, and Yiannis Tselekounis. 2016. Blockchain Mining Games. In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016*, Vincent Conitzer, Dirk Bergemann, and Yiling Chen (Eds.). ACM, 365–382. <https://doi.org/10.1145/2940716.2940773>
- [9] Satoshi Nakamoto. 2008. Bitcoin: A peer-to-peer electronic cash system. (2008).
- [10] Arvind Narayanan, Joseph Bonneau, Edward Felten, Andrew Miller, and Steven Goldfeder. 2016. *Bitcoin and cryptocurrency technologies: a comprehensive introduction*. Princeton University Press.
- [11] K. Nayak, S. Kumar, A. Miller, and E. Shi. 2016. Stubborn Mining: Generalizing Selfish Mining and Combining with an Eclipse Attack. In *2016 IEEE European Symposium on Security and Privacy (EuroSP)*. 305–320. <https://doi.org/10.1109/EuroSP.2016.32>
- [12] Ayelet Sapirshstein, Yonatan Sompolinsky, and Aviv Zohar. 2017. Optimal Selfish Mining Strategies in Bitcoin. In *Financial Cryptography and Data Security*, Jens Grossklags and Bart Preneel (Eds.). Springer Berlin Heidelberg, Berlin,

Heidelberg, 515–532.

- [13] Yonatan Sompolinsky and Aviv Zohar. 2018. Bitcoin's underlying incentives. *Commun. ACM* 61, 3 (2018), 46–53. <https://doi.org/10.1145/3152481>
- [14] Aviv Zohar. 2015. Bitcoin: Under the Hood. *Commun. ACM* 58, 9 (Aug. 2015), 104–113. <https://doi.org/10.1145/2701411>